Analysis of Lead Time Correlation under a Base-Stock Policy[1]

Tim Hellemans

Research Center for Operations Management, University KU Leuven, Belgium
Department of Mathematics and Computer Science, University of Antwerp, Belgium

Robert N. Boute

Research Center for Operations Management, University KU Leuven, Belgium
Technology and Operations Management Area, Vlerick Business School, Belgium.

Benny Van Houdt

Department of Mathematics and Computer Science, University of Antwerp, Belgium.

Abstract

We analyze the impact of lead time correlation on the inventory distribution, assuming a periodic review base-stock policy. We present an efficient method to compute the shortfall distribution for any Markovian lead time process, and we provide structural results when lead times are characterized by a 2-state Markov-modulated process. The latter reveals how lead time correlation increases the inventory variance and enables a closed form for the asymptotic behavior of the shortfall's variance in case the two possible lead time values are sufficiently different. We also establish upper and lower bounds on the inventory variance, which hold for any general time-homogeneous lead time process. Our results are complemented by a numerical experiment that indicates how commonly used approximations of the shortfall distribution mis-specify base-stock levels in the presence of lead time correlation. Not only does the inventory distribution increase in variance as the lead time correlation increases, it also becomes multi-modal.

Keywords: Inventory, Base-stock policy, Stochastic correlated lead-times, order crossovers

[1]To appear in *European Journal of Operations Research*

1. Introduction

The base-stock policy is a well known inventory policy, shown to be optimal when there is no fixed ordering cost and both holding and shortage costs are proportional to the volume of on-hand inventory or shortage (see e.g. Scarf, 1959, Chapter 13). The optimal base-stock level is commonly determined via the demand over the lead time, which is typically approximated using a uni-modal (e.g., normal or negative binomial) distribution. This approach has been validated to stochastic lead times, as long as there are no order crossovers (see e.g. Hadley and Whitin, 1963; Kaplan, 1970).

Order crossover occurs whenever replenishment orders do not arrive the sequence in which they are placed. There are several reasons why order crossovers exist: the use of (several) global suppliers (simultaneously), adoption of multiple and more variable transportation modes, the increased frequency of orders (driven by just-in-time inventory practices), etc. (Robinson et al., 2001; Srinivasan et al., 2011; Disney et al., 2016).

If such order crossover occurs, traditional inventory analysis overstates expected shortages and thus setting the base-stock levels based on the lead-time demand distribution will overestimate base-stock levels and lead to excess inventory. Instead, one should rely on the distribution of the shortfall, defined as the outstanding inventory (or inventory on order) at the start of a period (Robinson et al., 2001). The shortfall distribution determines the inventory distribution, which has a smaller variance than the lead time demand in case of independent and identically distributed (i.i.d.) lead times (Zalkind, 1976).

In this paper we focus on the impact of correlation in stochastic lead times, in combination with order crossovers. When lead times are correlated, the lead time demand is not impacted, but the shortfall distribution and hence the inventory distribution and optimal base-stock levels are.

1.1. Background and Motivation

Lead time correlation is only rarely studied in the literature, despite its prevalence in many real life supply chains. In transportation for instance, consecutive lead times may be positively correlated: if it takes long to deliver an order (for instance due to congestion), it is likely that the subsequent order will also face a long lead time. Negative correlation may occur where multiple modes of transportation or multiple suppliers are interchangeably

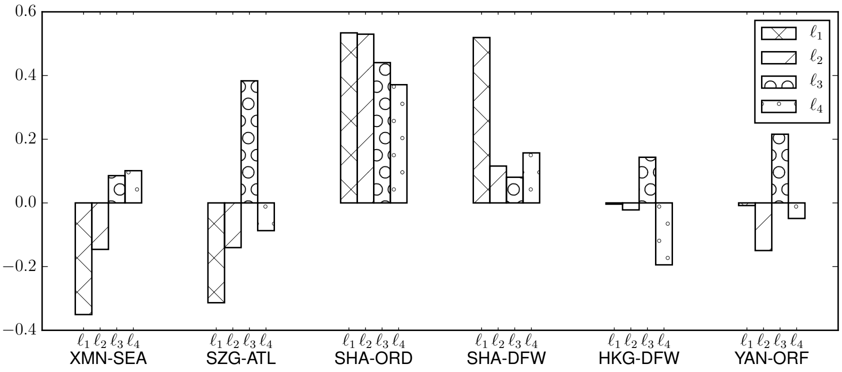

Figure 1: The lag-$1, 2, 3, 4$ correlations in consecutive lead times for six inter-port shipping routes (the abbreviations represent the port codes).

used with different lead times (e.g., a fast transportation mode or nearshore supplier versus a slow transportation mode or offshore supplier). We have analyzed logistics data of 65,536 inter-port shipments of global forwarders. These data reveal that consecutive inter-port lead times are always correlated, with correlation values that can be either positive or negative. Figure 1 reports the lag-$1, 2, 3, 4$ correlations of six of these routes (the abbreviations represent the port codes): the two routes with the most negative correlation (XMN-SEA, SZG-ATL), the two routes with the highest positive correlation (SHA-ORD, SHA-DFW) and the two routes with the least correlation among all routes (HKG-DFW, YAN-ORF). This dataset reveals that there is always correlation in the lead times, with the positive correlation generally stronger than the negative. The same dataset was considered in Disney et al. (2016) to demonstrate the existence of crossovers.

The existence of correlated lead times and order crossovers motivates us to analyze how such correlation in consecutive lead times impacts the inventory distribution (via its impact on the shortfall distribution) in a single-item inventory system with stochastic lead times and order crossovers under a simple periodic review base-stock policy. We acknowledge that in case of order crossovers, a base-stock policy is no longer optimal (Srinivasan et al., 2011). However, the base-stock policy is a well-studied policy with many familiar properties. Moreover, base-stock control policies are often studied in literature as they provide a benchmark on how much inventory is needed to provide a certain service level. In this sense it sharpens the focus on the higher-level business issue of inventory/service trade-off without getting into operational issues such as order sizes. This makes it an attractive model to analyze the effect of lead times on system performance.

Our approach exists in analyzing the shortfall distribution through the number of outstanding orders, as the latter is the only component of the shortfall distribution which depends on the correlation in lead times. We study how correlation in the lead time process impacts the shortfall distribution (and thus the optimal base-stock policy and associated costs) for a given set of lead times and a given steady state distribution of the lead time process. We know that for a given lead time distribution, the lead time demand is independent of the correlation in lead times. For general time-homogeneous lead time processes without crossovers, the shortfall distribution coincides with the lead time demand distribution and is thus also independent of the correlation in the lead time process. However, as soon as crossovers are prevalent, the shortfall distribution and its variance are strongly influenced by the correlation in consecutive lead times. To the best of our knowledge, we are not aware of any paper that provides structural insights on the impact of lead time correlation on the inventory system in the presence of order crossovers.

1.2. Related literature

A majority of the inventory literature has circumvented the difficulty due to order crossover either by simply ignoring it (e.g., stating that the probability of order crossover is small), or by constructing models so that crossovers are not possible or provides no benefit. However, recent literature has shown that order crossovers are prevalent in many real life supply chains, and ignoring order crossover has significant inventory cost implications. Even better, supply chain managers could exploit order crossover to reduce safety stocks, which means that order crossover is not necessarily baneful, but actually helpful in reducing inventory risk (Hayya et al., 2008).

Robinson et al. (2001) examine the effects of order crossover in a periodic review base-stock model. They show that optimizing base-stock levels based on the lead time demand distribution – rather than the shortfall – can lead to significantly higher inventory cost, even if the probability of order crossover is small. They present an iterative algorithm for computing the distribution of the number of orders outstanding, and formulae for the inventory shortfall distribution. As the shortfall not only has a smaller variance than the lead time demand distribution (this result was first derived by Zalkind (1976)), but also a different shape, they propose to use the negative binomial distribution to approximate the shortfall, rather than the common practice of approximating the distribution as normal. In a follow-up article, Bradley et al. (2005) develop an upper bound on the variance of the number of orders outstanding, which facilitates the computation of both the normal and negative binomial approximations. They show that the variance of the number of orders outstanding is bounded above by the standard deviation of lead time divided by $\sqrt{3}$. Wensing and Kuhn (2015) extend the focus towards periodic-review base-stock policies with arbitrary review period.

Hayya et al. (2008) introduce the concept of "effective lead time", defined as the time between the ith order placement and the ith order arrival, with the index i not tagged to a particular order. For a given lead time distribution, the occurrence of crossovers leads to an effective lead time distribution, whose mean is the same as that of the original lead time but whose variance is less than the original variance. As the effective lead times, and not the original lead times, determine the inventory distribution, supply chain managers could exploit order crossover to reduce safety stocks. Bischak et al. (2014) develop an approximation of the effective (reduced) lead time standard deviation using simulation analysis. They are one of the few that address correlated lead times (e.g., due to congestion) in their analysis. Muharremoglu and Yang (2010) provide a numerical method to determine optimal base-stock levels for a broad class of stochastic (including non-i.i.d.) lead time processes.

He et al. (1998) examine the impact of crossovers in a continuous review inventory system with constant demand and stochastic lead times and find that the optimal order quantity is lower (compared to ignoring it), but the reorder point may be more or less, depending on the parameters of the problem. Song and Zipkin (1996) investigate different approximations to the variance of the shortfall for continuous review inventory policies, which they model as being normally distributed.

Whereas the majority of the work dealing with order crossover correctly states that considering order crossover can lead to lower inventory levels and subsequently lower inventory costs, only few have investigated the optimal inventory policy for inventory systems with order crossover. Srinivasan et al.

(2011) show using dynamic programming that, in the presence of order crossover, the optimal inventory control policy is state-dependent and not only takes into account the inventory position, but also the age of pending orders and the probability distribution of their arrivals in the future. Due to the curse of dimensionality, the optimal policy is intractable for problems with large state spaces. Disney et al. (2016) find that, when there is order crossover, the so-called proportional base-stock policy, a linear generalization of the base-stock policy outperforms the traditional base-stock policy.

There is also a number of papers that discuss order crossovers due to (deterministic) changes in the lead time, e.g., due to a lead time reduction program or due to a change of supplier. This leads to expected crossovers, which can be anticipated upon. Axsäter (2011) deals with the transient inventory control to anticipate such lead time changes and Gaalman and Riezebos (2005) and Riezebos and Gaalman (2009) show that the standard base-stock policies are no longer optimal in case of expected crossovers. Also Riezebos and Zhu (2015) consider a similar problem. However, such lead-time changes are not that much related to the problem we consider.

1.3. Our Results and Positioning in the Literature

With the exception of Bischak et al. (2014) and Muharremoglu and Yang (2010), the literature studying order crossovers assumes i.i.d. lead times. We contribute to the literature by addressing the following questions:

- How does correlation in lead times impact the inventory distribution and the optimal base-stock levels?
- How do the current approximations (which assume i.i.d. lead times) perform in the presence of lead time correlation?

To the best of our knowledge, we are not aware of any paper that provides structural results on the impact of lead time correlation on the inventory distribution. That is also the reason why we restrict our analysis to a simple base-stock policy with i.i.d. demand and no correlation between demand and lead times. We acknowledge that the current literature on stochastic lead times does cope with more advanced techniques, such as correlated demand (Wang and Disney, 2017), correlation between demand and lead time (Boute et al., 2014) or stochastic lead time reduction (Hayya et al., 2011). However, none of these consider correlation in lead times. As it is not our ambition to include these generalizations, our main benchmark is Robinson et al. (2001), who study a base-stock policy with order crossovers and i.i.d. lead times.

Our main results can be summarized as follows. First, we present an efficient numerical method to compute the shortfall distribution through the distribution of the number of outstanding orders for general Markov modulated lead times. This extends the numerical method presented in Robinson et al. (2001) for i.i.d. lead times to include correlated lead times. Second, while this numerical scheme allows fast numerical evaluation, it does not establish structural results on how the correlation affects the shortfall distribution. We therefore complement this numerical method with structural results on the variance of the shortfall distribution. We prove for a 2-state Markov modulated lead time process that the variance of the shortfall distribution (and hence of the inventory distribution) grows as a function of the correlation (more specifically, the lag-1 autocorrelation) and we provide a closed form of the asymptotic behavior of the variance of the shortfall in case the two possible lead time values are sufficiently different. Third, we present for a general time-homogeneous process a tight upper bound on the variance of the shortfall distribution and a lower bound, of which the tightness depends on the steady state distribution and state space of the lead time process. This dependence can be characterized analytically for a 2-state Markov modulated lead time process, which enables a tight lower bound for a given steady state lead time distribution.

Finally, we present a numerical example which demonstrates that for strongly correlated lead time processes, the shortfall distribution tends towards the lead time demand distribution, which is multi-modal. As a consequence the (uni-modal) normal and negative binomial approximations are both no longer good estimators of the shortfall (and inventory) distribution. When we restrict to a 2-state Markov modulated lead time process we can analytically derive the convergence of the number of outstanding orders to the lead time as the correlation tends to 1, which entails the convergence in distribution of the shortfall to the lead time demand. Through numerical examples we conjecture that this convergence extends to general Markov modulated lead times. This result contributes to the existing literature, as Robinson et al. (2001) note that in the presence of i.i.d. lead times and order crossovers, the lead time demand distribution may be multi-modal, but the corresponding shortfall distribution is generally not; and as long as the inventory parameters are computed using the shortfall, the multi-modality of lead time demand does not pose a problem and uni-modal (like normal or negative binomial) approximations may still be accurate. However, we show that this is no longer valid when lead times are correlated, in which case the shortfall distribution may also be multi-modal and uni-modal approximations may be inaccurate.

The remainder of the paper is structured as follows. In the next section we provide a formulation of the problem. Section 3 is devoted to the study of the shortfall distribution under general time-homogeneous, Markov modulated and 2-state Markov modulated lead time process. In Section 4 we provide structural results to describe the variance of the shortfall, which is a prime determinant of the inventory-related costs. A numerical experiment is presented in Section 5 to illustrate our findings.

2. Problem Formulation

Throughout this paper, we use a periodic review model with infinite horizon, where we always assume to be in steady state. In a period, the following sequence of events occur: (1) previously placed orders are received in inventory; (2) the order quantity of the current period is determined; (3) instant arrival of this order in case the lead time is 0; (4) Uncertain demand is realized and satisfied from the inventory on hand; and (5) the shortfall and inventory are calculated, and backlog and holding costs are incurred based on the end-of-period inventory levels. This is schematically depicted in Figure 2 and we explain this in more detail below.

We denote the lead time process by $\{L_t\}$ where the value of L_t represents the number of periods it takes for an order placed at time $t \in \mathbb{N}$ to arrive in inventory. We assume that $\{L_t\}$ is time-homogeneous with steady state distribution π and a countable state space $\mathcal{S}$. In our analysis we mostly specify that $\{L_t\}$ follows an irreducible, time-homogeneous Markov process with transition matrix $\mathbf{L}$ (in some cases we will assume a general time-homogeneous process). This means that for every $s, s' \in \mathcal{S}$ and $t \in \mathbb{N}$ we have $\mathbb{P}\{L_{t+1} = s' \mid L_t = s\} = \mathbf{L}_{s,s'}$. Let π be its steady state distribution, i.e. a vector s.t. $\pi \cdot \mathbf{L} = \pi$, then for any $s \in \mathcal{S}$ and $t \in \mathbb{N} : \mathbb{P}\{L_t = s\} = \pi_s$. We use the elements of $\mathcal{S}$ as indexes; if for example $\mathcal{S} = \{1, 9, 27\}$, we write: $\pi = \begin{pmatrix} \pi_1 & \pi_9 & \pi_{27} \end{pmatrix}$. In this notation it is always assumed if $\mathcal{S} = \{s_1, \ldots, s_m\}$

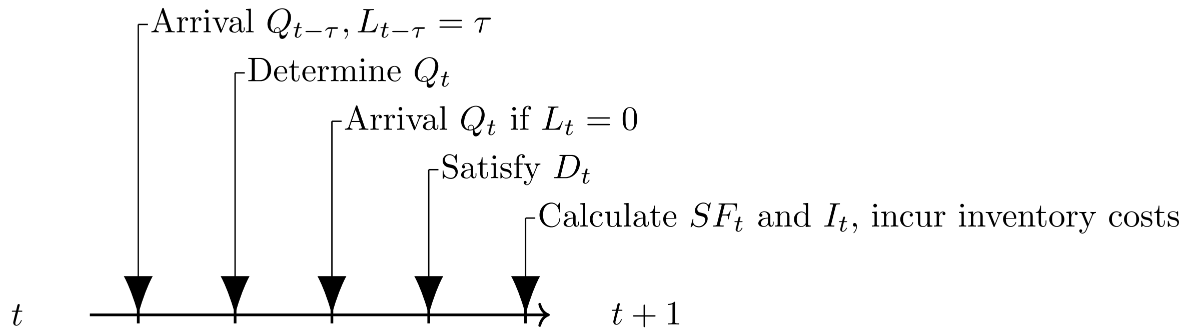

Figure 2: Schematic Representation of the series of events in one period.

and $s_1 < s_2 < \cdots < s_m$ (with possibly $m = \infty$) that the first row/column of $\mathbf{L}$ and the first element of π correspond to s_1 and so on.

We denote the demand process by $\{D_t\}$, which we assume to be an i.i.d. process independent of $\{L_t\}$. We denote Q_t the order quantity in period t. We assume a standard base-stock policy, which means that in period $t \in \mathbb{N}$ the order quantity $Q_t = D_{t-1}$. We define the shortfall SF_t as the amount by which the inventory level I_t is below the base-stock level S at the end of period t (after satisfying demand), so that $I_t = S - \mathrm{SF}_t$:

$$\mathrm{SF}_t := D_t + \sum_{l=0}^{\infty} D_{t-l-1} \cdot \delta\{L_{t-l} > l\},$$

with $\delta\{A\} = 1$ if A is true and 0 otherwise. This quantity represents the outstanding inventory still on order at the end of the period, plus the current period's demand. Under a time-homogeneous lead time process, the shortfall distribution does not depend on the time t and we simply denote the shortfall by SF (we use this notation for all variables whose distribution does not depend on t, e.g. the lead time L). We assume that a per-unit holding cost h and shortage cost b are levied against the positive (on-hand) and negative (backordered) parts of the inventory level at the end of each period. With $I = S - \mathrm{SF}$ the end-of-period inventory, we have $\mathrm{Var}(I) = \mathrm{Var}(\mathrm{SF})$ (which means that we can use the variance of the inventory level and the variance of the shortfall interchangeably), and the expected costs per period is then:

$$\mathbb{E}\left[h \cdot I^+ + b \cdot I^-\right],$$

where we denote $f^{\pm} := (0 \vee \pm f)$ for any function $f : \mathbb{R} \to \mathbb{R}$. Whereas the mean of the shortfall is not influenced by the correlation in lead times (see Lemma 2), its variance is. Moreover, the variance of the shortfall is a prime indicator of the expected costs: the higher the variance, the higher the inventory-related costs. This makes the variance of the shortfall an interesting object to study, which is the topic of Section 4.

3. Analysis of the shortfall distribution under correlated lead times

To analyze the distribution of the shortfall, we distinguish three cases. First we consider a general time-homogeneous lead time process, in which case we show that correlation in lead times only impacts the shortfall distribution if there are crossovers; in the absence of crossovers, the shortfall is not impacted by any correlation in lead times (see Section 3.1). Second, for Markov modulated lead times, we provide an efficient recursive scheme to compute the shortfall distribution (Theorem 1); this scheme can also take correlated demand into account. Note that similar schemes exist only for i.i.d. lead times (not for correlated lead times). Third, in case of a 2-state lead time process we characterize the distance between the shortfall and the lead time demand analytically (Proposition 1) and we show that the shortfall converges (in distribution) to the lead time demand as the lead time correlation tends to one (Corollary 1). Among others, this means that if the lead time demand distribution is multi-modal, the shortfall becomes multi-modal as the lead time correlation increases, and uni-modal approximations of the shortfall distribution are no longer accurate. This is different from the case where lead times are i.i.d., where the lead time demand distribution may be multi-modal, but the shortfall is generally not and uni-modal approximations of the shortfall work well (Robinson et al., 2001).

3.1. Time-Homogeneous Lead Time Process

Robinson et al. (2001) note that in the presence of order crossovers, the lead time demand, defined by

$$\mathrm{LTD}_t := \sum_{i=0}^{L_t} D_{t+i},$$

cannot be used to approximate the inventory distribution; instead the shortfall $\mathrm{SF} = S - I$ should be used, which is determined by the number of outstanding orders at time t, denoted by $V_t = \sum_{k=0}^{\infty} \delta\{L_{t-k} > k\}$. The shortfall SF for i.i.d. demand is given by:

$$\mathrm{SF} \stackrel{d}{=} \sum_{l=0}^{V} D_l,$$

where $\stackrel{d}{=}$ represents equality in distribution. If there are no crossovers, it can be seen that $L \stackrel{d}{=} V$ (Muharremoglu and Tsitsiklis, 2008, Proposition 3.9). This implies that in the absence of order crossovers, the correlation in the lead time process has no influence on the shortfall, as it does not have any effect on the lead time demand. In Appendix B we provide an example of the significance of this observation for Markov modulated lead times.

In what follows we will focus our analysis on the number of outstanding orders V, as this determines the shortfall distribution SF.

3.2. Markov Modulated Lead Time Process

When the lead time process $\{L_t\}$ is Markov modulated, we can find an $|\mathcal{S}|$ dimensional Markov transition matrix $\mathbf{L}$, s.t. $\{L_t\}$ is a Markov process with transition matrix $\mathbf{L}$ and steady state π ($\pi \cdot \mathbf{L} = \pi$). The shortfall distribution in the presence of crossovers can then be computed as follows. Define for any $s \in \mathcal{S}$ and natural numbers $k \leq n$:

$$g_{(s;k|n)} := \mathbb{P}\{L_{t-n} = s \text{ and } |\{i \in \{0, \dots, n-1\} \mid L_{t-i} > i\}| = k\},$$

which is the probability that of the last n orders, exactly k are still outstanding and the order placed n periods ago has lead time s. We can calculate these values recursively. As a base we have:

$$\forall s \in \mathcal{S} : g_{(s;0|0)} = \pi_s.$$

The recursive relation is given in the following Theorem.

Theorem 1. *For $n < \sup(\mathcal{S})$, $k \leq n+1$ and $s \in \mathcal{S}$ we have:*

$$g_{(s;k|(n+1))} = \sum_{s_n \leq n, s_n \in \mathcal{S}} \frac{\pi_s}{\pi_{s_n}} \mathbf{L}_{s,s_n} g_{(s_n;k|n)} + \sum_{s_n > n, s_n \in \mathcal{S}} \frac{\pi_s}{\pi_{s_n}} \mathbf{L}_{s,s_n} g_{(s_n;(k-1)|n)}.$$

Here we use the convention that $g_{(s;k|n)} = 0$ if $k > n$ or $k < 0$.

Proof. See Appendix C. Here we also introduce a method for dealing with non-i.i.d. demand. □

Using these values we can determine $g_{(k|n)} := \sum_{s\in\mathcal{S}} g_{(s;k|n)}$, which is the probability that of the last n orders, there are exactly k still outstanding. Let $g_k := \mathbb{P}\{V = k\}$ denote the probability of having exactly k outstanding orders. In case $\mathcal{S}$ is finite, this equals $g_{(k|m)}$ with $m = \max(S)$ and otherwise $g_k = \lim_{m\to\infty} g_{(k|m)}$.

Remark: Robinson et al. (2001) also provides a recursive scheme to compute the values of $g_{(k|m)}$ and g_k, however the recursion is much simpler for i.i.d. lead times as one no longer needs to keep track of the exact values of the lead time process; and $\mathbf{L}_{s,s_n} = \pi_{s_n}$ for all values of s, s_n, which allows a simple recursion to directly obtain the values $g_{(k|n)}$ (see Eq. (2) in Robinson et al. (2001)).

Once we have obtained the values of g_k we find the shortfall distribution:

$$\mathrm{SF} \overset{d}{=} \sum_{k=0}^{\infty} g_k D^{(k+1)},$$

where $D^{(k+1)}$ is the $k+1$-fold convolution of D, which can be easily computed. If the maximum lead time is $m = \sup(\mathcal{S}) < \infty$ the complexity of finding the distribution of the number of outstanding orders is $O(|\mathcal{S}|^2 m^2)$, which is bounded by $O(m^4)$.

3.3. Two-state Markov Modulated Lead Time Process

In addition to the above numerical procedure to determine the shortfall distribution in the presence of lead time correlation and order crossovers, we can find structural results how the lead time correlation impacts the shortfall distribution when we restrict to a 2-state Markov modulated lead time. We acknowledge that a 2-state lead time process may be an oversimplification of many practical settings; nevertheless it allows to obtain structural insights in the impact of lead time correlation, which we can then numerically validate to more general lead time processes. For this purpose we first provide an exact expression of the distance between the number of outstanding orders, V_t, and the lead time, L_t, as a function of the correlation in the lead time. We denote the lag-n correlation of the lead time process $\{L_t\}$ by:

$$\ell_n := \mathrm{Corr}(L_t, L_{t+n}),$$

and for simplicity we write ℓ for ℓ_1. For general lead time processes, the shortfall distribution depends on all the values $(\ell_n)_n$. However, for a 2-state Markov modulated lead time process, we can show that $\ell_n = \ell^n$, so that we only need to consider ℓ. Moreover, we can find a simple expression of $\mathbf{L}$ in function of this correlation ℓ, which will prove to be useful in our further analysis of the shortfall.

Lemma 1. *If the lead time process $\{L_t\}$ is given by an irreducible 2-state Markov process with state space $\{0, m\}$, steady state $(\alpha, 1-\alpha)$ and demand is i.i.d., it follows that:*

$$\ell = \mathrm{Det}(\mathbf{L}),$$

which entails for arbitrary $n \in \mathbb{N}$:

$$\ell_n = \ell^n.$$

Further, ℓ is independent of m and we can write the n'th power of the transition matrix as:

$$\mathbf{L}^n = \begin{pmatrix} (1-\alpha)\ell^n + \alpha & (1-\alpha)(1-\ell^n) \\ \alpha(1-\ell^n) & \alpha\ell^n + (1-\alpha) \end{pmatrix}, \begin{cases} \frac{\alpha}{\alpha-1} \leq \ell & \text{if } \alpha \in [0, \frac{1}{2}], \\ \frac{\alpha-1}{\alpha} \leq \ell & \text{if } \alpha \in [\frac{1}{2}, 1]. \end{cases}$$

Proof. See Appendix D. □

Remark: One can easily generalize Lemma 1 to include a state space $\{m', m\}$.

Using Lemma 1 we can express the distance of the number of outstanding orders to the lead time process, $\mathbb{E}[|V_t - L_t|]$, as a deterministic function of the state space, steady state and the correlation ℓ in case of a 2-state Markov modulated lead time process.

Proposition 1. *Let $\{L_t\}$ be a 2-state Markov chain with state space $\mathcal{S} = \{0, m\}$ for $m \in \mathbb{N}_0$ and steady state $\pi = (\alpha, 1-\alpha)$ for $\alpha \in [0, 1]$ and suppose we have i.i.d. demand. Then,*

$$\mathbb{E}[|V_t - L_t|] = \frac{2(\alpha-1)\alpha\left(\ell^m - \ell m + m - 1\right)}{\ell - 1}.$$

Proof. See Appendix E. □

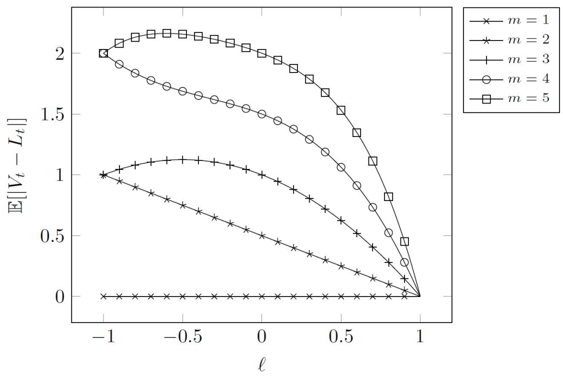

Figure 3: Convergence of $\mathbb{E}[|V_t - L_t|]$ to zero as $\ell \to 1$ for steady state $\pi = (1/2\ 1/2)$ and state space $\mathcal{S} = \{0, m\}$ for different values of m.

By applying l'Hôpital's rule, the right hand side of the above Proposition converges to zero when $\ell \to 1$. This means that the distribution of the number of outstanding orders converges to the lead time distribution as the lead time correlation approaches 1. Figure 3 illustrates the convergence speed of $\mathbb{E}[|V_t - L_t|]$ for $\alpha = 1/2, m = 1, 2, 3, 4, 5$. As m increases, the convergence is slower.

One can generalize the result in Proposition 1 to the case where lead times have a state space $\mathcal{S} = \{m', m\}$, but then one should compare $V_{t+m'}$ and L_t instead of V_t and L_t. Of course, as the distribution of V_t and L_t is independent of t, convergence in distribution of V to L still holds.

As the lead time demand and shortfall distribution are given by $\sum_{l=0}^{L} D_l$ and $\sum_{l=0}^{V} D_l$ respectively, Proposition 1 also gives an indication on the convergence of the shortfall distribution to the lead time demand distribution.

Corollary 1. *Under the same conditions as in Proposition 1 we have: the shortfall converges in distribution to the lead time demand as $\ell \to 1$, more specifically: Let* $\mathrm{SF}^{(n)}$ *denote the shortfall distribution in case the correlation is $\ell^{(n)}$, then for any sequence $\{\ell^{(n)}\}$ in $[0,1]$ s.t. $\ell^{(n)} \to 1$, we have* $\mathrm{SF}^{(n)} \to_d$ LTD *with $\to_d$ convergence in distribution.*

Proof. See Appendix F. □

Based on numerical analysis, we conjecture that this convergence remains valid for general Markov modulated lead times (see Section 5).

4. Properties of the Variance of the Shortfall under Correlated Lead Times

We now provide structural results how lead time correlation impacts the most important property of the shortfall distribution, namely its variance.

4.1. Time-Homogeneous Lead Time Process

Further on, whenever convenient, we denote for a random variable X its variance by σ_X^2 and its mean by μ_X. Denote for a set A and a number $x : I_A(x)$ the indicator function which equals one if $x \in A$ and zero otherwise. For a general time-homogeneous lead time process, we find that the mean of the shortfall does not depend on the correlation in the lead time process:

Lemma 2. *For a time-homogeneous demand and lead time process s.t.* $\mu_L < \infty$ *and* $\sum_{n=1}^{\infty} \mathbb{E}[D \cdot I_{\{V \geq n\}}] < \infty$*, we have* $\mu_{\text{SF}} = \mu_D \cdot (\mu_L + 1)$.

Proof. This result follows by applying Wald's identity and the fact that $\mu_V = \mu_L$ (which was first shown in Zalkind (1978)). □

Remark: This expression for μ_{SF} also appears in e.g. Disney et al. (2016) and Robinson et al. (2001) for the case of i.i.d. lead times. We see that it still holds for correlated lead times.

We now inspect how the variance of the shortfall behaves with respect to the correlation in the lead time.

Lemma 3. *For a time-homogeneous lead time process* $\{L_t\}$ *and i.i.d. demand process* $\{D_t\}$*, the variance of the shortfall distribution is given by:*

$$\sigma_{\text{SF}}^2 = (\mu_L + 1)\sigma_D^2 + \sigma_V^2 \mu_D^2.$$

Proof. This is a straightforward computation which can be found in Appendix G for the sake of completeness. □

Remark: This expression for σ_{SF}^2 also appears in Disney et al. (2016) for the case of i.i.d. lead times. We see that this result generalizes to correlated lead times.

Lemma 3 shows that the only part of σ_{SF}^2 that is influenced by the correlation in the lead time process, is the variance of the number of outstanding orders, σ_V^2. For Markov modulated lead times we can make use of g_k (obtained in Theorem 1) to determine the value of σ_V^2. When we restrict to a 2-dimensional Markov modulated lead time, we can express σ_V^2 analytically.

4.2. Markov Modulated Lead Time Process

By making use of g_k, we can numerically find the value of σ_V^2 in case of Markov modulated lead times:

$$\sigma_V^2 = \sum_{k=0}^{m} g_k k^2 - \mu_L^2,$$

with $m := \sup(\mathcal{S})$. In case $\mathcal{S}$ is unbounded, this is an infinite series.

4.3. Two-state Markov Modulated Lead Time Process

In this section we show some properties in case $\{L_t\}$ can only take two values, say $m, m' \in \mathbb{N}$. The main result from this section is Theorem 2, where we show that the variance of the shortfall increases in function of the correlation in the lead time. Further we provide explicit lower and upper bounds on the variance of the shortfall (Proposition 3) and its asymptotic value as $|m - m'| \to \infty$ (Proposition 4).

Assume without loss of generality that $m' < m$ and let $\overline{m} := m - m'$ (note that with only two possible lead times we can have crossovers except if $m' = m - 1$). The shortfall is given by:

$$\mathrm{SF}_t = D_t + \sum_{k=0}^{m'-1} D_{t-k-1} + \sum_{k=0}^{\overline{m}-1} \delta\{L_{t+m'+k} = m\} D_{t-m'-k-1}.$$

Indeed, at any time t, we know that the orders placed up to $m' - 1$ periods ago (reflected in the first term of SF_t), are yet to be delivered in inventory. The orders placed at least m' but no more than m periods ago (the second term of SF_t) are not yet delivered if and only if that order has lead time m. Define the process $\overline{L_t} := L_t - m'$, the above yields $\sigma_{\mathrm{SF}}^2 = m'\sigma_D^2 + \sigma_{\overline{\mathrm{SF}}}^2$, with $\overline{\mathrm{SF}}$ the shortfall distribution for the lead time process $\{\overline{L_t}\}$. As the part $m'\sigma_D^2$ does not depend on the correlation in the lead time process, we focus on $\sigma_{\overline{\mathrm{SF}}}^2$; for this process the lead times are given by 0 and $\overline{m}$. For ease of notation, we let $\{L_t\}$ be the lead time process with state space $\mathcal{S} := \{0, m\}$ and keep in mind that if we were to replace the zero lead time by a non-zero lead time, we have to add the part $m'\sigma_D^2$.

As discussed in Section 3.3, for 2-state Markov modulated lead times, the variance of the shortfall distribution only depends on ℓ as $\ell_n = \ell^n$. This enables us to find a non-decreasing function f s.t. if we fix m and the steady state $\pi = (\alpha, 1-\alpha)$, we have $\sigma^2_{\mathrm{SF}} = c + f(\ell, m)$, where c is independent of ℓ. This implies that the variance of the shortfall is a deterministic function of the correlation and the value of m; more specifically it is non-decreasing with respect to the correlation.

Proposition 2. *For a 2-state Markov modulated lead time process with lead times 0 and $m \in \mathbb{N}$ and i.i.d. demand process, we have:*

$$\sigma^2_V = \frac{\sigma^2_L}{m} + \frac{2\sigma^2_L}{m^2} \sum_{k=1}^{m-1} (m-k)\ell^k.$$

This entails:

$$\sigma^2_{\mathrm{SF}} = \sigma^2_{i.i.d.} + 2\left(\frac{\mu_D \sigma_L}{m}\right)^2 \sum_{k=1}^{m-1} (m-k)\ell^k,$$

with $\sigma^2_{i.i.d.} = \sigma^2_{const} + \frac{\mu_D^2 \sigma_L^2}{m}$ and $\sigma^2_{const} = (\mu_L + 1)\sigma^2_D$.

Proof. See Appendix H. □

Proposition 2 shows that σ^2_{SF} consists of three parts. The first two parts are known results in standard inventory theory: $(\mu_L + 1)\sigma^2_D$ represents the variance of the shortfall that is due to the *mean* number of orders that have not yet arrived and thus corresponds to the case of constant lead times; The second part $\frac{\mu_D^2 \sigma_L^2}{m}$ is due to introducing variability into the steady state of the lead time process, this part is always positive and corresponds to the case where we have an i.i.d. lead time process. The last part, $2\left(\frac{\mu_D \sigma_L}{m}\right)^2 \sum_{k=1}^{m-1} (m-k)\ell^k$, is new and is due to the correlation in the lead time process. It is however important to note that this part need not be positive; indeed, for negative values of ℓ, this part may be negative.

We now explicitly show that in case of a 2-state Markov modulated lead time process, this last term is non-decreasing with respect to the lag-1 correlation (and zero when the correlation is zero). We then extend these results to the general Markov modulated process numerically.

Define the function

$$f(x,n) := \frac{1}{n^2}\sum_{k=1}^{n-1}(n-k)x^k, n \in \mathbb{N}, x \in [-1,1].$$

Then according to Lemma 1, we find a closed-form result for the variance of the shortfall in function of the lag-1 correlation:

$$\sigma_{\mathrm{SF}}^2 = \sigma_{i.i.d.}^2 + 2\mu_D^2\sigma_L^2 f(\ell, m).$$

By analyzing f, we can show that σ_{SF}^2 increases in function of the correlation in the lead time process, which leads to Theorem 2.

Lemma 4. *The function $f(x,n)$ defined above is non-decreasing in function of x on $[-1,1]\times\mathbb{N}$, i.e. $\forall x \in [-1,1], \forall n \in \mathbb{N} : \frac{\partial f(x,n)}{\partial x} \geq 0$.*

Proof. See Appendix I. □

Theorem 2. *For a 2-state Markov modulated lead time process and i.i.d. demand, we have: $\frac{\partial \sigma_{\mathrm{SF}}^2}{\partial \ell} \geq 0$.*

Proof. This is a trivial consequence of the above discussion and Lemma 4. □

Remark: Even though we have shown that σ_{SF}^2 is non-decreasing with respect to the lag-1 correlation, it is clear from the definition of f that the slope of σ_{SF}^2 is much steeper on $[0,1]$ than on $[-1,0]$. Indeed, for $\ell \in [-1,0]$ all odd powers of ℓ are increasing while the even powers of ℓ are decreasing, whilst for $\ell \in [0,1]$ they are all increasing.

Theorem 2 implies that stronger correlation in lead times leads to more volatile inventory. This result also allows us to get a lower and upper bound for the shortfall variance, as they coincide to the cases with respectively minimal and maximal lead time correlation:

Proposition 3. *Assume we have a 2-state Markov modulated lead time process and i.i.d. demand. The variance of the shortfall is bounded from above by:*

$$\sigma_{\mathrm{SF}}^2 \leq \sigma_{i.i.d.}^2 + \mu_D^2\sigma_L^2 = \sigma_{\mathrm{LTD}}^2,$$

and this bound is tight. A lower bound (independent of the steady state π) is given by:

$$\sigma^2_{i.i.d.} + 2\mu_D^2\sigma_L^2 \cdot \frac{1}{4m^2}(-2m + (-1)^{m+1} + 1) \leq \sigma^2_{\mathrm{SF}},$$

and this bound is attained for even values of m and $\alpha = 1/2$. For even m and arbitrary α, this lower bound reduces to σ^2_{const}. For arbitrary values of m and α, the following lower bounds are tight:

$$\begin{cases} \sigma^2_{lower,1} := \sigma^2_{i.i.d.} - 2\mu_D^2\sigma_L^2 \cdot \frac{(m-\mu_L)\left(\mu_L\left(\left(1-\frac{m}{\mu_L}\right)^m - 1\right) + m^2\right)}{m^4} \leq \sigma^2_{\mathrm{SF}} & \textit{if } \alpha \leq \frac{1}{2}, \\ \sigma^2_{lower,2} := \sigma^2_{i.i.d.} - 2\mu_D^2\sigma_L^2 \cdot \frac{\mu_L\left(m\left(\left(\frac{\mu_L}{\mu_L-m}\right)^m + m - 1\right) - \mu_L\left(\left(\frac{\mu_L}{\mu_L-m}\right)^m - 1\right)\right)}{m^4} \leq \sigma^2_{\mathrm{SF}} & \textit{if } \alpha \geq \frac{1}{2}. \end{cases}$$

Proof. See Appendix J. □

Remark: These lower and upper bounds are tight for any steady state of the lead time. We note that the tight lower bound depends on this steady state, whilst the upper bound does not. This is a consequence of the fact that any steady state may correspond to a perfectly correlated process, whilst the perfectly negative correlated process has a unique steady state, namely the uniform distribution. For a general discussion on these bounds, see the proof of Proposition 3.

We use these expressions of σ^2_{SF} to analyze $\sigma^2_{\mathrm{SF}}/\sigma^2_{const}$ for $m \to \infty$, which gives an idea of the magnitude of the error made by assuming constant lead times whilst you actually have i.i.d. or correlated lead times.

Proposition 4. *Assume we have a 2-state Markov modulated lead time process and i.i.d. demand. Denote $\theta := \frac{\alpha\mu_D^2}{\sigma_D^2}$, the following convergence holds (here we take the limit for $m \to \infty$):*

- $\sigma^2_{i.i.d.}/\sigma^2_{const} \to 1 + \theta$
- $\sigma^2_{\mathrm{LTD}}/\sigma^2_{const} \to \infty$
- $\sigma^2_{lower,1}/\sigma^2_{const} \to 1 + (1 - 2\alpha)\theta$
- $\sigma^2_{lower,2}/\sigma^2_{const} \to 1 + (1 - 2(1 - \alpha))\theta$
- *In particular for $\alpha = 1/2$ we get $\sigma^2_{lower,1}/\sigma^2_{const} = \sigma^2_{lower,2}/\sigma^2_{const} \to 1$.*

Proof. See the Appendix K. □

Remark: The interpretation of Proposition 4 is as follows. Suppose we have a positively correlated lead time process with 2 sufficiently distinct lead times. If we estimate the variance of the shortfall (and with this, the base-stock level) assuming i.i.d. lead times, then we make an error which is no larger than:

$$\sigma^2_{\mathrm{LTD}}/\sigma^2_{i.i.d.} = \sigma^2_{\mathrm{LTD}}/\sigma^2_{const} \cdot \sigma^2_{const}/\sigma^2_{i.i.d.}.$$

Using Proposition 4 we find that this can become arbitrarily large! Hence, ignoring the lead time correlation potentially under-estimates the base-stock level. This misspecification increases with m and is unbounded.

We generalize these results in the more dimensional framework, but here we no longer have $\ell_n = \ell^n$ (even for Markov modulated lead times) and we have to work with the different lag-n correlations.

4.4. Extension to Time-Homogeneous Lead Time Processes

Lemma 3 shows that the only part of σ^2_{SF} that is influenced by the lead time correlation is σ^2_V, which equals:

$$\sigma^2_V = \sum_{k=0}^{\infty} \mathrm{Var}(\delta\{L > k\}) + 2 \sum_{0 \leq k < l} \mathrm{Cov}(\delta\{L_{t-l} > l\}, \delta\{L_{t-k} > k\}). \tag{1}$$

When we move from a 2-state time-homogeneous lead time process to a general time-homogeneous lead time process, a tractable analysis of the variance of the shortfall distribution is unfortunately elusive as:

- When L_t can only take 2 values, the terms $\delta\{L_t > k\}$ in Eq. (1), which equal 0 if $L_t \leq k$ and 1 otherwise, reduce to L_t/m for $0 < k < m$. However, when moving to a more-dimensional lead time process, such a simplification is no longer possible.

- Even if we found a way to write $\mathrm{Cov}(\delta\{L_{t-l} > l\}, \delta\{L_{t-k} > k\})$ as a function of $\{L_{t-k}\}$, in general it is no longer possible to find functions $\{f_m\}$ s.t. $\ell_m = f_m(\ell)$. Making these functions dependent on the steady state of the Markov process does not solve this problem (as illustrated by the example in Appendix L). This makes it highly unlikely that Eq. (1) can be written in function of the lag-1 correlation as we did in Proposition 2.

Nevertheless, we can show that the same lower and upper bound on the variance of the shortfall distribution (see Proposition 4) still hold when we extend from a 2-state to a general time-homogeneous lead time process. We also conjecture that more correlation in the lead time process implies a higher variance in the shortfall and we provide the intuition behind this conjecture. Particularly interesting is the fact that as correlation in the lead time increases, the shortfall distribution tends towards the lead time demand. Whereas we could prove this for a 2-state Markov modulated lead time in Proposition 1, we now conjecture this finding for a general time-homogeneous process.

In Eq. (1), only the last term is influenced by the lead time correlation:

$$\begin{aligned}\mathrm{Cov}(\delta\{L_{t-k} > k\}, \delta\{L_{t-l} > l\}) &= \mathbb{E}[\delta\{L_{t-k} > k\}\delta\{L_{t-l} > l\}] - \mathbb{P}\{L > k\}\mathbb{P}\{L > l\} \\ &= \mathbb{P}\{L > l\} \cdot (\mathbb{P}\{L_{l-k} > k \mid L_0 > l\} - \mathbb{P}\{L > k\}).\end{aligned} \tag{2}$$

The influence of the correlation in the lead times is captured in the term $\mathbb{P}\{L_{l-k} > k \mid L_0 > l\} - \mathbb{P}\{L > k\}$ which is the difference between on the one hand the probability of having an order with lead time longer than k, knowing that the order placed $l - k$ periods ago, has a lead time of at least l periods, and on the other hand the probability of having a lead time of at least k with no prior information. It is likely to expect that if the correlation in cumulative lead times increases, $\mathbb{P}\{L_{l-k} > k \mid L_0 > l\}$ increases while $\mathbb{P}\{L > l\}$ is not influenced. Moreover for positive correlation one would expect that generally $\mathbb{P}\{L_{l-k} > k \mid L_0 > l\} - \mathbb{P}\{L > k\} > 0$.

Also, similar to the 2-state lead time case, we find that in the more-dimensional case, we have a steeper ascend in the variance of the shortfall for $\ell \in [0, 1]$. This was implicitly proven for the 2-dimensional lead time process in Lemma 4, where we derive an exact expression for the variance of the shortfall in function of ℓ. Here, we see that as long as we have negative correlation, the probability $\mathbb{P}\{L_{l-k} > k \mid L_0 > l\}$ might behave unexpected, depending on the difference between l and k. Once the correlation is positive, these probabilities all increase in function of the correlation.

Unfortunately, a direct generalization of Theorem 2, which would be to state that $\frac{\partial \sigma^2_{\mathrm{SF}}}{\partial \ell} \geq 0$, is not possible (see Appendix L). However, from the intuition obtained above we can conjecture the following generalization in the more dimensional case:

Conjecture 1. *Suppose we have two time-homogeneous lead time processes $\{L_t\}$ and $\{L'_t\}$ with the same steady state s.t. for all $k \in \mathbb{N}$, we have:*

$$\mathrm{Cov}(L_0, L_k) \leq \mathrm{Cov}(L'_0, L'_k),$$

and i.i.d. demand. Then the variance of the shortfall distribution σ^2_{SF} corresponding to the lead time process $\{L_t\}$ is smaller than that corresponding to the the lead time process $\{L'_t\}$.

Although we cannot prove Conjecture 1, we at least have the following (weak) non-decreasing result:

Proposition 5. *Suppose we have i.i.d. demand and two time-homogeneous lead time processes $\{L_t\}$ and $\{L'_t\}$ with the same steady state s.t. we have for all $k < l \in \mathbb{N}$:*

$$\mathbb{P}\{L_{l-k} > k \mid L_0 > l\} \leq \mathbb{P}\{L'_{l-k} > k \mid L'_0 > l\},$$

then we have that the σ^2_{SF} corresponding to the process $\{L_t\}$ is smaller than that of the process $\{L'_t\}$.

Proof. This is immediately clear from Eqs. (1-2). □

We now establish bounds on σ^2_{SF} for general time-homogeneous lead time processes.

Theorem 3. *For a time-homogeneous lead time process $\{L_t\}$, we have: $0 \leq \sigma^2_V \leq \sigma^2_L$. In particular, this implies for i.i.d demand $\{D_t\}$, that σ^2_{SF} is bounded by:*

$$\sigma^2_{const} = (\mu_L + 1)\sigma^2_D \leq \sigma^2_{\mathrm{SF}} \leq (\mu_L + 1)\sigma^2_D + \sigma^2_L \mu^2_L = \sigma^2_{\mathrm{LTD}}.$$

Proof. See Appendix M. □

Remark: The upper bound is tight as it is attained by taking the limit $\ell \to 1$ for appropriately chosen transition matrices. In the two dimensional case we noticed that for Markov modulated lead times it was possible to obtain $\sigma^2_V = 0$ for even values of m, by taking perfectly negative correlated lead times (see the Proof of Proposition 3). We may wonder whether this is still possible in the multi-dimensional case by posing this as an algebraic question. We therefore numerically checked with state space $\{0, \ldots, m\}$ and π uniform for $m = 2, \ldots, 10$ that $\sigma_V^2 = 0$ is always attainable. Without going further into detail, we conjecture the answer is yes: for an appropriate state space and uniform steady state, it is always possible to attain $\sigma_V^2 = 0$ for an arbitrary number of possible distinct lead times. This means that for a uniformly distributed steady state and $\mathcal{S} = \{n, n+1, \ldots, n+k\}$ for some $n, k \in \mathbb{N}$, the lower bound found by setting $\sigma_V^2 = 0$ is always attained, however for other π and $\mathcal{S}$, one could find a lower bound which is more tight (as we did for the 2-dimensional case in Proposition 3).

This upper bound corresponds to the one found in Robinson et al. (2001) for i.i.d. lead times, which corresponds to the case without crossovers (where the shortfall equals the lead time demand). Note however, that the bounds provided by Bradley et al. (2005) for i.i.d. lead times:

- $\sigma_V^2 \leq \min\left\{\frac{\sigma_L}{\sqrt{3}}, \mu_L\right\}$;
- if $L \in [\mu_L - k\sigma_L, \mu_L + k\sigma_L]$ then $\left(\frac{k}{k^2+1}\right)\sigma_L \leq \sigma_V^2$,

no longer hold in the presence of lead time correlation, as σ_V^2 tends towards σ_L^2 when the correlation increases (see Proposition 1 for the 2-state lead time process). The lower bound does not generalize either as the variance of the shortfall may be zero for an appropriate correlation structure.

Similar to the 2-state lead time case, the upper bound for the variance of the shortfall corresponds to the setting with no order crossovers, in which case the shortfall is just the lead time demand $\text{LTD} = \sum_{k=0}^{L} D_k$, which is independent of the correlation in the lead time process.

5. Numerical Examples

In this section we numerically illustrate how the presence of lead time correlation impacts the variance of the shortfall and the optimal base-stock levels. We first consider the lead time data of the inter-port shipping route SHA-ORD, visualized in Figure 1. We then further investigate the impact of correlation in lead times for a wider range of correlation values using generated data in Section 5.2.

5.1. Real Data

To assess the impact of lead correlation using the lead time data of the inter-port shipping route SHA-ORD in Figure 1, we discretize the lead times to the state space $\mathcal{S} = \{10, 15, 20, 25, 30, 35\}$ (we do this by replacing each lead time with the value closest to it in $\mathcal{S}$). Let $A_{i,j}$ denote the number of occurrences of the transition from a lead time equal to i to a subsequent lead time which equals j; the associated transition matrix $\mathbf{L}$ is then obtained by $\mathbf{L}_{i,j} = A_{i,j} / \sum_j A_{i,j}$.

Figure 4 (left panel) shows how the distribution of the number of outstanding orders g_k under the correlated lead time process compares to its i.i.d equivalent (by assuming $\{L_t\}$ is i.i.d. with distribution π). Although the difference may not seem significant, we find that the variance of the number of outstanding orders under i.i.d. lead times is $\sigma_V^2 = 2.5$, whereas we have $\sigma_V^2 = 4.7$ when the lead time correlation is taken into account.

When we optimize the base-stock level by setting:

$$S^* := \operatorname{argmin}\{S \in \mathbb{N} \mid \mathbb{P}\{\mathrm{SF} \leq S\} \leq b/(b+h)\},$$

and assume an i.i.d. Poisson(10) demand and service level $b/(b+h) = 0.95$, we find that $S^* = 191$ and the corresponding safety stock is 46 for this (correlated) lead time process. In contrast, assuming i.i.d. lead times, we would set $S^* = 179$ or a safety stock of 34. Clearly, ignoring the correlation in lead times would set the base-stock and safety stock levels too low. Figure 4 (right panel) illustrates how this gap grows as demand increases.

We also performed some sensitivity analysis on the service levels. Figure 4 (bottom panel) shows that for high service levels, the safety stocks are higher under correlated lead times, but it is the opposite for low service levels. In the latter case, positive lead time correlation leads to lower base-stock and safety stock levels, despite its higher inventory variance. This is due to the shape of the distribution of the number of outstanding orders.

5.2. Generated Data

We assume an i.i.d. demand $\{D_t\}$ which is Poisson(10) distributed and a lead time process $\{L_t\}$ that follows a $5-$state Markov process with state space $\mathcal{S} = \{0, 7, 8, 9, 10\}$ and steady state $\pi = \left(\frac{1}{5}\ \frac{1}{5}\ \frac{1}{5}\ \frac{1}{5}\ \frac{1}{5}\right)$. We have run many numerical experiments and decided to include this setting because it allows a broad range of correlation values (incl. negative correlation as π

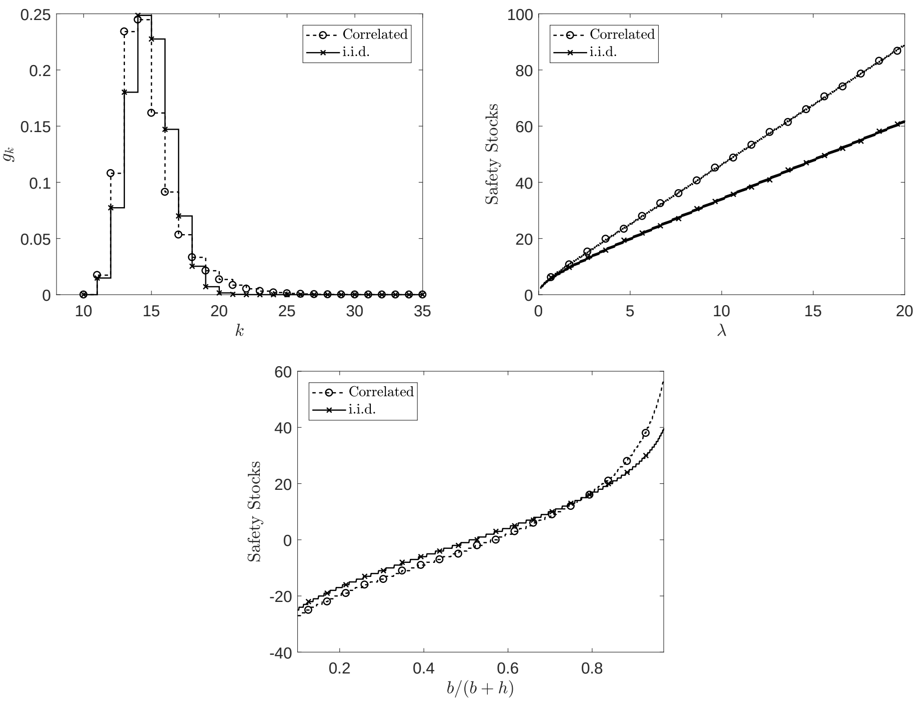

Figure 4: Top left: Distribution of the number of outstanding orders under correlated versus i.i.d. lead times. Top right: Comparison of safety stocks in function of λ, for $b/(b+h) = 0.95$ and an i.i.d. Poisson(λ) demand process. Bottom: comparison of safety stocks in function of service level (characterized by $b/(b+h)$) for i.i.d. Poisson(10) demand.

is uniformly distributed). Also, as we have essentially two ranges of lead times (a short lead time $L = 0$ with probability $1/5$ and a long lead time $L \in \{7, 8, 9, 10\}$ with probability $4/5$), the lead time demand is bi-modal. The latter is interesting as it demonstrates how the shortfall (and thus the inventory) distribution becomes bi-modal when we alter the lead time correlation.

We introduce a control parameter φ to control the correlation in lead times for a given uniform steady state π, such that as φ goes from -1 to 1, the lag-1 correlation ℓ of the associated process increases, with positive φ referring to positive correlation and negative φ indicating negative correlation. To this end, let for $a, b \in \mathbb{R} : \gamma_{a,b}(\varphi) := a(1-\varphi) + b\varphi$ be the linear function connecting a and b. Using these we define the transition matrix

$\mathbf{L}^{(\varphi)}$ of the Markov modulated lead time by letting:

$$\mathbf{L}_{i,j}^{(\varphi)} := \begin{cases} \gamma_{\pi_j,1}(-\varphi) & \text{if } \varphi \in [-1,0], i = |\mathcal{S}| - j + 1, \\ \gamma_{\pi_j,0}(-\varphi) & \text{if } \varphi \in [-1,0], i \neq |\mathcal{S}| - j + 1, \\ \gamma_{\pi_j,1}(\varphi) & \text{if } \varphi \in [0,1], i = j, \\ \gamma_{\pi_j,0}(\varphi) & \text{if } \varphi \in [0,1], i \neq j. \end{cases}$$

For any φ, the matrix $\mathbf{L}^{(\varphi)}$ is a transition matrix with steady state π. With this representation, $\mathbf{L}^{(-1)}$ is the anti-diagonal matrix, $\mathbf{L}^{(0)}$ is the transition matrix corresponding to the i.i.d. case and $\mathbf{L}^{(1)}$ corresponds to the perfect correlation case. As the steady state is uniformly distributed, the anti-diagonal matrix is a possible transition matrix, yet $\ell = -1$ is not achievable.

Figure 5 (left panel) illustrates how the correlation (we plot the lag-$1, 2, 3$ correlations) changes in function of the control variable φ for our specific example. For positive φ, the different lag correlations all increase in function of φ. However, on $[-1, 0]$ only the odd lag correlations increase. We see that we still have $\ell_n \approx \ell^n$, suggesting that the variance of the shortfall will have a steeper slope on $[0, 1]$ than on $[-1, 0]$. These lag correlations in turn control the variance of the shortfall distribution.

The right panel in Figure 5 shows how the value of φ impacts the variance of the shortfall. The variance of the shortfall *decreases* for φ going from -1 to 0, which is not in line with the results for the 2-state lead time case (which revealed that the variance of the shortfall increases in the lead time correlation). However, if we were to identify the states $7, 8, 9, 10$ as one state, the steady state becomes $[1/5\ 4/5]$, and in the 2-state lead time case the variance of the shortfall increases in the lag-1 correlation. Apparently, the fact that we have more lead time states influences this effect. The key takeaway is that for negative correlation the behavior of σ^2_{SF} is less predictable but also less significant than for positive φ. Indeed, through many numerical experiments and as well for the analytical results in the 2-state lead time process, we found that the increase in variance of the shortfall on $[0, 1]$ is substantially larger than the change on $[a, 0]$ where a represents the minimal attainable correlation. As positive correlation is also the most natural type of correlation in the lead time process, that region is therefore of most interest.

Assuming a backlog cost $b = 20$ and holding cost $h = 2$, we now evaluate

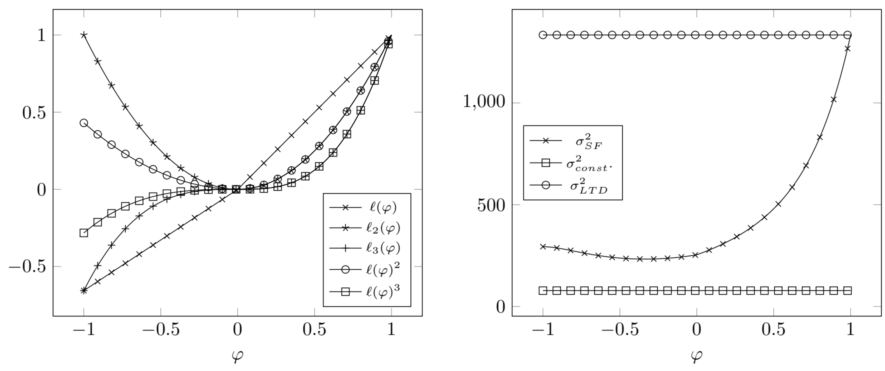

Figure 5: Left panel: the lag $1, 2, 3$ correlation associated to $\mathbf{L}^{(\varphi)}$ with state space $\mathcal{S} = \{0, 7, 8, 9, 10\}$ and $\pi = \left(\frac{1}{5}\ \frac{1}{5}\ \frac{1}{5}\ \frac{1}{5}\ \frac{1}{5}\right))$ in function of the parameter $\varphi \in [-1, 1]$. Right panel: the corresponding variance of the shortfall distribution for Poisson(10) distributed demand together with the theoretical upper/lower bounds obtained in Theorem 3, respectively given respectively given by σ^2_{LTD} and σ^2_{const}.

the optimal base-stock level and its corresponding inventory related costs in function of the parameter φ. The distribution of the shortfall can be determined in our example as follows: we first determine the distribution of the number of outstanding orders V by means of recursion using Theorem 1. The shortfall is then Poisson$((k + 1) \cdot 10)$ distributed with probability $\mathbb{P}\{V = k\}$ (as the sum of independent Poisson processes with parameters λ_1 and λ_2 is again Poisson with parameter $\lambda_1 + \lambda_2$).

We benchmark our exact procedure with several existing approximations that rely on the approximation of the shortfall distribution to determine the optimal base-stock level. The first two approximations take lead time correlation into account. The subsequent four approximations "ignore" the lead time correlation.

1. Fit the shortfall distribution with a normal distribution with mean μ_{SF} and variance σ^2_{SF} for each value of φ. We denote this the "normal" approximation.

2. Instead of the normal distribution, we fit the shortfall distribution to a negative binomial distribution, referred to as the "NB"approximation. This was suggested by Robinson et al. (2001).

3. We calculate the variance of the shortfall $\sigma^2_{i.i.d.}$ independent of the correlation in the lead time process and approximate the shortfall to a

normal random variable with mean μ_{SF} and variance $\sigma^2_{i.i.d.}$. We refer to this strategy as "i.i.d.".

4. We assume lead times are deterministic and approximate the shortfall by a normal random variable with variance σ^2_{const} and mean μ_{SF}. We refer to this strategy as "Const". This corresponds to the lower bound on the base-stock level.

5. We assume there are no crossovers and approximate the lead time demand to a normal variable with mean μ_{SF} and variance σ^2_{LTD}. We refer to this method as "LTD_n".

6. We assume there are no crossovers and use the exact lead time demand distribution, which is independent of the correlation, to approximate the shortfall. We refer to this strategy as "LTD_{ex}". This corresponds to the upper bound on the base-stock level.

Figure 6 summarizes the results of the base-stock calculations using these methods. It shows how the optimal base-stock levels generally increase in function of φ on $[0, 1]$, which is due to the fact that the variance of the shortfall distribution increases in function of φ. We see that the base-stock level increases dramatically as φ, and thus the lead time correlation, gets closer to one. Clearly, assuming i.i.d. lead times is appropriate only when the correlation is negligible ($\varphi \approx 0$). Assuming constant lead times always leads to a severe underestimation of the optimal base-stock level. We also observe that ignoring crossovers and using the exact lead time demand only works well for $\varphi \approx 1$. Finally, fitting a normal distribution to the lead time demand seems to always result in significant overestimations of the base-stock level. In other words, ignoring the correlation in lead times never leads to a good estimate for the optimal base-stock policy.

We also deduce that, whilst for small values of φ the negative binomial and normal distribution seem to be approximating the shortfall distribution fairly well, the same can not be said for large values of φ. Even though the correlation in lead times is taken into account, we find that when using the negative binomial and normal approximations, the optimal base-stock levels are overestimated. That is due to the fact that the increase in optimal base-stock levels when going from zero to perfect correlation is lower than would be expected from the increase in variance of the shortfall distribution. This

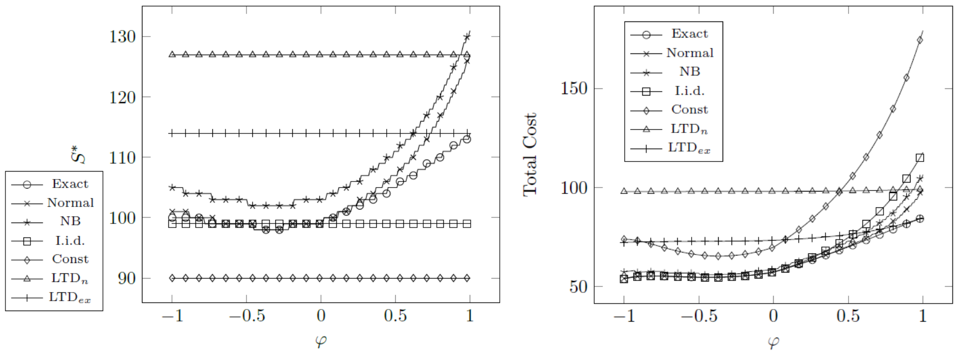

Figure 6: Left panel: The base-stock level for the different approximations of the shortfall distribution in function of φ with $b = 20, h = 2$. Right panel: The total costs of the corresponding base-stock policy with backlog cost $b = 20$ and holding cost $h = 2$.

is a consequence of the fact that the lead time demand is multi-modal and for strongly correlated lead times we have that the shortfall tends towards the lead time demand (in the two dimensional case we showed this in Proposition 1, and for $\ell = 1$ it holds that $\mathrm{SF} \overset{d}{=} \mathrm{LTD}$ in the general case). This effect occurs when the lead time values sufficiently differ from each other and there is strong correlation between subsequent lead times.

The right panel of Figure 6 shows the average inventory related costs,

$$\mathbb{E}[h \cdot (S^* - \mathrm{SF})^+ + b \cdot (S^* - \mathrm{SF})^-],$$

corresponding to these optimized base-stock levels S^* under the different approximations, when imposing them to our setting with correlated lead times. In general, we find that all approximations which do not take the correlation in lead times into account, lead to significant higher inventory costs. Clearly, the higher the misspecification of the base-stock level, the higher the total costs. The normal and negative binomial approximations, which both take the variance of the shortfall into account, perform best among all considered approximations. For high correlation however, also these approximations yield much higher total costs. That is due to the fact that these approximations overestimate the optimal base-stock level, as discussed before.

When we focus on the relative cost increase by using the normal/negative binomial distribution to fit the shortfall distribution (see Figure 7), we see that the cost increase is substantial for higher values of φ (up to a 20% cost

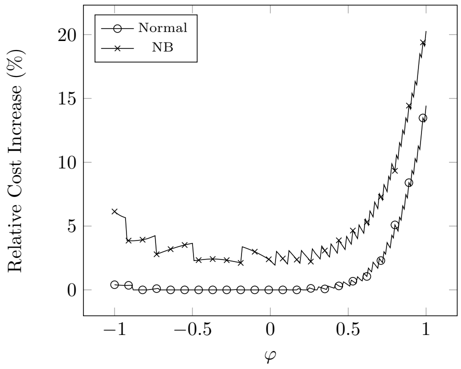

Figure 7: Relative cost increase by using a negative binomial or normal fit for the shortfall distribution.

increase in our example). This effect is reinforced for larger values of the service level, defined by $b/(b+h)$. Note that in this example, the normal approximation to the shortfall distribution yields better results than the negative binomial approximation. However, we have found this to be due to the choice of numeric experiment; both approximations do consistently give substantially higher costs, but which one is better depends on the numerical example considered.

Inspired by this observation, we attempted to find the root cause of this overestimation. By looking at the probability density function (pdf) of the shortfall (see Figure 8), we observe that if φ is sufficiently different from 0, the pdf of the shortfall becomes multi-modal (both for negative and positive correlation) and the normal/negative binomial distributions are no longer a good fit for the distribution of the shortfall. This effect is much more prevalent for high lead time correlation (high values of φ). Due to this multi-modality, the variance of the shortfall increases significantly, whereas the right tail of the distribution does not become fat. This results in an over-estimation of the base-stock level when uni-modal approximations of the shortfall are used, as is also evident from the left panel of Figure 6. The same multi-modality is observed (albeit to a lesser extent) for negative lead time correlation, which also leads to (albeit more modest) overestimations of the base-stock level (see Figure 6 left panel).

This numerical experiment revealed that in the presence of positive lead time correlation, ignoring the correlation and assuming i.i.d. lead times,

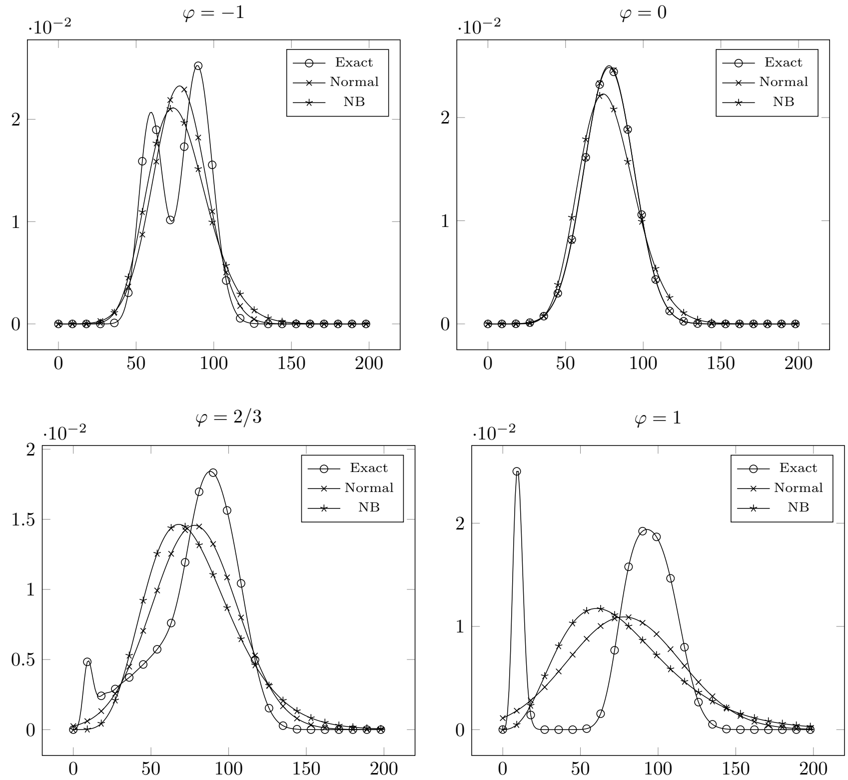

Figure 8: The probability density functions of the (exact) shortfall distribution compared to those of the normal and negative binomial estimates for different values of φ.

leads to an underestimation of the base-stock level as the correlation makes the inventory levels more volatile. Alternatively, when we take the correlation into account and fit a uni-modal distribution (e.g., normal or negative binomial) with the correct mean and variance, we again introduce an error, but in this case we overestimate the base-stock level due to the multi-modal character of the shortfall distribution. In case of strongly correlated lead times, it is advised to make use of (near-)exact methods like the one proposed in this paper, or the numerical method suggested by Muharremoglu and Yang (2010) when historical lead time data are available.

6. Conclusions

This paper studies the impact of correlation in stochastic lead times, in combination with order crossovers on inventories. When lead times are correlated, the lead time demand is not impacted, but the shortfall distribution and hence the inventory distribution and optimal base-stock levels are. With the exception of Bischak et al. (2014) and Muharremoglu and Yang (2010), the literature that takes order crossovers into account, assumes i.i.d. lead times. We contribute to the literature by providing structural results on the impact of lead time correlation on the inventory distribution. We also study how the current approximations (which assume i.i.d. lead times) perform in the presence of lead time correlation.

Our results show that higher correlation values generally imply higher inventory variance. This effect is significant for positive correlation (i.e., when going from uncorrelated to perfect correlation), and rather modest (and sometimes even reversed) when going from the most negative correlation to no correlation. We established lower and upper bounds on the variance of the shortfall: whereas the upper bound is tight, the lower bound is only tight for a certain choice of state space and steady state. We characterized the latter dependence for the 2-state Markov modulated lead time process, which enabled us to derive tight lower bounds for arbitrary steady state and state space.

As the optimal base-stock levels are highly dependent on the correlation in lead times, it should therefore be taken into account when setting base-stock levels. Due to the fact that the shortfall distribution, and thus also the inventory distribution, becomes multi-modal for high lead time correlation, the base-stock level increases less than one may first expect based on the variance of the shortfall. Given the non-negligible impact of the lead time correlation and the multi-modal distribution of the shortfall distribution for high lead time correlation, it is advised to apply exact methods like the one proposed in this paper.

7. Acknowledgements

This research was supported by the Research Foundation-Flanders (FWO-Vlaanderen), grant no. G069016N. The authors acknowledge with appreciation the generous contributions and constructive suggestions of the two anonymous referees.

References

Axsäter, S., 2011. Inventory control when the lead-time changes. Production and Operations Management 20 (1), 72–80.

Bischak, D. P., Robb, D. J., Silver, E. A., Blackburn, J. D., 2014. Analysis and management of periodic review, order-up-to level inventory systems with order crossover. Production and Operations Management 23 (5), 762–772.

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

Bradley, J. R., Robinson, W, L., 2005. Improved base-stock approximations for independent stochastic lead times with order crossover. Manufacturing & Service Operations Management 7 (4), 319–329.

Disney, S. M., Maltz, A., Wang, X., Warburton, R. D., 2016. Inventory management for stochastic lead times with order crossovers. European Journal of Operational Research 248 (2), 473–486.

Gaalman, G., Riezebos, J., 2005. Managing expected inventory order crossovers. Order 1, o2.

Hadley, G., Whitin, T. M., 1963. Analysis of inventory systems. Prentice Hall.

Hayya, J. C., Bagchi, U., Kim, J. G., Sun, D., 2008. On static stochastic order crossover. International Journal of Production Economics 114 (1), 404–413.

Hayya, J. C., Harrison, T. P., He, X. J., 2011. The impact of stochastic lead time reduction on inventory cost under order crossover. European Journal of Operational Research 211 (2), 274–281.

He, X. X., Xu, S. H., Ord, J. K., Hayya, J. C., 1998. An inventory model with order crossover. Operations Research 46, S112–S119.

Kaplan, R. S., 1970. A dynamic inventory model with stochastic lead times. Management Science 16 (7), 491–507.

Muharremoglu, A., Tsitsiklis, J. N., 2008. A single-unit decomposition approach to multiechelon inventory systems. Operations Research 56 (5), 1089–1103.

Muharremoglu, A., Yang, N., 2010. Inventory management with an exogenous supply process. Operations Research 58 (1), 111–129.

Riezebos, J., Gaalman, G., 2009. A single-item inventory model for expected inventory order crossovers. International Journal of Production Economics 121 (2), 601–609.

Riezebos, J., Zhu, S. X., 2015. Mrp planned orders in a multiple-supplier environment with differing lead times. Production and Operations Management 24 (6), 883–895.

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

Scarf, H., 1959. The optimality of (s, S) policies in the dynamic inventory problem.

Song, J.-S., Zipkin, P. H., 1996. Inventory control with information about supply conditions. Management science 42 (10), 1409–1419.

Srinivasan, M., Novack, R., Thomas, D., 2011. Optimal and approximate policies for inventory systems with order crossover. Journal of Business Logistics 32 (2), 180–193.

Wang, X., Disney, S. M., 2017. Mitigating variance amplification under stochastic lead-time: The proportional control approach. European Journal of Operational Research 256 (1), 151–162.

Wensing, T., Kuhn, H., 2015. Analysis of production and inventory systems when orders may cross over. Annals of Operations Research 231 (1), 265–281.

Zalkind, D., 1976. Further results for order-level inventory systems with independent stochastic leadtimes. Technical Report No. 76–6.

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

Appendix A. List Of Notations

$t \in \mathbb{N}$	Time period
$\{L_t\}$	Lead time process
$\mathcal{S}$	Countable state space of $\{L_t\}$
π	Steady state of $\{L_t\}$
$\mathbf{L}$	Transition matrix of $\{L_t\}$ if it is a discrete time Markov chain
$\{D_t\}$	Demand process
Q_t	Order Quantity placed in period t
SF_t	Shortfall at the end of period t
I_t	Inventory at the end of period t
S	Base-stock level
h	Per unit holding cost
b	Per unit shortage cost
f^+	Defined as $f^+(s) = f(s) \vee 0 = \max\{f(s), 0\}$
f^-	$(-f)^+$
LTD_t	Lead time demand at time t
V_t	Number of outstanding orders at time t
$\stackrel{d}{=}$	Equality in distribution
$\rightarrow_d$	Convergence in distribution
$g_{(s;k\|n)}$	Probability that of the last n orders, exactly k are outstanding and the order placed n periods ago has lead time s
$g_{(k\|n)}$	Probability that of the last n orders, k are still outstanding
g_k	Probability that there are k outstanding orders
ℓ_n	Correlation between L_t and L_{t+n}

ℓ	ℓ_1
$X^{(n)}$	The $n-$fold convolution of a random variable X
μ_X	The mean of a random variable X
σ_X^2	The variance of a random variable X

Appendix B. Example associated to Section 3.1

Assume a system with lead time process $\{L_t\}$ where no order crossovers occur and the possible lead times are $0, 1$ and 2. Assume further that $\{L_t\}$ follows a Markov process, with transition matrix $\mathbf{L}$. The shortfall distribution remains the same as long as the steady state does not change, for example if we take the steady state $\pi := \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{pmatrix}$ then we find that all transition matrices of the form:

$$\mathbf{L} = \begin{pmatrix} a & b & 1-a-b \\ 1-a & c & a-c \\ 0 & 1-b-c & b+c \end{pmatrix}, 0 \leq a, b, a+b, c, a-c, b+c \leq 1,$$

lead to the same shortfall distribution. More generally, as long as we have for any $s \in \mathbb{N}, s_1, s_2 \in \mathcal{S}, s_1 < s < s_2$ that also $s \in \mathcal{S}$ and $\mathbf{L}$ is an upper Hessenberg matrix then the shortfall distribution is independent of the chosen transition matrix given some fixed steady state. However once we allow the possibility of having crossovers, the shortfall distribution becomes dependent on the correlation in $\{L_t\}$.

Appendix C. Proof of Theorem 1

We first need to introduce some new notation. For $n \in \mathbb{N}$ we denote $q'_{(s,i_{n-1},\ldots,i_0)}, s \in \mathcal{S}, i_{n-1}, \ldots, i_0 \in \{0, 1\}$ as the probability that the order placed n periods ago has lead time j and for $k \in \{0, \ldots, n-1\}$ we have that the order placed k periods ago has arrived if $i_k = 0$ and has not yet arrived if $i_k = 1$.

Formally we have:

$$q'_{(s;i_{n-1},\ldots,i_0)} := \mathbb{P}\{L_{t-n} = j, L_{t-(n-1)} \sim_{i_{n-1}} n-1, \ldots, L_t \sim_{i_0} 0\}$$

where we denote:

$$\sim_i := \begin{cases} \leq & \text{if } i = 0 \\ > & \text{if } i = 1. \end{cases}$$

We have the following recursive relation for these values:

$$\begin{aligned}
q'_{(s;i_{n-1},\ldots,i_0)} &= \mathbb{P}\{L_{t-n} = s; L_{t-(n-1)} \sim_{i_{n-1}} n-1, \ldots, L_t \sim_{i_0} 0\} \\
&= \sum_{s_{n-1}\in\mathcal{S}, s_{n-1}\sim_{i_{n-1}}(n-1)} \mathbb{P}\{L_{t-n} = s, L_{t-(n-1)} = s_{n-1}, L_{t-(n-2)} \sim_{i_{n-2}} n-2, \ldots, L_t \sim_{i_0} 0\} \\
&= \sum_{s_{n-1}\in\mathcal{S}, s_{n-1}\sim_{i_{n-1}}(n-1)} \mathbb{P}\{L_{t-n} = s \mid L_{t-(n-1)} = s_{n-1}\} \cdot q'_{(s_{n-1},i_{n-2},\ldots,i_0)} \\
&= \sum_{s_{n-1}\in\mathcal{S}, s_{n-1}\sim_{i_{n-1}}(n-1)} \frac{\pi_s}{\pi_{s_{n-1}}} \mathbf{L}_{s,s_{n-1}} q'_{(s_{n-1};i_{n-2},\ldots,i_0)}.
\end{aligned}$$

Remark: In case demand is not assumed to be i.i.d., one should take into account which orders are still outstanding to find the shortfall distribution. One can do this by using $q_{(i_n,\ldots,i_0)} := \sum_{s\in\mathcal{S}} q'_{(s;i_n,\ldots,i_0)}$, which can be calculated by applying the above recursion.

We use this recursive relation to show the correctness of the formula for $g_{(s;k|n)}$. Let $A_k^{n+1} := \{(i_0,\ldots,i_n) \in \{0,1\}^{n+1} \mid \sum_{j=0}^n i_j = k\}$, we find:

$$\begin{aligned}
g_{(s;k|(n+1))} &= \sum_{(i_0,\ldots,i_n)\in A_k^{n+1}} q'_{(s;i_n,\ldots,i_0)} \\
&= \sum_{(i_0,\ldots,i_n)\in A_k^{n+1}} \sum_{s_n\in\mathcal{S}, s_n\sim_{i_n} n} \frac{\pi_s}{\pi_{s_n}} \mathbf{L}_{s,s_n} q'_{(s_n;i_{n-1},\ldots,i_0)} \\
&= \sum_{s_n\in\mathcal{S}, s_n\leq n} \frac{\pi_s}{\pi_{s_n}} \mathbf{L}_{s,s_n} g_{(s_n;k|n)} + \sum_{s_n\in\mathcal{S}, s_n>n} \frac{\pi_s}{\pi_{s_n}} \mathbf{L}_{s,s_n} g_{(s_n;(k-1|n))}.
\end{aligned}$$

Appendix D. Proof of Lemma 1

We denote the transition matrix associated to $\{L_t\}$ by $\mathbf{L} := \begin{pmatrix} a & 1-a \\ b & 1-b \end{pmatrix}$, $0 \leq a, b \leq 1$ and its steady state by $\pi := \begin{pmatrix} \alpha & 1-\alpha \end{pmatrix}$, $0 < \alpha < 1$. As we have $\pi\mathbf{L} = \pi$ and $\alpha \notin \{0,1\}$ we find that:

$$\mathbf{L} = \begin{pmatrix} a & 1-a \\ \frac{\alpha(1-a)}{1-\alpha} & \frac{\alpha a - 2\alpha + 1}{1-\alpha} \end{pmatrix},$$

with $\max\{\frac{2\alpha-1}{\alpha},0\} \leq a < 1$. Therefore:

$$\begin{aligned}\ell = \mathrm{Corr}(L_t, L_{t+1}) &= \frac{\mathbb{E}[L_t L_{t+1}] - \mu_L^2}{\sigma_L^2} \\ &= \frac{m^2(\alpha a - 2\alpha + 1) - m^2(1-\alpha)^2}{(1-\alpha)^2 - (1-\alpha)} = \frac{a-\alpha}{1-\alpha}.\end{aligned}$$

From this we already see that ℓ does not depend on m. It is easily checked that $\mathrm{Det}(\mathbf{L}) = \frac{a-\alpha}{1-\alpha}$. The second part now follows by defining the process $\{L_t'\}$ as $L_t' := L_{tn}$ for each $t \in \mathbb{N}$. Denote $\mathbf{L}'$ for the transition matrix corresponding to $\{L_t'\}$ and ℓ_n' for its lag$-n$ correlation. We find:

$$\ell_n = \ell_1' = \mathrm{Det}(\mathbf{L}') = \mathrm{Det}(\mathbf{L}^n) = \ell_1^n.$$

We saw above that $\ell = \frac{a-\alpha}{1-\alpha}$, using this and the fact that $\ell_n = \ell^n$ we instantly find:

$$\mathbf{L}^n = \begin{pmatrix} (1-\alpha)\ell^n + \alpha & (1-\alpha)(1-\ell^n) \\ \alpha(1-\ell^n) & \alpha\ell^n + (1-\alpha) \end{pmatrix}.$$

In order for this to be a transition matrix for $n = 1$ is $\alpha/(\alpha-1) \leq \ell$ if $\alpha \in [0, 1/2]$ and $(\alpha-1)/\alpha \leq \ell$ if $\alpha \in [1/2, 1]$. This completes the proof as any product of transition matrices is again a transition matrix.

Appendix E. Proof of Proposition 1

Applying Lemma 1 we find:

$$\begin{aligned}\mathbb{E}[|V_t - L_t|] &= \pi_0 \mathbb{E}[V_t \mid L_t = 0] + \pi_m \mathbb{E}[(m - V_t) \mid L_t = m] \\ &= \pi_0 \mathbb{E}\left[\sum_{k=0}^{m-1} \delta\{L_{t-k} > k\} | L_t = 0\right] + m\pi_m - \pi_m \mathbb{E}\left[\sum_{k=0}^{m-1} \delta\{L_{t-k} > k\} | L_t = m\right] \\ &= \pi_m \left(\sum_{k=0}^{m-1} (\mathbf{L}^k)_{m0}\right) + \pi_m \left(m - \sum_{k=0}^{m-1} (\mathbf{L}^k)_{mm}\right) \\ &= \frac{2(1-\alpha)\alpha\left(\ell^m - \ell m + m - 1\right)}{1-\ell}.\end{aligned}$$

Appendix F. Proof of Corollary 1

Suppose we have some sequence $\{\ell^{(n)}\}$ in $[0, 1]$ which converges to 1. We use the notation $V^{(n)}$ and $\mathrm{SF}^{(n)}$ for the associated number of outstanding orders and shortfall distribution.

We show that the characteristic function (CF) of the shortfall converges to the CF of the lead time demand (this is equivalent to weak convergence). For a random variable X, the CF is defined as $\phi_X(t) := \mathbb{E}[e^{itX}]$, which yields:

$$\begin{aligned}
\phi_{\mathrm{SF}^{(n)}}(t) &= \mathbb{E}[e^{it\mathrm{SF}^{(n)}}] \\
&= \mathbb{E}\left[\prod_{k=0}^{V^{(n)}} e^{itD_k}\right] \\
&= \mathbb{E}\left[\prod_{k=0}^{V^{(n)}} \mathbb{E}[e^{itD_k} \mid V^{(n)}]\right] \\
&= \mathbb{E}[\phi_D(t)^{V^{(n)}}].
\end{aligned}$$

Analogously one finds $\phi_{\mathrm{LTD}}(t) = \mathbb{E}[\phi_D(t)^L]$. We should therefore show that for any $a \in \mathbb{C}$ we have $\mathbb{E}[a^{V^{(n)}}]$ converges to $\mathbb{E}[a^L]$ for $n \to \infty$. As we have convergence of $V^{(n)}$ to L in L^1 norm, this convergence also holds in distribution. Since $V^{(n)}$ and L are both bounded (let M be an upper bound for both these values) this entails the sought convergence. Indeed letting $f(x) := a^x$ for $x \leq M$ and $f(x) = M^x$ for $x > M$ we find: $\mathbb{E}[a^{V^{(n)}}] = \mathbb{E}[f(V^{(n)})] \to \mathbb{E}[f(L)] = \mathbb{E}[a^L]$.

Appendix G. Proof of Lemma 3

We can calculate the variance of the shortfall in function of V. By applying the law of total variance in the second equality, we find:

$$
\begin{aligned}
\sigma_{\mathrm{SF}}^2 &= \mathrm{Var}\left(\sum_{l=0}^{V_t} D_l\right) \\
&= \mathbb{E}\left[\mathrm{Var}\left(\sum_{l=0}^{V_t} D_l \mid V_t\right)\right] + \mathrm{Var}\left(\mathbb{E}\left[\sum_{l=0}^{V_t} D_l \mid V_t\right]\right) \\
&= \sum_{k=0}^{\infty} \mathbb{P}\{V=k\}\sigma_D^2(k+1) + \sum_{k=0}^{\infty}((k+1)\mu_D)^2\mathbb{P}\{V=k\} - \mu_D^2(\mu_L+1)^2 \\
&= \sigma_D^2\mathbb{E}[V+1] + \mu_D^2\mathbb{E}[(V+1)^2] - \mu_D^2(\mu_L+1)^2 \\
&= \sigma_D^2(\mu_V+1) + (\sigma_V^2+\mu_V^2+2\mu_V+1)\mu_D^2 - \mu_D^2(\mu_L+1)^2 \\
&= \sigma_D^2(\mu_L+1) + (\sigma_V^2+\mu_L^2+2\mu_L+1)\mu_D^2 - \mu_D^2(\mu_L+1)^2 \\
&= (\mu_L+1)\sigma_D^2 + \sigma_V^2\mu_D^2.
\end{aligned}
$$

Appendix H. Proof of Proposition 2

For the variance of the number of outstanding orders, we find:

$$
\begin{aligned}
\sigma_V^2 &= \mathrm{Var}\left(\sum_{k=0}^{m-1} \delta\{L_{t-k} > k\}\right) \\
&= \sum_{k=0}^{m-1} \mathrm{Var}(\delta\{L>k\}) + 2\sum_{0\le k<l<m} \mathrm{Cov}(\delta\{L_{t-l}>l\}, \delta\{L_{t-k}>k\}).
\end{aligned}
$$

As we have $L \in \{0, m\}$ we have for any $k \in \{0, \dots, m-1\}, t \in \mathbb{N} : \delta\{L > k\} = \frac{L}{m}$. This allows us to express:

$$
\sigma_V^2 = \frac{\sigma_L^2}{m} + \frac{2}{m^2}\sum_{0\le k<l<m} \mathrm{Cov}(L_{t-l}, L_{t-k}).
$$

In case we work with i.i.d. lead times we find (using Lemma 3):

$$
\sigma_{\mathrm{SF}}^2 = (1+\mu_L)\sigma_D^2 + \frac{\mu_D^2\sigma_L^2}{m} = \sigma_{const}^2 + \frac{\mu_D^2\sigma_L^2}{m},
$$

which is the reason why we denote this value by $\sigma^2_{i.i.d.}$. In the correlated case we note that:

$$\mathrm{Cov}(L_{t-l}, L_{t-k}) = \sigma_L^2 \mathrm{Corr}(L_0, L_{l-k}) = \sigma_L^2 \ell_{l-k}.$$

This yields:

$$\sigma_V^2 = \frac{\sigma_L^2}{m} + \frac{2\sigma_L^2}{m^2} \sum_{k=1}^{m-1} (m-k)\ell_k.$$

From this we can infer:

$$\sigma^2_{\mathrm{SF}} = \sigma^2_{i.i.d.} + 2\left(\frac{\mu_D \sigma_L}{m}\right)^2 \sum_{k=1}^{m-1} (m-k)\ell_k.$$

Appendix I. Proof of Lemma 4

Fix some arbitrary $n \in \mathbb{N}$, it obviously suffices to show the claim for

$$p(x) = \sum_{k=1}^{n} (n+1-k)x^k.$$

Taking the sum from 0 to $n-1$ and taking the derivative yields:

$$p'(x) = \sum_{k=0}^{n-1} (n-k)(k+1)x^k.$$

It now suffices to show that $p'(x) \geq 0$ on $[-1, 1]$. One can check that we have:

$$p'(x) = \frac{nx^{n+2} - (2+n)x^{n+1} + (2+n)x - n}{(x-1)^3},$$

it thus suffices to show that $\frac{nx^{n+2}-(2+n)x^{n+1}+(2+n)x-n}{x-1} \geq 0$ for $x \in [-1, 1]$. This polynomial is equal to $n(x^{n+1}+1) - 2(x^n + \cdots + x)$ thus we need to show that

$$\frac{(x^{n+1}+1)}{2} \geq \frac{x^n + \cdots + x}{n}.$$

For $x \in [-1, 0]$ we have $\frac{(x^{n+1}+1)}{2} \geq 0$ while $x^n + \cdots + x = x(1-x^n)/(1-x) \leq 0$, hence we may assume $x \in (0, 1]$. We do this part by induction skipping the case $n = 1$. We have:

$$\frac{x^{n+1} + \cdots + x}{n+1} = \frac{nx\frac{x^n+\cdots+x}{n} + x}{n+1} \leq \frac{nx\frac{x^{n+1}+1}{2} + x}{n+1}.$$

It remains to show that the inequality $(nx\frac{x^{n+1}+1}{2}+x)/(n+1) \leq (x^{n+2}+1)/2$ holds. Subtracting the left hand side from the right hand side we obtain:

$$\frac{x^{n+2} - ((n+2)x - n - 1)}{2(n+1)},$$

which is positive as x^{n+2} is convex and the tangent line at 1 is given by $y = ((n+2)x - n - 1)$.

Appendix J. Proof of Proposition 3

From the set of possible lead time processes we look at some special cases, in particular the highest/lowest possible correlation and how they compare to having constant lead times. From these we will already see that the impact of positive correlation will be bigger than that of negative correlation.

In Lemma 1 we showed that:

$$\mathbf{L} = \begin{pmatrix} (1-\alpha)\ell + \alpha & (1-\alpha)(1-\ell) \\ \alpha(1-\ell) & \alpha\ell^n + (1-\alpha) \end{pmatrix}, \begin{cases} \frac{\alpha}{\alpha-1} \leq \ell & \text{if } \alpha \in [0, \frac{1}{2}], \\ \frac{\alpha-1}{\alpha} \leq \ell & \text{if } \alpha \in [\frac{1}{2}, 1]. \end{cases}$$

We graphically represent this definition domain in Figure J.9 (i.e. $\{(\alpha, \ell) \in \mathbb{R} \times \mathbb{R} \mid \mathbf{L} \text{ is the transition matrix of some Markov process}\}$ is shown). We discuss the boundaries a, b, c, d and e of the domain of (α, ℓ) and the special line f.

(a, b) These correspond to having a steady state of $\begin{pmatrix} 1 & 0 \end{pmatrix}$ resp. $\begin{pmatrix} 0 & 1 \end{pmatrix}$, i.e. we have a lead time process which is constant.

(c) This part corresponds to having near perfect correlation (we get arbitrarily close to correlation 1), i.e. $\mathbf{L} = \begin{pmatrix} 1-\varepsilon & \varepsilon \\ a \cdot \varepsilon & 1 - a \cdot \varepsilon \end{pmatrix}$ here $\varepsilon > 0$ and $a > 0$ are s.t. $0 < a \cdot \varepsilon < 1$ this ε. For low values of ε, $\mathbf{L}$ corresponds to having strong correlation, a can be chosen to control the steady state of $\mathbf{L}$. This means that if we observe the system at an arbitrary point it has lead time 0 with probability α (one can find an appropriate value for a s.t. $\mathbf{L}$ has steady state $\pi = \begin{pmatrix} \alpha & 1-\alpha \end{pmatrix}$) and m with probability $1 - \alpha$ but as we go further in the future the lead time remains unchanged with high probability. We can thus see the

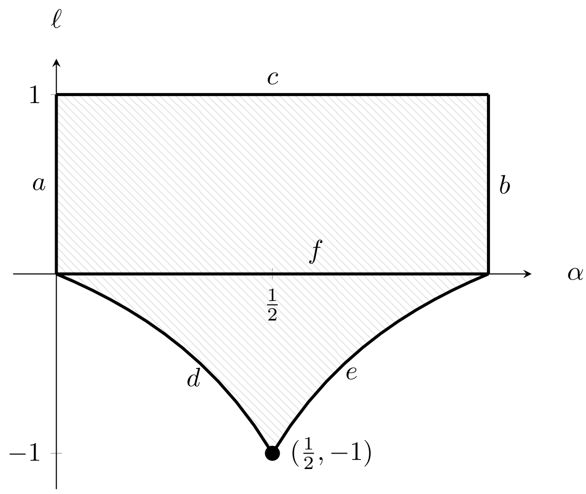

Figure J.9: Definition domain of $\mathbf{L}$ in function of α and ℓ.

lead time process as being an arbitrary, but unknown constant value. As we have shown that $\frac{\partial \sigma^2_{\mathrm{SF}}}{\partial \ell} \geq 0$, we have for any fixed α that taking the limit $\ell \to 1$ gives an upper bound for σ^2_{SF}. We find (by using the continuity of $f(x, m)$ in x):

$$\begin{aligned}
\sigma^2_{\mathrm{SF}} &\leq \lim_{\ell \to 1} \left(\sigma^2_{i.i.d.} + 2\mu_D^2 \sigma_L^2 \cdot f(\ell, m)\right) \\
&= \sigma^2_{i.i.d.} + 2\mu_D^2 \sigma_L^2 \cdot f(1, m) \\
&= \sigma^2_{i.i.d.} + \left(\frac{m-1}{m}\right) \mu_D^2 \sigma_L^2 \\
&= \sigma^2_{\mathrm{LTD}}
\end{aligned}$$

Note that the influence of the correlation in the lead times is completely contained in $\mu_D^2 \sigma_V^2$, thus the part of σ^2_{SF} which is due to the correlation in the lead time process can not become larger than $\mu_D^2 \sigma_L^2$. This upper bound corresponds to the increase of variation going from the i.i.d. case to perfect correlation. Moreover one can easily verify that $\sigma^2_{\mathrm{LTD}} = (\mu_L + 1)\sigma_D^2 + \mu_D^2 \sigma_L^2$ which corresponds to the upper bound we found. These results are generalized in the multi-dimensional setting, see Theorem 3.

(d, e) These correspond to the minimal possible correlation for every value of α, note that perfect negative can only be obtained in case of $\alpha = \frac{1}{2}$ since in this case we constantly switch between the two possible lead times so both occur an equal amount of times. Evaluating in these bounds gives us a lower bound for σ^2_{SF}. Fix some $\alpha \in (0,1)$ then we need to distinguish two cases.

$\alpha \leq \frac{1}{2}$: In this case we find:

$$\begin{aligned}
\sigma^2_{\text{SF}} &\geq \sigma^2_{i.i.d.} + 2\mu_D^2\sigma_L^2 \cdot f\left(\frac{\alpha}{\alpha-1}, m\right) \\
&= \sigma^2_{i.i.d.} + 2\mu_D^2\sigma_L^2 \cdot \frac{1}{m^2}\alpha\left((\alpha-1)\left(\left(\frac{\alpha}{\alpha-1}\right)^m - 1\right) - m\right) \\
&= \sigma^2_{i.i.d.} - 2\mu_D^2\sigma_L^2 \cdot \frac{(m-\mu_L)\left(\mu_L\left(\left(1-\frac{m}{\mu_L}\right)^m - 1\right) + m^2\right)}{m^4}
\end{aligned}$$

$\alpha \geq \frac{1}{2}$: In this case we find:

$$\begin{aligned}
\sigma^2_{\text{SF}} &\geq \sigma^2_{i.i.d.} + 2\mu_D^2\sigma_L^2 \cdot f\left(\frac{\alpha-1}{\alpha}, m\right) \\
&= \sigma^2_{i.i.d.} + 2\mu_D^2\sigma_L^2 \cdot \frac{1}{m^2}(\alpha-1)\left(\alpha\left(\left(\frac{\alpha-1}{\alpha}\right)^m - 1\right) + m\right) \\
&= \sigma^2_{i.i.d.} - 2\mu_D^2\sigma_L^2 \cdot \frac{\mu_L\left(m\left(\left(\frac{\mu_L}{\mu_L-m}\right)^m + m - 1\right) - \mu_L\left(\left(\frac{\mu_L}{\mu_L-m}\right)^m - 1\right)\right)}{m^4}
\end{aligned}$$

Note that both these cases include the case $\alpha = \frac{1}{2}$ and that for this α we find:

$$\begin{aligned}
\sigma^2_{\text{SF}} &\geq \sigma^2_{i.i.d.} + 2\mu_D^2\sigma_L^2 \cdot \frac{1}{4m^2}\left(-2m + (-1)^{m+1} + 1\right) \\
&= \sigma^2_{const} + \frac{\mu_D^2\sigma_L^2}{2m^2}((-1)^{m+1} + 1). \qquad \text{(J.1)}
\end{aligned}$$

Here we see that $\frac{\mu_D^2\sigma_L^2}{2m^2}((-1)^{m+1} + 1)$ is zero whenever m is even, we see that in this case we get $\sigma^2_{\text{SF}} = \sigma^2_{const}$ which corresponds to the case $\sigma^2_V = 0$. Indeed, for m even and perfect negative correlation we find that σ^2_V is zero, this might seem odd at first but trivial when we draw a picture (see Figure J.10). We see that the general formulation of the lower bound is a lot more complex than that of the upper bound, in the more dimensional case, we will use the simple lower bound σ^2_{const},

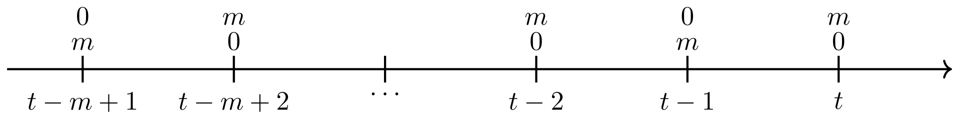

Figure J.10: As we have perfect negative correlation the only two possibilities for our lead time process are depicted in the above drawing (namely $0m \dots 0m$ and $m0 \dots m0$) and in both cases the number of outstanding orders is given by $\frac{m}{2}$.

which clearly always holds and can here be achieved for $\alpha = 1/2$ and m even.

(f) This line corresponds to having $\ell = 0$, i.e. the case where we assume $\{L_t\}$ is an i.i.d. sequence, we get the transition matrix:

$$\mathbf{L} = \begin{pmatrix} \alpha & 1-\alpha \\ \alpha & 1-\alpha \end{pmatrix}.$$

As noted before, we get $\sigma^2_{\mathrm{SF}} = \sigma^2_{i.i.d.}$.

Appendix K. Proof of Proposition 4

We now look at $\sigma^2_{\mathrm{SF}}/\sigma^2_{const}$ for some special cases of σ^2_{SF} namely the i.i.d. case, the lower bound and the upper bound. Let us denote these ratios by $R_{i.i.d.}, R_{lower}$ and R_{LTD}. In particular we get simple expressions for the limit of m to infinity, which corresponds to the case where we have 2 possible, sufficiently different, lead times. For the i.i.d. case we find:

$$R_{i.i.d.}(m,\alpha) := \frac{\sigma^2_{i.i.d.}}{\sigma^2_{const}} = 1 + \frac{\mu_D^2 \sigma_L^2}{m\sigma^2_{const}}.$$

We introduce the notation $\theta_m := \frac{\mu_D^2 \sigma_L^2}{m\sigma^2_{const}} \to \alpha\frac{\mu_D^2}{\sigma_D^2} =: \theta$, we thus see that $R_{i.i.d.}(m,\alpha) \underset{m\to\infty}{\longrightarrow} 1+\theta$ (note that the difference in going from constant lead time to i.i.d. lead time increases as α increases). For the upper bound we find that the ratio can not be greater than:

$$R_{\mathrm{LTD}}(m,\alpha) := R_{i.i.d.}(m,\alpha) + m\theta_m,$$

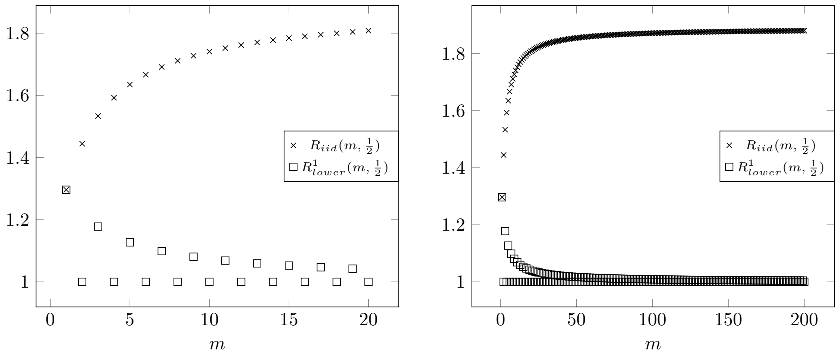

Figure K.11: For $\sigma_D^2 = 9, \mu_D = 4$ ($\theta = \frac{8}{9}$): The ratio of the variance of the shortfall for i.i.d. resp. perfectly negative correlated lead times to constant lead times in function of the maximum lead time m.

which tends to infinity at rate θ for $m \to \infty$. For the lower bound we find that R_{lower} splits into two parts depending on the value of α:

$$\begin{cases} R^1_{lower}(m,\alpha) = R_{i.i.d.}(m,\alpha) + \frac{2\theta_m}{m}\alpha\left((1-\alpha)\left(1-\left(\frac{\alpha}{\alpha-1}\right)^m\right)-m\right) & \text{if } \alpha \leq \frac{1}{2} \\ R^2_{lower}(m,\alpha) = R_{i.i.d.}(m,\alpha) - \frac{2\theta_m}{m}(1-\alpha)\left(\alpha\left(\left(\frac{\alpha-1}{\alpha}\right)^m-1\right)+m\right) & \text{if } \alpha \geq \frac{1}{2}. \end{cases}$$

We also obtain convergence, namely:

$$\begin{cases} R^1_{lower}(m,\alpha) \underset{m\to\infty}{\longrightarrow} 1 + (1-2\alpha)\theta \\ R^2_{lower}(m,\alpha) \underset{m\to\infty}{\longrightarrow} 1 + (1-2(1-\alpha))\theta \end{cases}$$

We take a look at the special case $\alpha = \frac{1}{2}$ and $\ell = -1$. We find:

$$R^1_{lower}(m, \frac{1}{2}) = 1 + \frac{1}{2}(1 + (-1)^{m+1})\frac{\theta_m}{m} \to 1.$$

This is the only case for which the ratio of the lower bound tends to one. The part $((-1)^{m+1}+1)$ makes the ratio go up/down every subsequent value of m. This effect is triggered by the variance of the amount of outstanding orders. See Figure K.11 for the convergence of $R_{i.i.d.}$ and R_{lower} with $\sigma_D^2 = 9, \mu_D = 4$ and $\alpha = 1/2$.

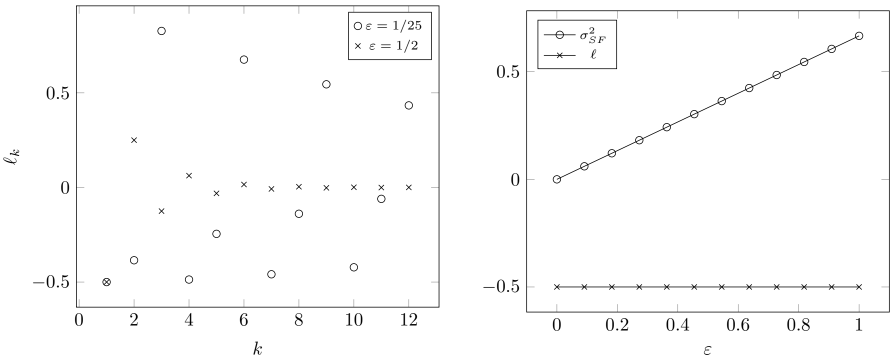

Figure L.12: The different lag−k correlations for the Markov processes given in Appendix L for $\varepsilon = 1/25, 1/2$.

Appendix L. It is no longer possible to find functions $(f_m)_{m\in\mathbb{N}_0}$ s.t. $\ell_m = f_m(\ell)$

Consider for any $\varepsilon \in (0,1)$ the transition matrix:

$$\mathbf{P}_\varepsilon := \begin{pmatrix} 0 & 1-\varepsilon & \varepsilon \\ \varepsilon & 0 & 1-\varepsilon \\ 1-\varepsilon & \varepsilon & 0 \end{pmatrix},$$

with state space $\mathcal{S} = \{0,1,2\}$. Then we clearly have independent of ε the steady state $\pi = \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{pmatrix}$. One can easily calculate that independent of ε we have $\ell = -1/2$ while $\ell_2 = -1/2 - 3(-1+\varepsilon)\varepsilon$. This means that we have found multiple (even uncountably infinite) Markov processes with the same steady state which all have the same lag−1 correlation but different lag−2 correlation (see Figure L.12). This example also shows that we do not have that the variance of the shortfall is non-decreasing in function of the lag−1 correlation, we see in Figure L.12 that we can change the variance of the shortfall by changing ε whilst the lag−1 correlation remains fixed.

Appendix M. Proof of Theorem 3

The lower bound is trivial. For the upper bound, assume we have a countable state space $\mathcal{S}$ and steady state π. We calculate:

$$\begin{aligned}\sigma_V^2 &= \mathrm{Var}\left(\sum_{k=0}^{\infty} \delta\{L_{t-k} > k\}\right) \\ &= \sum_{k=0}^{\infty} \mathrm{Var}(\delta\{L_{t-k} > k\}) + 2 \cdot \sum_{k<l} \mathrm{Cov}(\delta\{L_{t-k} > k\}, \delta\{L_{t-l} > l\})\end{aligned}$$

For $k < l$ we note that:

$$\begin{aligned}\mathbb{E}[\delta\{L_{t-k} > k\}\delta\{L_{t-l} > l\}] &\leq \mathbb{E}[\delta\{L_{t-l} > l\}] \\ &= \mathbb{E}[\delta\{L_t > k\}\delta\{L_t > l\}],\end{aligned}$$

which shows that $\mathrm{Cov}(\delta\{L_{t-k} > k\}, \delta\{L_{t-l} > l\}) \leq \mathrm{Cov}(\delta\{L_t > k\}, \delta\{L_t > l\})$. Using this inequality and the fact that we have for any t, $L_t = \sum_{k=0}^{\infty} \delta\{L_t > k\}$, we find:

$$\sigma_V^2 \leq \mathrm{Var}\left(\sum_{k=0}^{\infty} \delta\{L_t > l\}\right) = \sigma_L^2.$$

This completes the proof by applying Lemma 3.