European Journal of Operational Research 270 (2018) 198–217

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Production, Manufacturing and Logistics

Continuous inventory control with stochastic and non-stationary Markovian demand

Walid W. Nasr*, Ibrahim J. Elshar

Faculty of Engineering and Architecture, Industrial Engineering and Management, American University of Beirut (AUB), Beirut 1107 2020, Lebanon

ARTICLE INFO

Article history:
Received 10 May 2017
Accepted 19 March 2018
Available online 28 March 2018

Keywords:
Supply chain management
Phase-type
Time-dependent
Markovian
Non-stationary demand

ABSTRACT

Non-stationary demand is common in many industrial settings and accounting for the non-stationarity in the demand process significantly complicates the analysis of inventory policies. This work presents an efficient computational framework, which utilizes a Markovian representation, to model and solve for the stochastic and non-stationary performance measures of an inventory system. The non-stationary and stochastic characteristics of the demand process are captured by an approximate Phase-type distribution. The differential equations corresponding to the Markovian representation are presented along with an algorithmic approach to numerically solve for the non-stationary performance measures. Time-dependent (s_t, S_t) continuous replenishment policies with a fixed ordering cost are investigated over a finite time horizon. The trade-off between the computational complexity and cost effectiveness of the policies are investigated numerically under different cost and demand distribution parameters. The numerical study also investigates the accuracy of using the time-dependent Phase-type distribution to capture key descriptors of the non-stationary demand process.

1. Introduction

Accounting for the time-dependent components of a supply chain system is motivated by the wide range of industries reporting time-dependency in their demand process. Neale and Willems (2009) argue that non-stationary demand is the rule rather the exception. The authors refer to Hewlett-Packard, Microsoft, Dell, Kraft Foods among other companies which have to account for time-dependency and seasonality when evaluating their supply chain policies. Although supply chain managers frequently encounter non-stationary demand and are aware of the inefficiencies of ignoring time-dependency, the added value of the existing tools which account for non-stationarity is offset by their added level of complexity (Ehrenthal, Honhon, & Van Woensel, 2014). The literature on stochastic non-stationary demand in the context of inventory systems is scarce in comparison to stationary demand, which is mainly due to the complicated nature of time-dependent models (Tunc, Kilic, Tarim, & Eksioglu, 2011). A main contribution of this work is to present a computational framework to efficiently capture the time-dependent behavior of inventory systems.

The majority of the literature on non-stationary systems deals with calculating the optimal expected cost or providing near optimal reorder policies and provides little insight on the behavior of the inventory system over time. The policies obtained from the existing models involve solving a recursive or sequential cost function over multiple periods. This makes the system less transparent and it becomes difficult for a supply chain practitioner to quantify and predict the impact of environmental changes on the system and conduct sensitivity analysis. Furthermore, estimating the analytical performance of inventory systems with non-stationary demand are mostly intractable and scarce in the inventory literature (Gershwin, Dallery, Papadopoulos, & Smith, 2012). In this work, we present an approach which closely evaluates the time dependent behavior of the performance measures of the inventory system. This is achieved by describing a Markovian representation of the inventory system and presenting an efficient numerical approach to calculate the non-stationary probability distribution of the inventory measures. Consequently, for any specified reorder policy, the behavior of the performance measures is observed over all points in time. For example, the mean and variance of the on-hand inventory and backorders are continuously calculated for all points in time and over any specified time interval. Another commonly used measure of system performance is the stockout probability. The model presented calculates the stockout probability over time and identifies the points in time where shortages have a high probability of occurring. Furthermore, it is possible to identify the specific points in time where the expected inventory reaches its peak as well as calculate the probability the inventory/backorders

* Corresponding author.
E-mail addresses: wn12@aub.edu.lb (W.W. Nasr), ije04@aub.edu.lb (I.J. Elshar).

https://doi.org/10.1016/j.ejor.2018.03.023
0377-2217/

exceeds a certain threshold. Obviously, such a detailed description of system behavior results in a more informed and time-phased decision making process.

Although Monte-Carlo simulation appears to be a straightforward approach to investigate the continuous behavior of the system over time, Monte-Carlo simulations can be tedious and computationally prohibitive for time dependent analysis (Nasr & Taaffe, 2013). Simulation is more suitable for stationary systems where the entire simulation run contributes to the accuracy of a single point estimate. In the time-dependent case, every replication yields one data point for a certain point in time. To compute the system measures at time t, multiple independent replications of the simulation are required to obtain point estimates within acceptable confidence intervals. Furthermore, the few available tools to simulate stochastic and time-dependent point-processes (from which we mention Gerhardt & Nelson, 2009a; Ma & Whitt, 2016) significantly add to the complexity of the simulation model.

In this work we utilize a time dependent Phase-type, Ph_t, distribution to model the stochastic and time-dependent demand process. Obviously, the functionality of the model is dependent on the flexibility and accuracy of the Ph_t process to capture the stochastic and time-dependent behavior of the demand process. There is an abundance of existing approaches which fit Ph distributions to stationary point processes. The existing literature also illustrates the flexibility and accuracy of the stationary Ph distribution when utilized as an approximate process (see Buchholz, Kriege, & Felko, 2014; Gerhardt & Nelson, 2009b for a thorough review on fitting point processes with Markovian distributions). The literature on fitting time-dependent Markovian processes is scarce in comparison to the stationary case. Recently, Nasr and Taaffe (2013) present an approach which fits a Ph_t process to the time-dependent and stochastic characteristics of a point process. This is achieved by fitting the parameters of a balanced two-level mixture of Erlangs (2-MECO) distribution to the rate and standard deviation of the approximated time-dependent counting process.

This work centers around describing the Markovian representation of the inventory system as well as the numerical solution of the corresponding differential equations. Accordingly, the time-dependent performance of the inventory system for a certain (s_t, S_t) policy can be closely observed and evaluated. An additional contribution of this work is calculating the expected cost for a given policy where we investigate the computational efficiency of the model in searching for a reorder policy to minimize expected cost.

2. Literature review

Classical papers by Scarf (1960) and Veinott, Jr (1966) studied the optimality of the (s, S) ordering policy for single product inventory systems with stochastic demand, periodic review, proportional backlog and holding costs, and fixed ordering cost. Scarf (1960) proved the optimality of the (s, S) policy for a finite horizon stationary multi-period problem with convex holding/backlog cost. Scarf's (1960)) proof immediately extends to non-stationary costs and demand distributions. This resulted in a stream of literature directed towards solving for the optimal or near optimal time varying parameters of the reorder policy. This proved to be a computationally challenging problem and we present a brief review of the existing approaches.

2.1. Computational approaches for non-stationary (s, S) policies

The work in Morton and Pentico (1995) utilizes the simplifying assumption of linear ordering costs and conducts an extensive comparison of four heuristic approaches to compute the near optimal solution of the time-varying (s, S) reorder parameters. This illustrates the computational complexity involved in calculating the optimal time dependent parameters for the relatively simpler model which does not account for a fixed ordering cost. A heuristic approach to compute near optimal (s, S) parameters is presented in Bollapragada and Morton (1999) for a single stage inventory problem facing a general non-stationary demand with a fixed ordering cost and proportional backordering and holding costs. The authors assume no lead time and approximate the demand process by a piecewise constant demand processes where a myopic heuristic calculates a static (s, S) reorder policy for the corresponding approximate stationary system for every period. The resulting stationary (s, S) policy for every period is used as an approximation for the dynamic replenishment (s_t, S_t) reorder policy. The efficiency of the heuristic is compared to the optimal solution obtained by solving the corresponding dynamic program. Tunc et al. (2011) investigate the cost of implementing stationary inventory policies when demand is non-stationary. The authors compare the performance of the optimal non-stationary (s, S) reorder policy with the optimal stationary (s, S) reorder policy. The authors illustrate that the cost of ignoring the non-stationarity of the demand process is significant for the majority of the cases. Furthermore, the numerical study in Tunc et al. (2011) shows that the magnitude of the sub-optimality of stationary policies depends heavily on the variation of the demand pattern. Non-stationary inventory systems with demand distribution uncertainty are considered in Qiu, Sun, and Lim (2017). The uncertainty in demand distribution is frequently encountered in real world systems due to limited information or the short life cycle of products. The uncertainty in demand distribution is captured by box and ellipsoid uncertainty sets. Effective and tractable dynamic programming approaches are presented to solve for the reordering policies.

2.2. Inventory systems with Markovian demand

In the context of Markovian demand, Song and Zipkin (1993) present a continuous review system where the demand process follows a Markov Modulated Poisson Process (MMPP). The optimality of an (s, S) reorder policy which is dependent on the state of the MMPP is illustrated for fixed and proportional ordering costs. The work in Sethi and Cheng (1997) further extend the generality of the optimality of the state dependent (s, S) policy to include cyclical and seasonal demand for both finite and infinite non-stationary systems. Computing the optimal or near optimal state dependent (s, S) reorder policy for stationary Markovian systems systems also proved to be a computational complex problem (Beyer, Cheng, Sethi, & Taksar, 2010). The work in Nasr and Maddah (2015) centers around the computational aspect of calculating the steady state performance measures of a state dependent policy with an MMPP demand process where the main source of computational complexity is a result of solving a system of linear equations. Accounting for the time-dependent components significantly adds to the computational complexity and the tools used for stationary systems can lead to costly estimates. In contrast to the existing approaches which assume stationary Markovian demand and center around approximating the steady state value of a performance measure, the time-dependent approach presented in this work approximates the value of the performance measure over all points in time.

2.3. Reordering policies for non-stationary and stochastic demand

The model in Silver (1978) assumes time-dependent demand forecasts where the forecast errors are normally distributed. For a fixed ordering cost and proportional carrying cost, Silver (1978) develops a sequential optimization procedure to determine the optimal ordering period, when to order, and the order size.

Bookbinder and Tan (1988) extend the model presented in Silver (1978) and assume that the demand distribution is known and backorders are filled with the arrival of the next supply. The authors compare Static and Dynamic policies for different scenarios where demand forecasts can be updated with a rolling horizon approach. Recently, Lim and Wang (2017) present a target-oriented optimization approach for multi-product inventory systems with stochastic demand. The demand for each product is characterized by an uncertainty set. The authors illustrate that the target-oriented optimization approach significantly outperforms traditional optimization approaches especially under demand distribution ambiguity or inaccurate demand distributions. Tarim and Kingsman (2006) and Tarim and Smith (2008) consider non-stationary and stochastic demand and present algorithmic approaches to analyze (R^n, S^n) and (R, S) reorder policies, respectively. Inventory systems with stochastic and time-dependent demand are considered in Levi, Pl, Roundy, and Shmoys (2007). The authors introduce a new marginal accounting scheme which considers the cost of over-ordering and under-ordering when a decision is made at every period. The decision is unaffected by future decisions which simplifies the problem significantly. The work is extended to uncapacitated systems in Levi, Roundy, Shmoys, and Truong (2008). The authors address inventory systems with complex demand structure which causes the state space of dynamic programs (DPs) to explode. The authors state that it is unlikely that efficient algorithms that solve the DPs to optimality are developed. This further motivates developing computationally efficient heuristics to solve for the reorder policies.

The remainder of the paper proceeds as follows. In Section 3, we review the notation of the time-dependent Ph_t process and present a computational approach to calculate characteristics of the counting process. The mathematical model and a computational framework to solve for the time-dependent inventory measures are presented in Section 4. Heuristics to search for the reorder policy are considered in Section 5. Numerical examples are presented in Section 6 to further illustrate the contribution of the computational approaches presented in this work. A summary discussion, future recommendations and conclusions are presented in Section 7.

3. The Ph_t process

We review the notation of the time dependent Ph_t process in Section 3.1. A computational framework to obtain key characteristics of the time-dependent demand process is presented in Section 3.2.

3.1. The \boldmath $\mathrm{Ph_t}$ demand process

The Ph distribution (as defined by Neuts, 1981) belongs to the family of Markovian distributions which have received extensive attention in the literature due to their flexibility in approximating stochastic point processes. We choose the Ph distribution to model demand processes since it has the property of being dense; i.e., the family of Ph distributions can approximate any distribution shape arbitrarily closely (Asmussen, 2008). Commonly used distributions which are also special cases of the Ph distribution include the exponential, Erlang (models low variability distributions), Hyper-exponential (models high variability distributions), among other probability distributions. The work in Johnson and Taaffe (1990) illustrates that with the Ph distribution, it is possible to match the first k moments of any distribution with a finite-phase Ph distribution. Johnson and Taaffe (1989) also show that the first k moments can be represented by a finite mixture of common-ordered Erlangs. The work in Telek and Horvath (2007) shows that fitting up to the kth moments requires a Ph distribution with at least $2k-1$ phases. Obviously, accounting for a larger number of moments increases the validity and accuracy of the model, but can significantly add to the computational complexity involved in analyzing the model. We refer to Buchholz et al. (2014) for more details on the Ph distribution and existing fitting approaches. In addition to the moment matching flexibility, the Markovian representation of the Ph and its time-dependent generalization, Ph_t, have the computationally attractive properties of linear first-order differential equations. The numeric solution of the differential equations allows us to capture the time-continuous behavior of a point process.

A commonly used notation of an m-phase Ph_t distribution describes a continuous-time Markov Chain with m transient states and one absorbing state. The parameters of the time-dependent Ph_t distribution are denoted by the infinitesimal generator of the continuous-time Markov process,

$$\underset{(m+1,\,m+1)}{\mathbf{Q(t)}} = \begin{bmatrix} \underset{(m\times m)}{\boldsymbol{\lambda(t)}} & \underset{(m\times 1)}{\boldsymbol{\mu(t)}} \\ \underset{(1\times m)}{\mathbf{0}} & 0 \end{bmatrix},$$

where $\lambda_{n,\nu}(t)$ is the time-dependent transition rate from transient state n to transient state ν at time t for $n,\nu = 1,\ldots,m$ and $n \neq \nu$. The parameter $\mu_n(t)$ is the time-dependent transition rate at time t from transient state n to the absorbing state for $n = 1,\ldots,m$. It follows that the diagonal entries of $\boldsymbol{\lambda(t)}$ are calculated as $\lambda_{n,n}(t) = -\left(\sum_{\nu=1,\nu\neq n}^{m} \lambda_{n,\nu}(t) + \mu_n(t)\right)$ for $n = 1,\ldots,m$. Let $\alpha_n(t)$ denote the probability the Markov process starts in state n at time t for $n = 1,\ldots,m$.

The non-stationary inter-demand epochs are modeled by a time-dependent Ph_t process with m transients states and one absorbing state. Let $\{A(t)\colon t \geq 0\}$ denote the state of the Ph_t process at time t. The time-dependent probabilities of being in state n of the Ph_t process are calculated by numerically solving the corresponding Kolmogorov Forward Equations (KFEs). The KFEs are presented below where the derivative is with respect to t and $P(.)$ denotes Probability(.).

$$P'(A(t)=n) = \sum_{\nu=1}^{m} \left(\lambda_{\nu,n}(t) + \alpha_n(t)\,\mu_\nu(t)\right) P(A(t)=\nu) \quad \text{for } n = 1,\ldots,m. \tag{1}$$

3.2. The time-dependent demand count process

In this section, we investigate the Ph_t counting process and we present a numerical approach to calculate the probability distribution of the non-stationary count process over any time interval $[t, t+\tau]$, for $t \geq 0$ and $\tau > 0$. Specifically and for the purpose of this work, we require the probability distribution of the count process over any time interval with duration $\tau = L$, where L is the lead time. Denote the demand-count over the time interval $[t, t+\tau)$ by $\{D_t(\tau)\colon t \geq 0,\ \tau > 0\}$. At any point in time t, we numerically solve for the probability distribution of the demand count over the time interval $[t-L, t)$ if $t \geq L$, and $[0,\ t)$ if $t < L$; i.e, we numerically solve for $P(D_{t-L}(L) = d)$ if $t \geq L$, and $P(D_0(t) = d)$ if $t < L$, for $d = 0, 1, 2, \ldots, \infty$.

The state space of the demand count process along with the state of the Ph_t process constitute a continuous-time Markov Chain. At any point in time t where $t \geq L$, the resulting state space probabilities of the demand-count process over the time interval $[t-L, t)$ and the state of Ph_t process are expressed as,

$$P_{d,n}(t) = P(D_{t-L}(L) = d, A(t) = n) \text{ for } d = 0,1,2,\ldots,\infty \text{ and } n = 1,2,\ldots,m.$$

Similarly for $0 \leq t < L$,

$$P_{d,n}(t) = P(D_0(t) = d, A(t) = n) \text{ for } d = 0,1,2,\ldots,\infty$$

Algorithm 1 Time-Dependent Demand-Count algorithm.

For $t > 0$,

Step 0. Initialization: set $t_s = \max(t - L, 0)$ and $n = 1$

Step 1. (Initialize the set of HSDE in (2))

Set $\mathrm{P}(D_{t_s}(0) = 0, A(t_s) = n) = P(A(t_s) = n)$

Step 2. Numerically solve the HSDEs in (2) from $\tau = 0$ to $\tau = \min(t, L)$.

▷ Note: the solution of the HSDEs for $v = 1, \ldots, m$ and $d = 0, \ldots, d_{max}$ yields,

$\mathrm{P}(D_{t_s}(\min(t, L)) = d, A(t) = v \mid A(t_s) = n)$

Step 3. For $d = 0, 1, \ldots, d_{max}$,

$$\text{Set } \mathrm{P}(D_{t_s}(\min(t, L)) = d | A(t_s) = n) = \sum_{v=1}^{m} \mathrm{P}(D_{t_s}(\min(t, L)) = d, A(t) = v | A(t_s) = n)$$

Step 4. If $n < m$, set $n = n + 1$ and go to Step 1

and $n = 1, 2, \ldots, m$.

To solve for the joint probabilities, $P_{d,n}(t)$ for $t > 0$, we consider the corresponding KFEs over the time interval $[t_s, t_s + \tau)$ where the derivative is with respect to τ and $t_s = \max(t - L, 0)$. The KFEs for $t > 0$, $d = 0, 1, 2, \ldots, \infty$ and $n = 1, 2 \ldots, m$,

$$\begin{aligned} P'(D_{t_s}(\tau) = d, A(t_s + \tau) = n) &= \sum_{v=1}^{m} \lambda_{v,n}(t_s + \tau)\, P(D_{t_s}(\tau) = d, A(t_s + \tau) = v) \\ &+ \sum_{v=1}^{m} \alpha_n(t_s + \tau)\, \mu_v(t_s + \tau)\, P(D_{t_s}(\tau) \\ &= d - 1, A(t_s + \tau) = v)\, \mathrm{I}_{(d>0)}, \end{aligned} \tag{2}$$

where the indicator function $I_{(.)} = 1$ if the relation in (.) is satisfied and 0 otherwise.

For $t \geq L$, The probabilities $P_{d,n}(t) = P(D_{t-L}(L) = d, A(t) = n)$ are calculated by numerically integrating the homogeneous set of differential equations (HSDEs) representing the KFEs (Eq. (2)) from $\tau = 0$ to $\tau = L$. Similarly for $0 \leq t < L$, the probabilities $P_{d,n}(t) = P(D_0(t) = d, A(t) = n)$ are calculated by numerically integrating the KFEs in Eq. (2) from $\tau = 0$ to $\tau = t$. Notice that there is no upper limit on the state space of the demand counts as expressed in Eq. (2). We set an upper bound on the demand counts, d_{max}, where the $P(D_t(L) > d_{max}) < \epsilon$ for $t \in [0, T - L]$ and ϵ is a sufficiently small number. The algorithm to numerically solve for the demand-count probability distribution for a time interval $[0, T]$ is presented below (Algorithm 1). We refer to the algorithm by the Time-Dependent Demand-Count (TDDC) algorithm.

Note that the TDDC algorithm at Step 3 calculates the demand count probability distribution conditioned on the state of the Phase-type process at the start of the time interval. Specifically at Step 3, the algorithm calculates $\mathrm{P}(D_0(t) = d | A(0) = n)$ for $0 \leq t < L$ and $P(D_{t_s}(L) = d | A(t_s) = n)$, for $t \geq L$. For $d = 0, 1 \ldots, d_{max}$,

Table 1
Moments of the demand count process.

Period:	1	2	3	4
1st moment	20	40	10	50
2nd moment	490	1707	140	2600

the demand-count over lead time distribution is calculated as follows:

$$P(D_0(t) = d) = \sum_{n=1}^{m} \Big(\mathrm{P}(D_0(t) = d | A(0) = n) \times P(A(0) = n) \Big) \quad \text{for } 0 \leq t < L, \tag{3}$$

$$P(D_{t_s}(L) = d) = \sum_{n=1}^{m} \Big(P(D_{t_s}(L) = d | A(t_s) = n) \times P(A(t_s) = n) \Big) \quad \text{for } t \geq L. \tag{4}$$

Example 3.2: Consider a demand process with a time-dependent rate as presented in Fig. 1. The demand rate is assumed to be piecewise constant over four consecutive time periods of duration 10 time units. The time-dependent variability exhibited by the demand process is captured by the second moment of the counting process and is given in Table 1. Notice that a piecewise constant demand rate is presented in this example for illustrative purposes. An inventory system with a continuous-time demand rate is investigated in Section 6.

A Ph_t process which matches the demand-count moments of Table 1 is described in Fig. 2.

The parameters of the fitted Ph_t process are $\lambda_{1,2}(t) = \mu_2(t) = 4\, r(t)\, \alpha(t)$ and $\lambda_{3,4}(t) = \lambda_{4,5}(t) = \mu_5(t) = 6\,(1 - \alpha(t))\, r(t)$. The parameter $\alpha(t)$ varies by period according to Table 2.

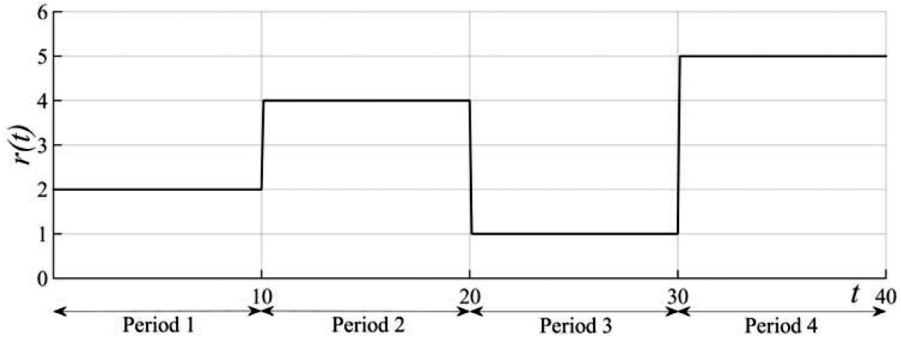

Fig. 1. Instantaneous demand rate.

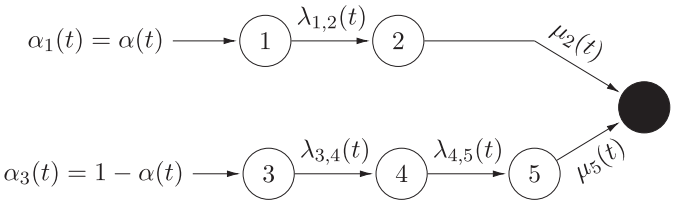

Fig. 2. Ph_t demand process – Example 3.2.

Table 2
$\alpha(t)$ – Example 3.2

Period	1	2	3	4
$\alpha(t)$	0.9339	0.8891	0.9350	0.8583

The probability distribution of the demand count process over any time interval $[\max(t-L,0),t]$ and $t \geq 0$, is calculated by executing the TDDC-algorithm. The mean and standard deviation of the resulting count process are presented in Fig. 3.

The existing algorithmic approaches to calculate the stationary moments of the counting process assume constant coefficients of the Markovian distribution and do not account for time-dependency. The work in Narayana and Neuts (1992) presents approaches to calculate the first two moments of the counting process of Markovian distributions. The computational method is based on the uniformization approach and assumes constant coefficients. Approaches to calculate the kth moment of the counting process are presented in Nielsen, Nilsson, Thygesen, and Beyer (2007). The work in Nasr, Charanek, and Maddah (2017) also presents a compact exponential equation to calculate the moments of the counting process of Markovian processes. We are not aware of any existing computational techniques to compute the probability distribution of the time-dependent counting process without resorting to numerical integration.

4. Model

In this section, we present the state space of the Markovian representation of the inventory system and the correspond KFEs. We also present the equations to solve for the time-dependent moments of the inventory measures. Consider a single item, finite horizon, and continuous review inventory system where the inter-demand epochs follow a time-dependent Ph_t process. The reorder policy is a time-dependent (s_t, S_t) policy. The model assumes a deterministic lead time, $L \geq 0$, where shortages are allowed and back-ordered. We assume a fixed ordering cost (ω \$ per order) and proportional holding (h \$ per item per unit time) and shortage costs (b \$ per item per unit time).

4.1. Inventory position

Denote the state of the inventory position at time t by $\{\text{IP}(t): t \geq 0\}$. The lowest achievable IP(t) level for $t \in [0, T]$ is $\ell = \min_{t\in[0,T]}(s_t) + 1$ and the upper limit on IP(t) is $u = \max_{t\in[0,T]}(S_t)$. Augmenting the states of the Ph_t demand process, $\{A(t): t \geq 0\}$, with the state of the inventory position, $\{\text{IP}(t): t \geq 0\}$, at time t, results in a Markovian process with state probabilities,

$$P_{i,n}(t) = P(\text{IP}(t) = i, A(t) = n)), \text{ for } i = \ell, \ldots, u \text{ and } n = 1, \ldots, m.$$

The corresponding KFEs where the derivative is with respect to t for $n = 1, \ldots, m$ and $i > s_t$,

$$P'_{i,n}(t) = \sum_{v=1}^{m} \lambda_{v,n}(t)\, P_{i,v}(t) + \alpha_n(t)\Big(\sum_{v=1}^{m} \mu_v(t)\, P_{i+1,v}(t)\Big) I_{(i\neq u)} + \alpha_n(t)\Big(\sum_{v=1}^{m}\sum_{j=\ell}^{s_t+1} \mu_v(t)\, P_{j,v}(t)\Big) I_{(i=S_t)}, \quad (5)$$

where the indicator function $I_{(.)} = 1$ if the relation in (.) is satisfied and 0 otherwise. For $t \geq 0$, $n = 1, \ldots, m$ and $i \leq s_t$,

$$P'_{i,n}(t) = \sum_{v=1}^{m} \lambda_{v,n}(t)\, P_{i,v}(t). \quad (6)$$

Notice that Eqs. (5) and (6) assume that a replenishment order is placed when a demand epoch is realized and the inventory position is less than or equal to s_t.[1] The KFEs denoted by Eqs. (5) and (6) are initialized by assigning values for the joint inventory position and the state of the Ph_t process at $t = 0$. For example, initializing the inventory position to a value of i_o can be achieved by setting $\sum_{n=1}^{m} P_{i_o,n} = 1$ and $P_{i,n} = 0$ for $i \neq i_o$, $n = 1, \ldots, m$. The joint probability, $P_{i,j}(t) = P(\text{IP}(t) = i, A(t) = j)$, is calculated by numerically integrating the set of KFEs in Eqs. (5) and (6) over the time interval $[0, T]$. The probability distribution of the time-dependent inventory position for $t \in [0, T]$ is calculated as,

$$P(\text{IP}(t) = i) = \sum_{n=1}^{m} P_{i,n}(t) \text{ for } i = \ell, \ldots, u. \quad (7)$$

The time-dependent kth moment of the inventory position for $t \in [0, T]$,

$$E[\text{IP}^k(t)] = \sum_{i=\ell}^{u} i^k\, P(\text{IP}(t) = i). \quad (8)$$

[1] We refer to Appendix C.1 for an investigation of the behavior of the inventory system when an order is placed as soon as the inventory position is less than s_t.

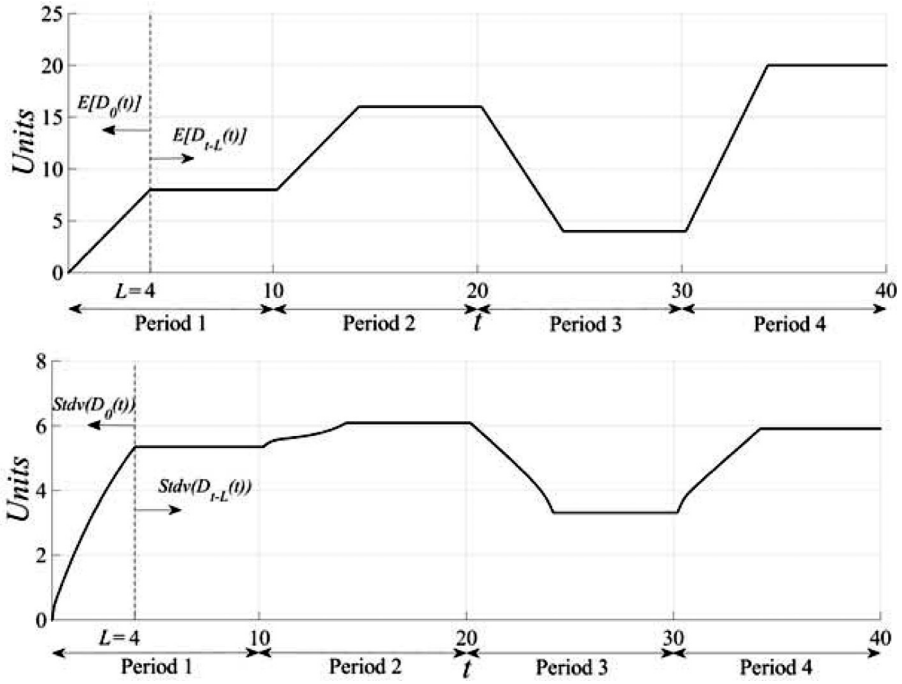

Fig. 3. Mean and standard deviation of the demand over lead time.

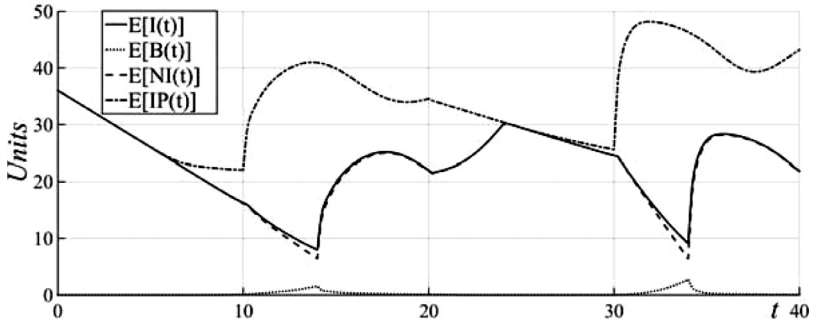

Fig. 4. Expected inventory measures over time.

The TDDC algorithm calculates the time-dependent probability distribution of the demand count process and, similarly, the time-dependent probability distribution of the inventory position is calculated by numerically solving Eqs. (5) and (6). This is illustrated by considering the Ph_t demand process of Example 3.2 for a given dynamic reorder policy, $(\bar{s}, \bar{S})$ where $\bar{s} = \{8, 16, 4, 20\}$ and $\bar{S} = \{36, 54, 24, 65\}$ over the four time periods. Utilizing Eq. (8) for $k = 1, 2$, we obtain the plots for the mean and standard deviation of the inventory position as shown in Figs. 4 and 5

4.2. Net inventory

In this section, we build on the equations of the inventory position to solve for the probability distributions of the net inventory. Let $\{\mathrm{NI}(t)\colon t \geq 0\}$ be the time-dependent net inventory. For a fixed lead time, L, the net inventory at time $t \geq 0$ is expressed in terms of the demand counting process and the inventory position as follows:

$$\mathrm{NI}(t) = \mathrm{IP}(0) - D_0(t) \qquad \text{for } 0 \leq t \leq L, \tag{9}$$

$$\text{and} \quad \mathrm{NI}(t) = \mathrm{IP}(t-L) - D_{t-L}(L) \qquad \text{for } t > L. \tag{10}$$

By conditioning on the state of the Ph_t process at time t, we obtain the conditional probability distribution of the net inventory for $0 \leq t \leq L$ and $i = \ell - d_{max}, \ldots, u$,

$$P(\mathrm{NI}(t) = i | A(0) = n) = \sum_{j=\ell}^{u} \Big(P(\mathrm{IP}(0) = j | A(0) = n) \times P(D_0(t) = j - i | A(0) = n) \Big), \tag{11}$$

$$\Rightarrow P(\mathrm{NI}(t) = i, A(0) = n) = \sum_{j=\ell}^{u} \Big(P(\mathrm{IP}(0) = j, A(0) = n) \times P(D_0(t) = j - i | A(0) = n) \Big). \tag{12}$$

Similarly for $t > L$ and $i = \ell - d_{max}, \ldots, u$,

$$P(\mathrm{NI}(t) = i | A(t-L) = n) = \sum_{j=\ell}^{u} \Big(P(\mathrm{IP}(t-L) = j | A(t-L) = n) \times P(D_{t-L}(L) = j - i | A(t-L) = n) \Big), \tag{13}$$

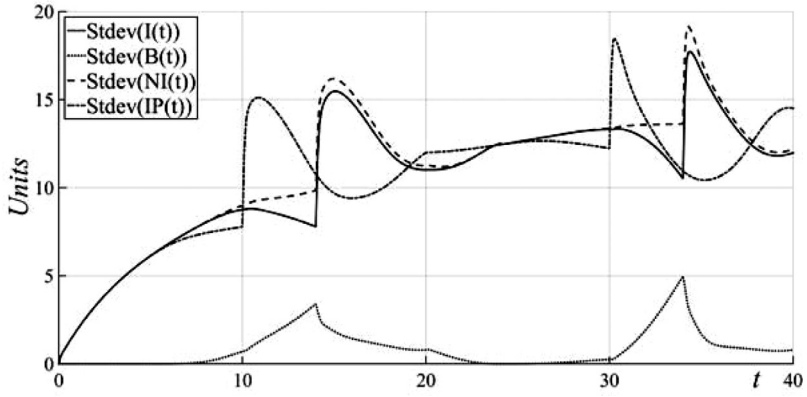

Fig. 5. Standard deviation of inventory measures over time.

$$\Rightarrow P(\mathrm{NI}(t)=i, A(t-L)=n) = \sum_{j=\ell}^{u}\Big(P(\mathrm{IP}(t-L)=j, A(t-L)=n) \times P(D_{t-L}(L)=j-i \mid A(t-L)=n)\Big). \quad (14)$$

Utilizing the joint probability distributions in Eqs. (12) and (14), the time-dependent probability distribution of the net inventory is calculated for $i=\ell-d_{max},\ldots,u$,

$$P(\mathrm{NI}(t)=i)=\sum_{n=1}^{m}P(\mathrm{NI}(t)=i, A(0)=n) \quad \text{for} \quad 0\le t\le L, \quad (15)$$

$$\text{and} \quad P(\mathrm{NI}(t)=i)=\sum_{n=1}^{m}P(\mathrm{NI}(t)=i, A(t-L)=n) \quad \text{for} \quad t>L. \quad (16)$$

The time-dependent kth moment of the net inventory position at time $t\ge 0$,

$$\mathrm{E}[\mathrm{NI}^k(t)]=\sum_{i=\ell-d_{max}}^{u} i^k\, P(\mathrm{NI}(t)=i). \quad (17)$$

4.3. Performance measures – on hand and backordered inventory

The performance measures we consider in this section are the on-hand inventory, backordered inventory, and the shortage probability. These performance measures can be directly derived from the distribution of the net inventory. Let $\{I(t): t\ge 0\}$ be the on hand inventory which is denoted by the positive component of the net inventory. Accordingly, $I(t)=\mathrm{NI}(t)^+$ where $\mathrm{NI}(t)^+=\max\{\mathrm{NI}(t),0\}$. The probability distribution of the net inventory for $t\ge 0$,

$$P(I(t)=i)=\begin{cases} P(\mathrm{NI}(t)=i) & \text{for } i=1,\ldots,u,\\ \sum_{j=\ell-d_{max}}^{0} P(\mathrm{NI}(t)=j) & \text{for } i=0.\end{cases} \quad (18)$$

Similarly, the number backordered at time $t\ge 0$, $\{\mathrm{B}(t): t\ge 0\}$, is the negative component of the net inventory. Let $\mathrm{B}(t)=\mathrm{NI}(t)^-$ where $\mathrm{NI}(t)^-=\max\{-\mathrm{NI}(t),0\}$. For $t\ge 0$,

$$P(\mathrm{B}(t)=i)=\begin{cases} P(\mathrm{NI}(t)=-i) & \text{for } i=1,\ldots,(d_{max}-\ell),\\ \sum_{j=0}^{u} P(\mathrm{NI}(t)=j) & \text{for } i=0.\end{cases} \quad (19)$$

The service level is another system performance measure which is directly calculated from the probability distribution of the net inventory. Let $\{\mathrm{S}(t): t\ge 0\}$ be the time-dependent service level which represents the probability that there are no shortages at time t. The time-dependent service level is calculated as,

$$\mathrm{S}(t)=P(\mathrm{NI}(t)>0)=1-P(\mathrm{B}(t)=0). \quad (20)$$

Consider the demand process of Example 3.2 and the reorder policy $(\bar{s},\bar{S})=(\{8,16,4,20\},\{36,54,24,65\})$. The mean and standard deviation of the on-hand inventory levels $I(t)$ and backorders $\mathrm{B}(t)$, are shown below in Figs. 4 and 5.

4.4. Ordering count process

A performance measure that is not directly calculated from the net inventory probability distributions is the number of orders placed. We present the computational framework to calculate the probability distribution of the ordering count process. The ordering process is accounted for as a system performance measure since the model assumes a fixed ordering cost. The ordering process is obviously a counting process and in this section we consider the distribution and moments of the number of orders placed over a time interval. Define $\{\mathrm{R}(t): t\ge 0\}$ as the number of orders by time t; i.e., the number of orders over the time interval $[0, t)$. For a given reorder policy, the number of orders placed by time t is dependent on the state of the demand process and the inventory position. Accordingly, a Markovian representation of the counting process at time t includes the state of the Ph_t demand process, $\{A(t): t\ge 0\}$, and the state of the inventory position, $\{\mathrm{IP}(t): t\ge 0\}$. The state probabilities of the resulting Markov representation is presented below,

$$P_{r,i,n}=P(\mathrm{R}(t)=r,\ \mathrm{IP}(t)=i,\ A(t)=n), \quad \text{for } r\ge 0,\ i=\ell,\ldots,u \text{ and } n=1,\ldots,m.$$

The corresponding KFEs where the derivative is with respect to t for $n=1,\ldots,m$, $s_t<i\le u$, and $r>0$,

$$P'_{r,i,n}(t)=\sum_{\nu=1}^{m}\lambda_{\nu,n}(t)\,P_{r,i,\nu}(t)+\alpha_n(t)\Big(\sum_{\nu=1}^{m}\mu_\nu(t)\,P_{r,i+1,\nu}(t)\Big)I_{(i\ne u)} +\alpha_n(t)\Big(\sum_{\nu=1}^{m}\sum_{j=\ell}^{s_t+1}\mu_\nu(t)\,P_{r-1,j,\nu}(t)\Big)I_{(i=S_t)}\,I_{(r>0)}. \quad (21)$$

For $n=1,\ldots,m$, $\ell\le i\le s_t$ and $r\ge 0$,

$$P'_{r,i,n}(t) = \sum_{\nu=1}^{m}\lambda_{\nu,n}(t)\,P_{r,i,\nu}(t). \quad (22)$$

A computationally extensive approach to obtain the probabilities $P_{r,i,n}$ for $r\ge 0$, $i=\ell,\ldots,u$ and $n=1,\ldots,m$, is to numerically solve the differential equations in (21) and (22). Such

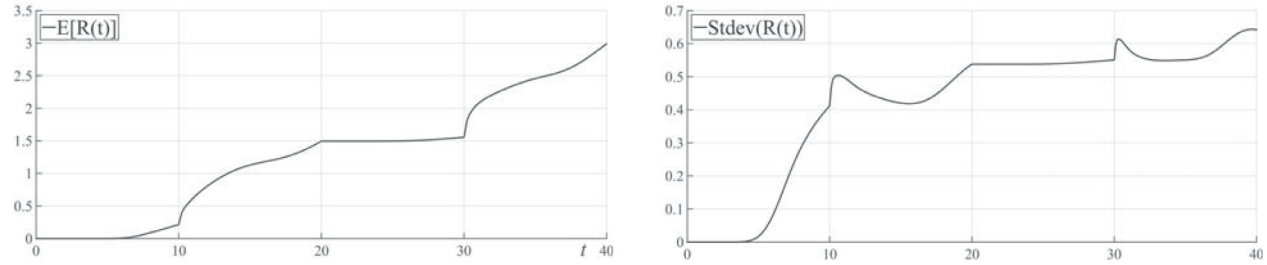

(a) Mean (b) Standard deviation

Fig. 6. $R(t)$ mean and standard deviation plots.

an approach requires truncating the infinite set of differential equations in (21) and (22) by setting an upper limit on the number of orders in a time interval. For example, an upper limit on the number of orders is r_{max} where $P(\mathrm{R}(T) > r_{max})$ is sufficiently small. A more efficient approach to capture the behavior of the ordering count process is to derive a finite set of moment differential equations to numerically calculate the moments of $\mathrm{R}(t)$ for $t \geq 0$. Next, we present the finite set of moment differential equations and the computational framework required to calculate the kth moment of the reorder process.

For $n = 1, \ldots, m$, $s_t < i \leq u$ and $k \geq 1$, denote the joint moments of the number of orders by time $t \geq 0$ as follows:

$$E[R^k(t); i, n] = E[R^k(t); \mathrm{IP}(t) = i, A(t) = n] = \sum_{r=0}^{\infty} r^k P_{r,i,n}(t). \quad (23)$$

Let $E'[R^k(t); i, n]$ be the derivative, with respect to t, of the moments presented in Eq. (23). The differential equations to calculate the kth moment of the number of orders in a time interval are derived from Eqs. (21)–(23) for $n = 1, \ldots, m$ and $s_t < i \leq u$,

$$\begin{aligned} E'[R^k(t); i, n] &= \sum_{r=0}^{\infty} r^k P'_{r,i,n}(t) \\ &= \sum_{\nu=1}^{m} \lambda_{\nu,n}(t) E[R^k(t); i, \nu] + \alpha_n(t) \Big(\sum_{\nu=1}^{m} \mu_\nu(t) E[R^k(t); i+1, \nu] \Big) I_{(i \neq u)} \\ &\quad + \alpha_n(t) \Big(\sum_{\nu=1}^{m} \sum_{j=\ell}^{s_t+1} \mu_\nu(t) \Big(\sum_{z=1}^{k} \binom{k}{z} E[R^z(t); j, \nu] + P_{j,\nu}(t) \Big) \Big) I_{(i=S_t)}, \end{aligned} \quad (24)$$

where

$$\binom{k}{z} = \frac{k!}{z!\,(k-z)!}.$$

For $n = 1, \ldots, m$, and $\ell \leq i \leq s_t$,

$$E'[R^k(t); i, n] = \sum_{\nu=1}^{m} \lambda_{\nu,n}(t)\, E[R^k(t); i, \nu]. \quad (25)$$

Setting $k = 1, 2$ in Eqs. (24) and (25), the first and second moment differential equations are obtained for $n = 1, \ldots, m$ and $s_t < i \leq u$. The mean and standard deviation of the number of orders is found by solving Eqs. (24) and (25) for $k = 1, 2$. As an illustrative example we consider the demand process of Example 3.2 and the reorder policy $(\bar{s}, \bar{S}) = (\{8, 16, 4, 20\}, \{36, 54, 24, 65\})$. The resulting mean and standard deviation of the ordering process over $t \in [0, 40]$ are plotted in Fig. 6.

The inventory performance measures of the examples considered in this paper are calculated by numerically integrating the corresponding HSDE. We refer to Appendix C.2 for a numerical illustration of the computational advantages of numerically solving the HSDEs over running Monte-Carlo simulations. Another approach to solve the HSDEs is to calculate the matrix exponential solution at discrete points in time. Accordingly, the system parameters are assumed constant within the intervals separating the discrete points in time. To achieve accurate approximations, the distance between the discrete points in time should be small enough to capture the fluctuation of the time-dependent parameters. We do not describe the matrix exponential approaches further since they can require significantly more computation for time-dependent systems. A more computationally practical alternative is to numerically solve the HSDE. We refer to Moler and Van Loan (2003) for approaches to solve HSDE, with constant coefficients, using matrix exponentials.

4.5. Expected cost function

The cost parameters are the holding cost, h ($ per item per unit time), the backordering cost, b ($ per item per unit time), and a fixed ordering cost ω, ($ per order). The expected cost by time T for a time dependent policy (s_t, S_t) is calculated as follows:

$$\Phi_{s_t,S_t}(T) = h \int_0^T E[I(t)]\, dt + b \int_0^T E[\mathrm{B}(t)]\, dt + \omega E[R(T)]. \quad (26)$$

In many real world settings, it is difficult to quantify a value for b. In such a case, the stockout probability is commonly used as a service level constraint. A byproduct of the computational model presented in this work is the calculation of the time-dependent probability distribution of the number backordered. Next, we assume a time-dependent but piecewise constant reorder policy.

5. Heuristic policies

In this section, we introduce heuristics to solve to for the reordering policies. We compare the performance of the heuristics in terms of reducing the expected cost and the corresponding computational efficiency. We also note that finding the optimal solution using dynamic programming (DP) approaches is computationally impractical. We refer to Levi et al. (2008) for a discussion on inventory systems with complex demand structures that cause the state space of dynamic programs to explode. The authors argue that developing efficient algorithms to solve the corresponding DPs to optimality is unlikely unless strong assumptions on the state space and underlying Markov process are considered. The computational complexity extends to finding reasonable lower bounds on the optimal cost. This motivates developing heuristics to solve for cost-efficient reordering policies.

We consider heuristics to minimize the expected cost over a time interval of duration T, $\Phi_{\bar{s},\bar{S}}(T)$, as expressed in Eq. (26). Notice that although we solve for the reorder policy over a fixed time interval T, the heuristics can be utilized in the context of a rolling horizon where the reorder policy is updated at any point in time t, $0 < t \leq T$. This would account for an updated demand forecast as well as the realization of the state of the inventory measures at

time t. The heuristics assume a piecewise constant reorder policy where the reorder levels change over specified time periods represented by consecutive and non-overlapping time intervals with duration δ. If the number of periods is N, then a reorder policy is denoted by $(\bar{\mathbf{s}}, \bar{\mathbf{S}})$ where $\bar{\mathbf{s}} = \{s_1, s_2, \ldots, s_N\}$ and $\bar{\mathbf{S}} = \{S_1, S_2, \ldots, S_N\}$.

5.1. Stationary approximation heuristic

The proposed heuristic solves for a stationary ordering policy at the start of each time period. Solving for reorder policies of the stationary problem for different instances of the demand parameters is considered in Bollapragada and Morton (1999). The resulting stationary reordering policies are tabulated for the different demand parameters. The non-stationary problem is solved by looking up the stationary reorder policies from the generated table. We utilize a similar approach where stationary reorder policies are considered for every time period. The stationary reorder policies for the demand instances are calculated by assuming the demand process over the lead time follows a Normal distribution. We refer to the heuristic by the stationary approximation (SA) heuristic.

Let $(\bar{\mathbf{s}}', \bar{\mathbf{S}}') = \{s'_1, s'_2, \ldots, s'_N\}, \{S'_1, S'_2, \ldots, S'_N\}$ represent the stationary reorder policies for the N time periods. For $\theta = 1, \ldots, N$,

$$s'_\theta = E[D_{\delta(\theta-1)}(L)] + z_\theta \,\mathrm{Stdv}(D_{\delta(\theta-1)}(L)), \tag{27}$$

where $D_{\delta(\theta-1)}(L)$ is the demand count over the time interval $[\delta(n-1), \delta(\theta-1)+L)$ and z_θ is the solution to the equation,

$$\frac{Q_\theta}{\mathrm{Stdv}(D_{\delta(\theta-1)}(L))}\left(\frac{h}{b+h}\right) = \int_{z_\theta}^{\infty}(u - z_\theta)\frac{1}{\sqrt{2\pi}}e^{-u^2/2}du, \tag{28}$$

where Q_θ is the order quantity based on the Economic Order Quantity (EOQ) model,

$$Q_\theta = \sqrt{\frac{2\,\omega}{h}\left(\frac{E[D_{\delta(\theta-1)}(L)]}{L}\right)}, \text{ for } \theta = 1, \ldots, N.$$

The equality in Eq. (28) assumes stationary demand and is based on the normality assumption of the demand counting process over the lead time. We refer to Silver, Pyke, and Peterson (1998) for more information on Eq. (28). The reorder up to level for $\theta = 1, \ldots, N$,

$$S'_\theta = s'_\theta + Q_\theta. \tag{29}$$

5.2. Line search heuristic

The algorithm is initiated with the reorder policy obtained from Eqs. (27) and (29) of the stationary approximation heuristic. The heuristic proceeds with a line search on the value of s_1 while fixing the remaining parameters. A line search is then performed on the resulting order policy where S_1 is varied and the remaining parameters remain unchanged. Similarly and for $n = 2, \ldots, N$, the heuristic proceeds with a line search on s_n then S_n. This process is repeated until the heuristic converges to a local optimum. A detailed representation of the line search (LS) heuristic is presented in Algorithm 2. The LS heuristic is initialized by the SA Heuristic as described in Step 0 of Algorithm 2.

Obviously, the LS heuristic does not guarantee global optimality but converges at a local optimal. The main source of computational complexity is solving the HSDE in Eqs. (5) and (6) to calculate the probability distribution of inventory position at time t as well as the HSDE in Eqs. (24) and (25) for $k = 1$ to calculate the expected number of orders by time t. The number of HSDE generated by Eqs. (5) and (6) is $(u - \ell + 1)\,m$. Similarly, the number of HSDE in Eqs. (24) and (25) for $k = 1$ is $(u - \ell + 1)\,m$. Accordingly, the computational complexity is a function of the number of states in the Ph_t process, m, and the range on the possible realizations of the inventory position, $u - \ell + 1$. In this work, MATLAB is utilized to solve the HSDE where the numerical solution is based on the Runge–Kutta formulae as described in Dormand and Prince (1980).

6. Numerical examples

In this section, we numerically examine the computational approaches and heuristics. We begin by illustrating the flexibility of the Ph_t process in matching the time-dependent characteristics of a point process which also serves as a motivation for modeling demand by a Ph_t process. In Section 6.1 we utilize the approach presented in Nasr and Taaffe (2013) to fit the time dependent characteristics of the demand process to a Ph_t distribution. A base demand process is utilized which is assumed to exhibit a continuous-time rate, $r(t)$. We implement the SA heuristic in Section 6.2 where the results obtained are utilized to initialize the LS algorithm in Section 6.3. We also summarize the solution of the heuristics for different cost parameters and we illustrate the implementation of the time-dependent service level constraint. In Section 6.4, we test the effectiveness of the Ph_t fitting distribution in providing accurate approximations for the time-dependent demand process.

6.1. The demand process and fitted $\mathrm{Ph_t}$ distribution

Consider a demand process exhibiting the following instantaneous demand rate $r(t)$ over an interval of 40 time units (Fig. 7),

$$r(t) = 1 + \frac{t}{10} + 0.75\,\sin(0.2\,\pi\,t). \tag{30}$$

The variability of the demand process is captured by the second moment of the counting process over consecutive and non-overlapping intervals of duration 10 time units denoted by periods 1, 2, 3 and 4 (Table 4).

Notice that fitting the variability in the counting process captures the variability in the duration of the inter-demand epochs as well as the auto-correlation (see Gusella, 1991; Nelson & Gerhardt, 2011; Whitt, 1982). This serves as a motivation behind selecting the variability in the counting process as a key characteristic to match when calculating the parameters of the approximating Ph_t process. We generate a Ph_t to match $r(t)$ and the second moment as given in Eq. (30) and Table 3, respectively. The fitted Ph_t is a balanced mixture of Erlangs with parameters presented below and illustrated in Fig. 8. The parameters m_1 and m_2 denote the number of phases in the fitted Ph_t distribution and are set to $m_1 = 2$ and $m_2 = 3$. We refer to Nasr and Taaffe (2013) for more information on the fitting process.

$$\mu_i(t) = \begin{cases} 2\,\alpha(t)\,r(t)\,m_1, & \text{if } i = m_1 \\ 2\,(1-\alpha(t))\,r(t)\,m_2, & \text{if } i = m_1 + m_2 \\ 0, & \text{otherwise} \end{cases}$$

$$\lambda_{i,i+1}(t) = \begin{cases} 2\,\alpha(t)\,r(t)\,m_1, & \text{if } 1 \le i < m_1 \\ 2\,(1-\alpha(t))\,r(t)\,m_2, & \text{if } m_1 < i < m_1 + m_2 \end{cases}$$

The probability distribution of the demand count over the lead time, $D_{t-L}(L)$ for $t \ge L$ and $D_0(t)$ for $t < L$, is obtained by implementing the TDDC Algorithm as described in Section 3.2. Accordingly, the mean and standard deviation of the demand counting process are plotted in Fig. 9.

6.2. Numerical example – stationary approximation heuristic

Consider a base case inventory system with the following cost parameters. A fixed ordering cost $\omega = \$80$/order, a back-ordering cost $b = \$3$/unit/unit time, and holding cost $h = \$1$/unit/unit time. The demand process is assumed to follow the fitted Ph_t distribution with parameters calculated in Section 6.1. The SA heuristic calculates the re-order policy for each time period (Eqs. (27) and (29)) as follows for $\theta = 1, 2, 3, 4$,

Algorithm 2 Line search heuristic.

Step 0. Initialization:

▷ Load the probability distribution of $D_{\max(t-L,0)}(t)$ for $t \geq 0$ - from TDDC algorithm

▷ Set $(\bar{\mathbf{s}}', \bar{\mathbf{S}}') = (\{s'_1, s'_2, \ldots, s'_N, \}, \{S'_1, S'_2, \ldots, S'_N\})$ - Equations (27) and (29)

▷ Initialize $(\bar{\mathbf{s}}, \bar{\mathbf{S}}) = (\bar{\mathbf{s}}', \bar{\mathbf{S}}')$

Step 1. Set $\theta = 1$

Step 2. Set $s_\theta = \underset{s_\theta}{\operatorname{argmin}} \Phi(\{s_1, s_2, \ldots, s_\theta, s'_{\theta+1}, \ldots, s'_N\}, \bar{\mathbf{S}})(T)$

Step 3. Set $S_\theta = \underset{S_\theta}{\operatorname{argmin}} \Phi(\bar{\mathbf{s}}, \{S_1, S_2, \ldots, S_\theta, S'_{\theta+1}, \ldots, S'_N\})(T)$

Step 4. If $\theta < N$, set $\theta = \theta + 1$ and go to Step 2

Step 5. If $(\bar{\mathbf{s}}, \bar{\mathbf{S}}) = (\bar{\mathbf{s}}', \bar{\mathbf{S}}')$ then set $(\bar{\mathbf{s}}^*, \bar{\mathbf{S}}^*) = (\bar{\mathbf{s}}, \bar{\mathbf{S}})$ and terminate

else set $(\bar{\mathbf{s}}', \bar{\mathbf{S}}') = (\bar{\mathbf{s}}, \bar{\mathbf{S}})$ and go to Step 1

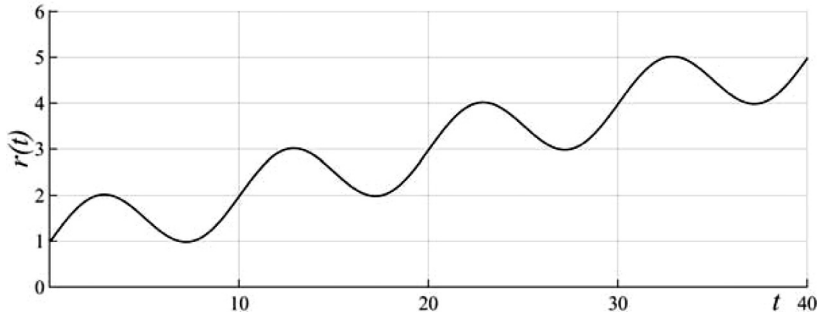

Fig. 7. Instantaneous demand rate.

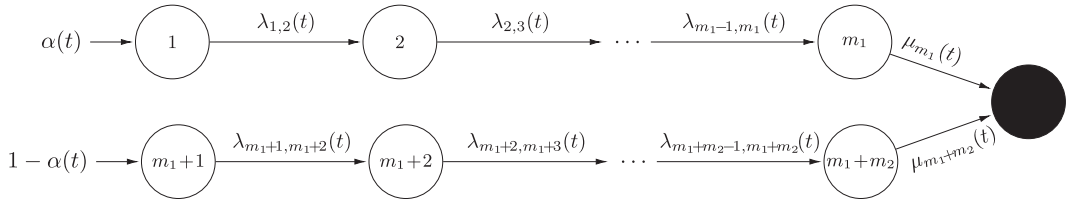

Fig. 8. Balanced 2-MECO – state numbers and phase transition rates.

Table 3
Second moments of the demand count process.

Period	1	2	3	4
2nd moment	240	650	1260	2070

Table 4
$\alpha(t)$.

Period	1	2	3	4
$\alpha(t)$	0.7637	0.7621	0.7614	0.7611

$$\left.\begin{array}{l} s'_\theta = E[D_{10(\theta-1)}(L)] + z_\theta \,\mathrm{Stdv}(D_{10(\theta-1)}(L)) \\ S'_\theta = s'_\theta + Q_\theta \end{array}\right\}$$

$$\rightarrow \begin{cases} \bar{\mathbf{s}}' = \{7,\ 11,\ 15,\ 19\} \\ \bar{\mathbf{S}}' = \{23,\ 31,\ 39,\ 46\}. \end{cases}$$

Implementing $(\bar{\mathbf{s}}', \bar{\mathbf{S}}') = (\{7,\ 11,\ 15,\ 19\}, \{23,\ 31,\ 39,\ 46\})$ as the reorder policy across the four time periods, we obtain the time-dependent probability distribution of the inventory measures.

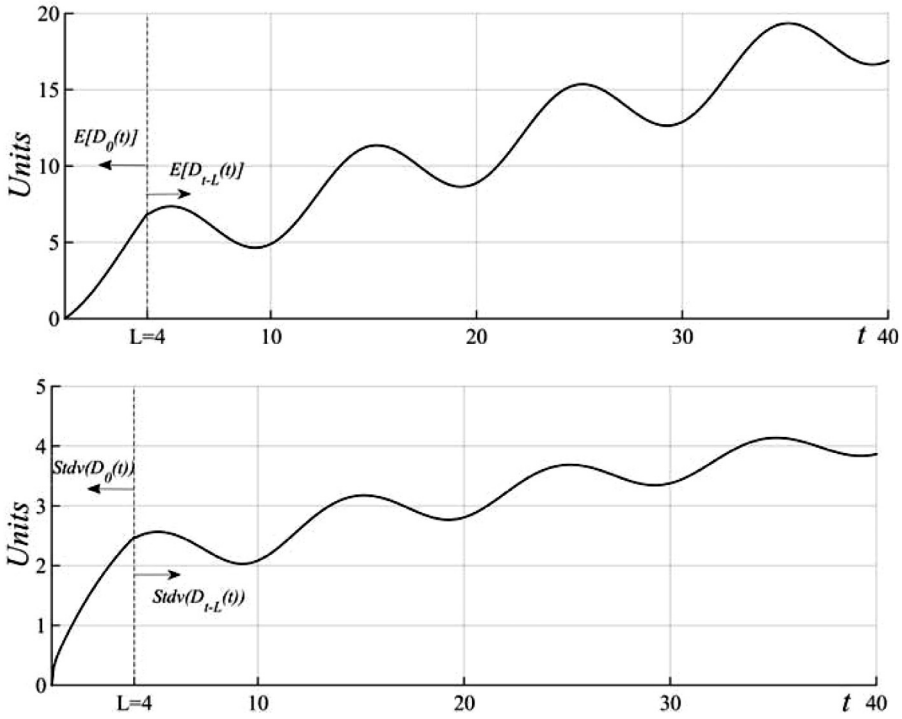

Fig. 9. Mean and standard deviation of the demand over lead time.

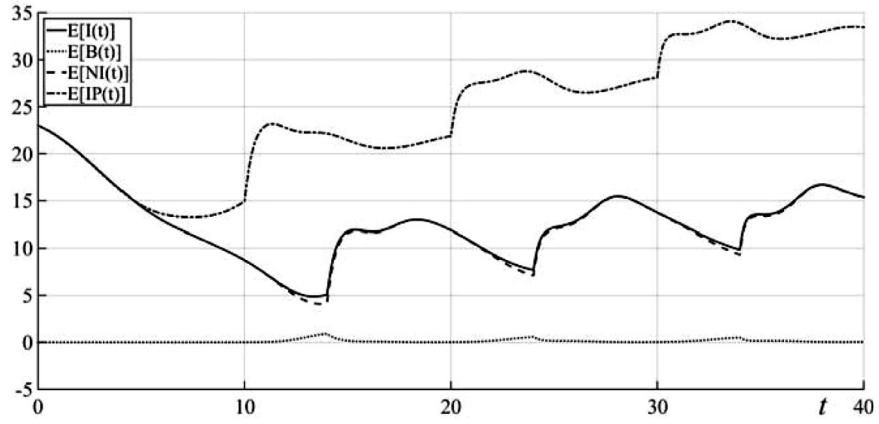

Fig. 10. Expected inventory measures – $(\tilde{\mathbf{s}}', \tilde{\mathbf{S}}')$ policy.

The time-dependent mean of I(t), B(t), NI(t) and IP(t) are calculated from the corresponding probability distributions and plotted in Fig. 10. A detailed plot of the mean and standard deviation of each inventory measure is presented in Appendix A, Fig. 14.

The total cost of implementing the $(\tilde{\mathbf{s}}', \tilde{\mathbf{S}}')$ policy by time $T = 40$ is \$960.35. A detailed plot of the evolution of the cost measures is presented in Fig. 11. As shown in Fig. 11, the holding and ordering costs are the main contributors to the total cost. The reordering policy $(\tilde{\mathbf{s}}', \tilde{\mathbf{S}}')$ results in a very low expected number of backorders as can be seen from Figs. 10 and 11.

6.3. Numerical example – line search heuristic and service level constraint

In this section, we improve on the reorder policy by running the line search heuristic. We also perform sensitivity analysis on the probability of encountering shortages by adjusting the reorder levels. We consider the base case cost parameters and initialize the LS heuristic with $(\tilde{\mathbf{s}}', \tilde{\mathbf{S}}')$. The resulting reorder policy converges to $(\tilde{\mathbf{s}}^*, \tilde{\mathbf{S}}^*) = (\{2, 4, 6, 5\}, \{18, 29, 38, 43\})$. The time-dependent expected costs are presented in Fig. 12 where the total cost by time $T = 40$ is \$806.33. The reorder policy $(\tilde{\mathbf{s}}^*, \tilde{\mathbf{S}}^*)$ improves on the initial policy (saving of 16% in total cost) but results in an increase in the expected number of backorders.

An output of the computational approach is the probability distribution of the number of backorders. Accordingly, we can closely observe the probability of shortages by evaluating $P(\mathrm{B}(t) > 0)$. Plot 1 of Fig. 13 shows the probability of encountering backorders if the initial policy is implemented. From a managerial perspective, Plot 1 clearly illustrates that the points in time in the neighborhood of $t = 14$, 24 and 34 are at the highest risk of experiencing stockouts. Notice that the orders arriving during the time intervals [10,14), [20,24) and [30,34) are placed in the previous time period under a different reorder policy. This explains the increasing behavior of the shortage probability in Plot 1 of Fig. 13 over the time intervals

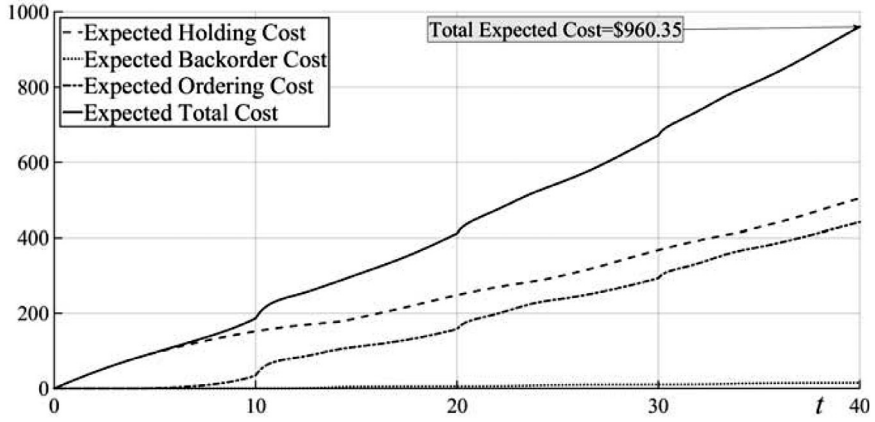

Fig. 11. Expected costs – $(\bar{\mathbf{s}}', \bar{\mathbf{S}}')$ policy.

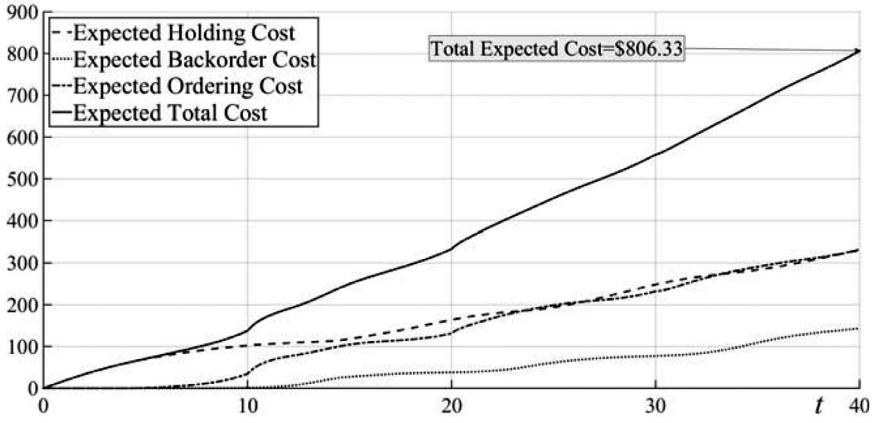

Fig. 12. Expected costs – $(\bar{\mathbf{s}}^*, \bar{\mathbf{S}}^*)$ policy.

[10,14), [20,24) and [30,34). Consequently, the points in time 14, 24 and 34 represent the earliest possible arrival of an order subject to the updated reorder policy.

The reorder policy $(\bar{\mathbf{s}}^*, \bar{\mathbf{S}}^*)$ results in a lower expected total cost but increases the stockout probabilities, Plot 2 of Fig. 13. In the case where the value of b is not quantifiable, a time-dependent service level constraint, as expressed in Eq. (20), can be implemented. Plot 3 of Fig. 13 illustrates the implementation of the time-dependent service level constraint for $S(t) \geq 0.9$ and $t \in [0, 40]$. The mean and standard deviation of the inventory measures, I(t), B(t), IP(t) and R(t), are plotted in Appendix A, Figs. 15 (without a service level constraint) and 16 (with a service level constraint). As expected, implementing the service level constraint significantly increases the on-hand inventory and reduces the number backordered as illustrated in Fig. 16. Implementing the service level constraint and accounting for the backordering cost, b, results in an increase of 90% in total cost (increase from \$806.33 to \$1532.97). We conduct the line search heuristic for a set of cost parameters where we vary $\omega = 20, 30, 40, 60, 80, 100$, $b = 3, 4, 5, 10$ and $h = 0.5, 0.75, 1$, which results in a total of 72 cases. The numerical examples are performed on an Intel Core i7 Processor with a processor speed of 3.5 giga hertz. The average time to solve the LS heuristic over the 72 cases is 25.10 seconds with a standard deviation of 10.66 seconds. The 72 cases of the LS heuristic results in 6051 instances of solving for the expected cost, $\Phi_{s_t, S_t}(T)$ of Eq. (26). The average solution time over the 6051 instances is 0.37 seconds with a standard deviation of 0.06 seconds. Since the SA heuristic involves solving a single reordering policy (Eqs. (27) and (29)), the computational time of the SA heuristic is significantly lower than the LS heuristic.

The results of the LS, SA and Static heuristics are reported in detail in Table 11 of Appendix B. The average cost of implementing the LS, SA and Static reorder policies over the 72 cases is \$628.30, \$684.80 and \$719.26, respectively. The average percentage difference between the SA and LS heuristic is 7.85% with a minimum of 1.63% and a maximum of 16.55%. The tabulated results show that the percentage difference between the SA and LS heuristic is most significant for high holding cost and low backordering cost. The average cost for $h = 1$ and $b = 3$ is 15.60% whereas the average cost when $h = 0.5$ and $b = 10$ is 2.25%. Determining the reorder policy using the SA heuristic requires solving Eqs. (27) and (29) and is a computationally convenient alternative to solving the LS heuristic. Accordingly, the SA heuristic would be a computationally efficient and cost effective heuristic in comparison to the LS heuristic for low holding cost and high backordering cost. The average percentage difference between the Static heuristic and the LS heuristic is 13.00% with a minimum of 7.24% and a maximum of 19.89%. The comparison shows that the incurred cost of using the Static heuristic is consistently significant and is not a cost effective alternative to the LS heuristic.

6.4. Numerical examples – fitting the demand process

In this section, we investigate the flexibility and limitation of the 2-MECO fitting approach in capturing the characteristics of the demand point process. This is motivated by the established re-

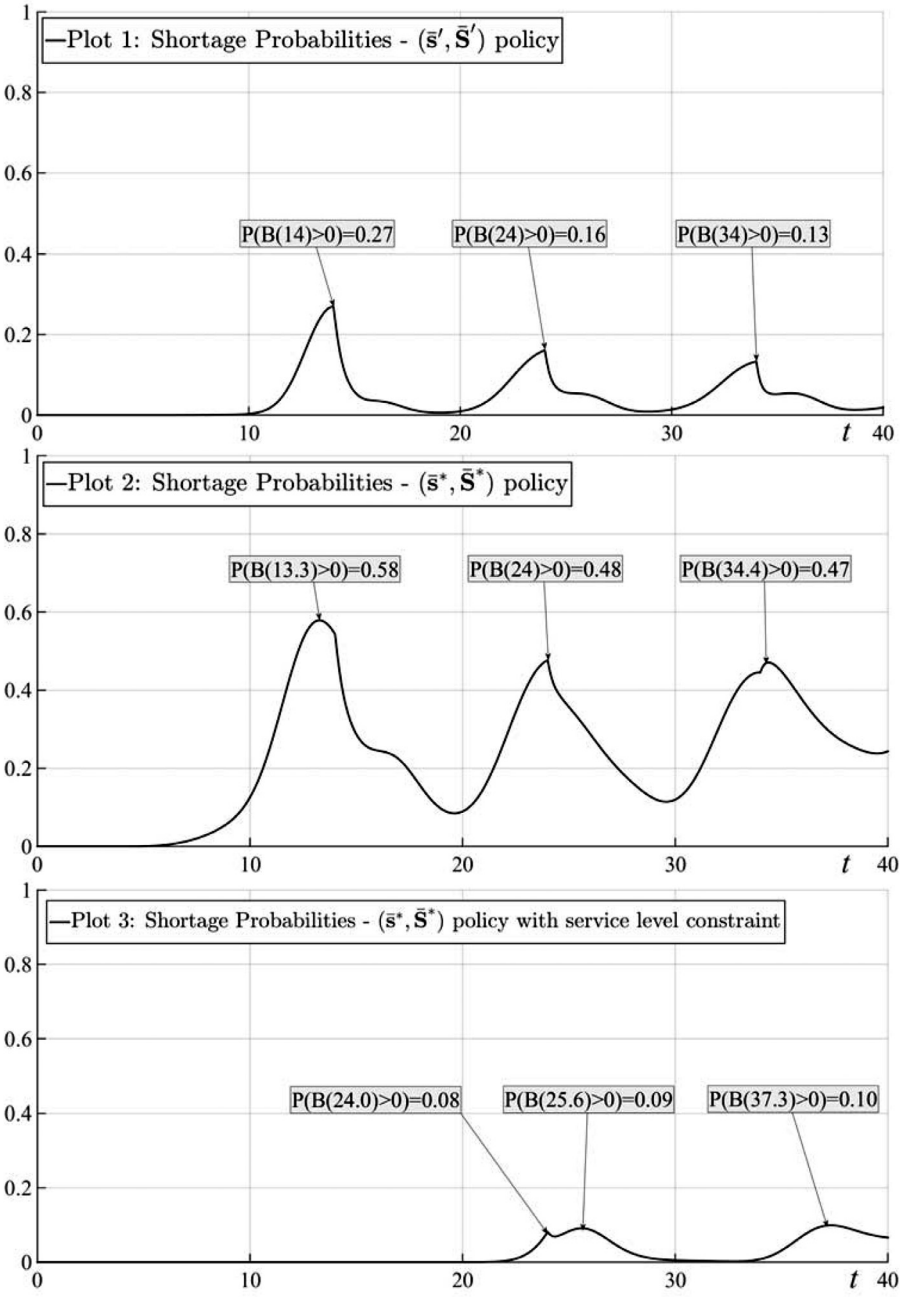

Fig. 13. Shortage probabilities.

sults in queueing theory that capturing the first two order descriptors of a point process can lead to different system measures, Andersen, Neuts, and Nielsen (2004). Ignoring third moments specifically can also lead to different queueing measures (Johnson & Luhman, 1994; Maddah, Nasr, & Charanek, 2017; Sahin & Perrakis, 1976; Whitt, 1984). Accordingly, we consider a source demand process which is more general than the renewal 2-MECO distribution. The demand process is represented by a Markovian Arrival Process (MAP), which exhibits auto-correlation between demand inter-arrival epochs. We consider a set of 10 cases where we keep the first moment fixed, but vary the dependency in the demand source process which also increases the variability in the counting process as well as the third moment of the counting process. The parameters of the source MAP are presented below where decreasing the parameter γ increases the dependency in the demand process as well as the variability in the counting process while maintaining the demand rate of $r(t)$.

$$\mathbf{D}_0 = \gamma \begin{pmatrix} -\nu_1 & \frac{3}{2} & \frac{1}{2} \\ \frac{3}{4} & -\nu_2 & \frac{3}{4} \\ \frac{1}{2} & \frac{3}{2} & -\nu_3 \end{pmatrix} \text{ and}$$

$$\mathbf{D}_1 = \begin{pmatrix} c_1\, r(t) & 0 & 0 \\ 0 & c_2\, r(t) & 0 \\ 0 & 0 & c_3\, r(t) \end{pmatrix}, \tag{31}$$

where $\nu_i = \sum_{j=1, j\neq i}^{3} \mathbf{D}_0(i,j) + c_i\, r(t)$, for $i = 1, 2, 3$. The parameters c_1, c_2 and c_3 are set to 0.1, 0.7 and 2.5, respectively. We refer to Lucantoni, Hellstern, and Neuts (1990) for more information on the MAP representation. A measure of the auto-correlation is

Table 5
Source cases – MAP.

Case	γ	$E[D_1^2]$	$E[D_2^2]$	$E[D_3^2]$	$E[D_4^2]$	$\text{Sk}(D_1)$	$\text{Sk}(D_2)$	$\text{Sk}(D_3)$	$\text{Sk}(D_4)$	$\bar{lag}_1$	$\bar{CVC}(\%)$
1	10	241.5	654	1267.8	2082.9	0.2975	0.2621	0.2326	0.2042	0.0069	1.54
2	5	243	658	1275.6	2095.7	0.3233	0.2958	0.2591	0.2424	0.0194	4.73
3	1	254.4	688.9	1335.5	2194.4	0.5000	0.4381	0.4221	0.4050	0.0966	30.44
4	0.75	258.9	701.1	1359.3	2233.6	0.5428	0.4834	0.4621	0.4473	0.1144	37.48
5	0.5	267.5	724.3	1404.6	2308.4	0.5981	0.5610	0.5376	0.5218	0.1400	47.57
6	0.25	289.5	785.3	1523.9	2505.5	0.7468	0.6988	0.6796	0.6661	0.1870	63.50
7	0.125	321.1	873.2	1696.4	2790.4	0.8622	0.8248	0.8067	0.7971	0.2414	76.38
8	0.1	332.2	904.4	1757.7	2891.9	0.8959	0.8557	0.8375	0.8267	0.2598	79.79
9	0.05	364.9	995.8	1937.2	3188.3	0.9510	0.9127	0.8949	0.8854	0.3131	88.12
10	0.025	389.3	1064.1	2071.2	3409.6	0.9749	0.9360	0.9189	0.9095	0.3538	93.42

Table 6
Fitting comparison.

Case	1	2	3	4	5	6	7	8	9	10
Δ_{SM} (%)	1.59	1.78	2.63	2.69	2.61	2.14	1.97	2.05	2.4	2.72
Δ_{Sk} (%)	53.33	58.18	68.47	69.48	71.01	74.87	82.04	84.59	90.26	92.73

Table 7
Inventory results for the MAP source demand.

Case	s_1	s_2	s_3	s_4	S_1	S_2	S_3	S_4	Cost_{MAP}
1	4	9	12	14	14	24	30	36	572.13
2	4	9	12	14	14	24	30	37	585.16
3	5	9	13	14	14	26	33	39	668.73
4	5	9	13	15	15	26	34	41	695.98
5	5	10	14	15	15	27	35	42	740.94
6	5	10	14	15	15	28	38	47	827.8
7	6	10	16	14	16	30	39	53	904.79
8	6	10	16	14	16	30	40	54	924.29
9	6	10	18	15	16	31	40	57	968.68
10	6	11	20	16	16	31	41	60	993.38

Table 8
Inventory results for the $P\tilde{h}_t$ approximating demand process.

Case	s_1	s_2	s_3	s_4	S_1	S_2	S_3	S_4	Cost_{Ph}
1	4	8	12	14	14	23	30	37	560.79
2	4	8	12	14	14	24	30	37	571.57
3	5	9	13	15	14	25	32	39	642.43
4	5	10	14	15	14	26	33	40	666.53
5	5	10	14	16	15	27	35	41	706.41
6	6	11	16	18	15	28	37	45	781.96
7	6	13	19	21	16	30	39	48	839.57
8	6	13	19	22	16	30	40	49	851.59
9	7	14	21	25	16	32	41	52	870.51
10	7	15	22	26	16	33	42	54	869.75

calculated by considering the lag_1 auto-correlation for the MAP process exhibiting a constant rate $\bar{r} = \int_0^T r(t)\,dt$. Another measure of correlation is denoted by the coefficient of variability due to correlation, $CVC = \frac{2\sum_{k=1}^{\infty} lag_k}{1+2\sum_{k=1}^{\infty} lag_k}$. The CVC accounts for the lag_k for $k \geq 1$. Decreasing γ also increases the third moment of the counting process which as measured by the skewness of the counting process. The Skewness parameter is calculated as the ratio of the third moment and the standard deviation, and is commonly used to test the asymmetry of the distribution of a counting process. The Skewness measure for each interval is calculated as, $\text{Sk}(D_i) = E[D_i^3]/\text{Stdev}(D_i)$.

We fit the rate $r(t)$ of Eq. (30) and the second moments D_i for $i = 1, 2, 3, 4$ of Table 5 to a 2-MECO process. Obviously, the 2-MECO is a renewal process and does not capture the correlation measures of the approximating distribution, $\bar{lag}_1$ and $\bar{CVC}$. We test the performance of the 2-MECO distribution in matching the second moment over any time interval of duration $L = 4$. In Appendix D, we plot the second moment of the source MAP process over the lead time, $D_t(L)$ for $0 \leq t \leq T$, and the second moment of the approximating 2-MECO process, $\tilde{D}_t(L)$ for $0 \leq t \leq T$, for Cases 1, 4, 7 and 10. The plots illustrate that the fitting approach provides an accurate fit for the second moment of the counting process. We calculate the average percentage difference, Δ_{SM}, for all 10 cases in Table 6 where,

$$\Delta_{SM} = \int_0^T \frac{\mid D_t(L) - \tilde{D}_t(L) \mid}{D_t(L)}\, d_t \times 100. \tag{32}$$

The 2-MECO consistently provides accurate approximations for the second moment of the counting process across the 10 cases as illustrated by Table 6 and the corresponding plots in Fig. 18.

Table 9
Comparison – source and fitted demand processes.

Case	Cost_{MAP}	Cost_{Ph}	$\Delta_1(\%)$	Cost	$\Delta_2(\%)$
1	572.13	560.79	−1.98	572.27	0.02
2	585.16	571.57	−2.32	585.39	0.04
3	668.73	642.43	−3.93	669.04	0.05
4	695.98	666.53	−4.23	696.61	0.09
5	740.94	706.41	−4.66	741.46	0.07
6	827.8	781.96	−5.54	831.61	0.46
7	904.79	839.57	−7.21	917.08	1.36
8	924.29	851.59	−7.87	937.94	1.48
9	968.68	870.51	−10.13	993.24	2.53
10	993.38	869.75	−12.44	1022.1	2.89

We also calculate the average percentage difference, Δ_{Sk}, when matching the Skewness parameter where,

$$\Delta_{Sk} = \sum_{i=1}^{4} \frac{\text{Sk}(D_i) - \text{Sk}(\tilde{D}_i)}{\text{Sk}(D_i)} \times 100, \tag{33}$$

and $\tilde{D}_i$ is the demand count of the approximating 2-MECO process over the ith time interval. The difference in the Skewness parameter is 53.33% for Case 1 and increases to reach 92.73% in Case 10.

The results of Table 6 illustrate that even though the 2-MECO provides accurate approximations for the first and second moments, it does not provide good approximations for the third moments (as measured by Δ_{Sk}).

Next, we numerically investigate the impact of increasing the correlation and skewness on the performance of the inventory system. The cost parameters considered are $\omega = \$25/\text{order}$, $b = \$5/\text{unit/unit time}$, and $h = \$1/\text{unit/unit time}$. The performance of

Table 10
Comparison with the stationary approximation and static heuristics.

	Stationary approximation heuristic										Static heuristic			
Case	s_1	s_2	s_3	s_4	S_1	S_2	S_3	S_4	Stationary cost	$\Delta_{SA}(\%)$	s	S	Static cost	$\Delta_{Static}(\%)$
1	7	11	15	19	16	23	29	34	613.13	7.17	10	26	693.74	21.26
2	7	11	15	19	16	23	29	34	623.91	6.62	10	26	707.52	20.91
3	7	11	16	21	16	23	29	36	711.67	6.42	11	27	796.9	19.17
4	7	12	16	21	16	23	30	37	736.75	5.86	11	27	826.44	18.74
5	7	12	17	23	17	24	31	38	788.94	6.48	11	28	874.95	18.09
6	8	13	19	25	17	25	33	41	883.18	6.69	11	29	970.41	17.23
7	9	15	21	27	18	26	34	43	982.25	8.56	12	30	1060	17.15
8	9	15	21	28	18	27	35	43	1000.44	8.24	12	30	1084.14	17.29
9	9	16	22	29	19	27	36	45	1052.71	8.67	12	30	1143.01	18.00
10	9	16	23	30	19	28	37	45	1081.06	8.83	12	30	1180.62	18.85

the inventory system is considered when the input demand is the source process (MAP) as well as the 2-MECO approximating process. The solution of the line search heuristic is considered for the MAP demand source process as well as the approximating $P\tilde{h}_t$ (2-MECO) in Tables 7 and 8, respectively. The results illustrate that as the dependence and skewness of the demand process increases, the incurred cost is significant. This is illustrated in Tables 7 and 8 where the percentage cost difference between Cases 1 and 10 is 73.63% for the MAP source process and 55.09% for the $P\tilde{h}_t$ process.

A comparison of the results of Tables 7 and 8 is presented in Table 9 where Δ_1 is the percentage difference between Cost_{MAP} and Cost_{Ph},

$$\Delta_1 = \frac{\text{Cost}_{Ph} - \text{Cost}_{MAP}}{\text{Cost}_{MAP}} \times 100. \tag{34}$$

As γ decreases, the cost of the source system, Cost_{MAP}, and the approximating system, Cost_{Ph}, increases for 1.98% (Case 1) to 12.44% (Case 10). Another practical measure is to evaluate the cost of implementing the ordering policies obtained from the $P\tilde{h}_t$ demand process (as presented Table 8), with an inventory system exhibiting the actual source demand process. Let Cost denote the cost of implementing the ordering policies of Table 8 in the inventory system exhibiting the actual source demand process. Let

$$\Delta_2 = \frac{\text{Cost} - \text{Cost}_{MAP}}{\text{Cost}_{MAP}} \times 100. \tag{35}$$

Accordingly, the percentage difference Δ_2 measures the incurred cost of utilizing the ordering policy obtained from the approximating system.

From a managerial perspective, the approximating policy underestimates the total cost (as measured by Δ_1) and the absolute difference increases as the dependence and skewness in the source demand process increases. The additional cost incurred from utilizing the approximating system's reordering policy (as measured by Δ_2) is less sensitive to the increase in dependence and skewness.

We compare the performance of the SA and Static heuristics relative to the LS heuristic across the 10 cases in Table 9. The percentage difference between the SA and LS heuristic is denoted by Δ_{SA} and the percentage difference between the Static heuristic and the LS heuristics is denoted by Δ_{Static}. The percentage difference Δ_{SA} is within the range 5.86% and 8.83% over the 10 cases. Similarly, the percentage difference Δ_{Static} is within 17.15% and 21.26%. We observe that the LS is consistently better than both heuristics. Although the cost increases as the variability in the counting process increases, the percentage difference between the heuristics is relatively stable (Table 10).

We also observe that simplifying the mathematical model by assuming a deterministic demand process leads to inaccurate and misleading results. This is a consequence of the stochastic component of the demand process. For example, Table 7 shows that the cost increases from 572.13 to 993.38 as a result of increasing the variability in the counting process. Denoting the solution of the deterministic case as a lower bound would result in a value which is less than 572.13. Consequently, obtaining a lower bound on the optimal cost by assuming a deterministic demand process is not a practical approach. The resulting lower bound is not a tight lower bound and would provide little insight on the optimal solution. To the best of our knowledge, obtaining practical lower bounds for the complex environment we consider in this paper (which includes Markovian demand and non zero lead time) is unattainable.

7. Conclusion and future work

This work presents a numerical approach to calculate the time-dependent probability distribution of the system performance measures which include the number of backorders, on-hand inventory, inventory position, and the ordering count process. An alternative to the approach presented in this work is to run Monte-Carlo simulations where computing point estimates within sufficiently small confidence intervals at every point in time can be computationally extensive. The analysis assumes a finite time-horizon, but the policy can be re-calculated or re-observed by simply re-initializing the set of differential equations at any future point in time. Re-initializing the differential equations at any point in time allows for a rolling time horizon analysis as well as investigating what if scenarios. Accordingly, the model can account for updated demand forecasts as they become available.

A main contribution of the computational approach is to provide a manager with a tool to conduct sensitivity analysis and closely observe the performance measures across the planning horizon for different reorder policies. Consequently, this results in a time-phased decision making process which is not primarily based on minimizing expected cost. The computational approaches presented also allow for an efficient calculation of the expected cost. This is illustrated by the numerical examples where we implement a search heuristic on a wide range of cost parameters.

Notice that the numerical examples assume piecewise constant reorder policies. In the case where a continuous-time reorder policy is realistic, the reorder policy does not have to be restricted to a piecewise constant policy. The model can account for any time-dependent reorder policy, i.e., the reorder levels s_t and S_t can take on any continuous-time function. Future work can include extensions to the computational framework presented in this paper. For example, the work can be extended to account for periodic review models. In such a case, the differential equations which describe the inventory behavior still hold but are re-initialized at the start of every period. The initial conditions for a period are determined

by the ending inventory from the previous period and the reorder levels.

The majority of the literature on stochastic inventory control accounts for the mean and the variance of the demand process. Although possible within the computational framework presented in this work, implementing fitting techniques to fit higher order descriptors of the Ph_t process would overshadow the contribution of the mathematical model presented in this work. Future work can consider fitting techniques to capture more descriptors of the demand process. Another extension can account for the more general time-dependent Markovian Arrival Process (MAP_t) to model the demand process. The work presented in this paper can be easily extended to the MAP_t case by simply adjusting the KFEs to account for the additional parameters which generalize a Ph_t to a MAP_t. Such an implementation can be considered when efficient approaches to fit time-dependent MAP_ts to time-dependent and correlated point processes are available. To the best of our knowledge such approaches are not available.

Acknowledgments

The authors gratefully acknowledge the financial support provided by the American University of Beirut Research Board under grant No 103187.

Appendix A

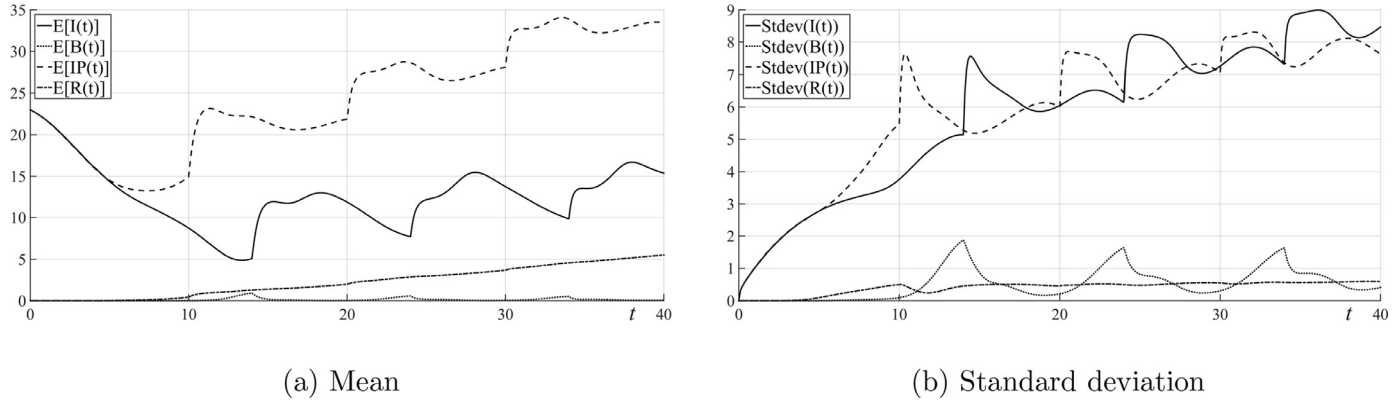

Fig. 14. Inventory measures –$(\bar{\mathbf{s}}', \bar{\mathbf{S}}')$ policy.

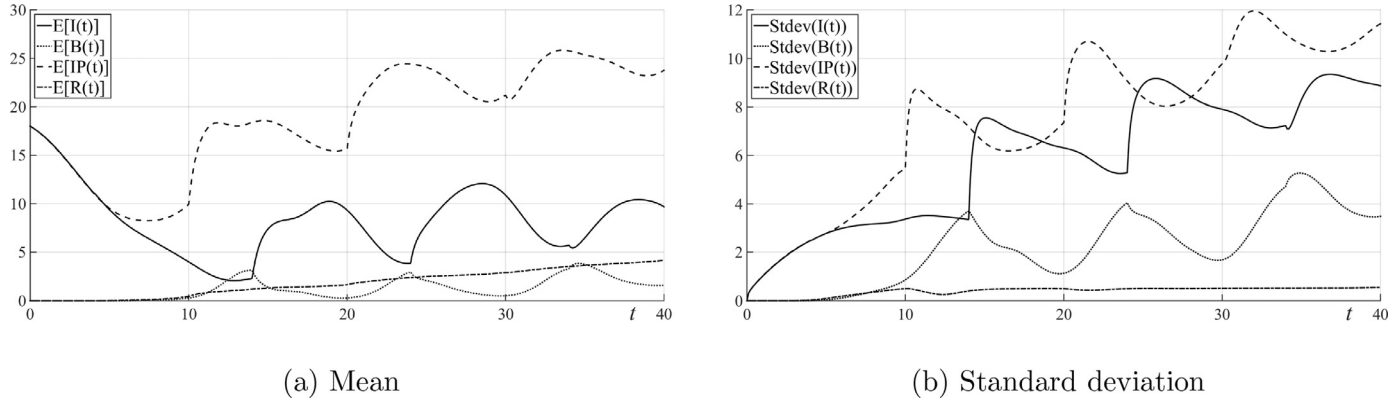

Fig. 15. Inventory measures – $(\bar{\mathbf{s}}^*, \bar{\mathbf{S}}^*)$ policy.

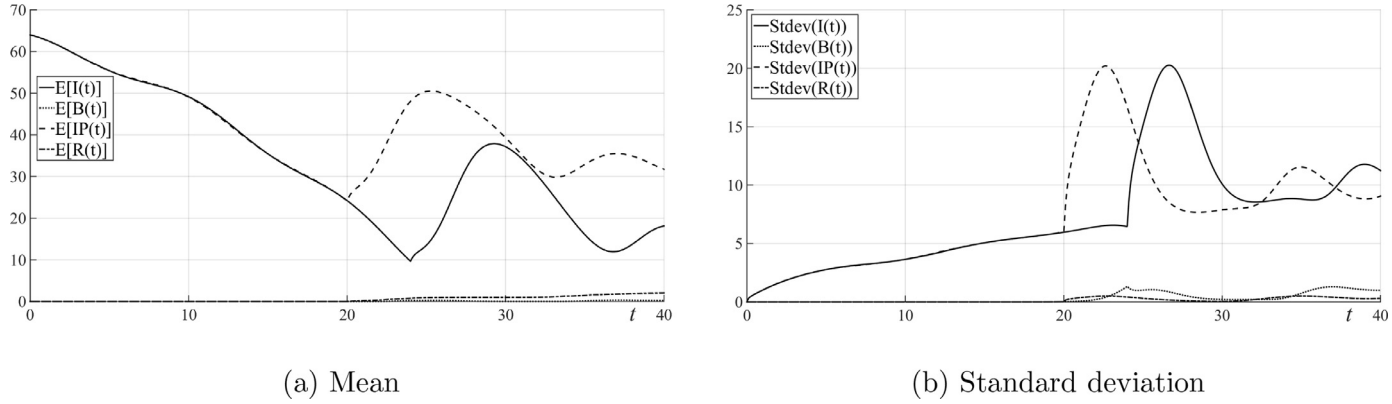

Fig. 16. Inventory measures – $(\bar{\mathbf{s}}^*, \bar{\mathbf{S}}^*)$ policy with service level constraint.

Appendix B

Table 11
Computational results for different inventory cost parameters.

ω	b	h	$\bar{s}^s$	$\bar{S}^s$	SH cost(S)	$\bar{s}'$				$\bar{S}'$				SA cost(S)	$\bar{s}^*$				$\bar{S}^*$				LS cost(S)	Savings (%) Static	Savings(%) SA
100	10	1	11	36	1180.0	7	11	15	19	25	34	42	49	1096.2	6	9	13	13	22	35	42	52	1048.0	11.2	4.4
		0.75	10	42	1009.0	7	11	15	19	28	37	46	54	944.5	7	9	13	14	24	39	46	59	912.2	9.6	3.4
		0.5	13	46	807.6	7	11	15	19	33	43	53	62	760.8	7	9	12	15	29	47	54	58	738.9	8.5	2.9
	5	1	7	33	1077.4	7	11	15	19	25	34	42	49	1073.6	2	6	10	9	20	32	39	48	960.4	10.9	10.5
		0.75	7	39	929.3	7	11	15	19	28	37	46	54	925.2	3	7	10	10	22	35	45	54	847.9	8.8	8.4
		0.5	9	42	757.0	7	11	15	19	33	43	53	62	750.3	4	7	10	12	26	45	51	55	700.3	7.5	6.7
	4	1	5	33	1041.3	7	11	15	19	25	34	42	49	1069.0	2	4	8	7	19	31	41	48	930.5	10.6	13.0
		0.75	6	38	901.7	7	11	15	19	28	37	46	54	921.4	2	6	9	8	21	35	44	52	823.6	8.7	10.6
		0.5	8	41	739.7	7	11	15	19	33	43	53	62	748.2	3	6	9	11	25	43	50	55	686.2	7.2	8.3
	3	1	2	33	988.0	7	11	15	19	25	34	42	49	1064.5	2	3	7	5	19	30	40	46	888.4	10.1	16.5
		0.75	4	36	863.4	7	11	15	19	28	37	46	54	917.5	2	4	7	6	20	35	43	49	790.7	8.4	13.8
		0.5	6	39	716.0	7	11	15	19	33	43	53	62	746.1	6	3	7	2	29	47	63	48	648.7	9.4	13.1
80	10	1	12	34	1087.0	7	11	15	19	23	31	39	46	995.2	6	9	13	14	20	33	41	46	955.5	12.1	4.0
		0.75	12	37	935.4	7	11	15	19	26	35	43	50	857.4	7	9	14	14	23	36	43	54	831.5	11.1	3.0
		0.5	12	44	742.8	7	11	15	19	30	40	49	57	699.2	8	9	13	15	26	46	54	57	674.8	9.2	3.5
	5	1	8	31	990.4	7	11	15	19	23	31	39	46	970.3	3	6	10	10	18	31	38	45	873.4	11.8	10.0
		0.75	9	35	862.3	7	11	15	19	26	35	43	50	835.8	4	7	11	10	21	33	41	51	770.0	10.7	7.9
		0.5	9	41	694.2	7	11	15	19	30	40	49	57	683.6	6	7	10	12	24	43	50	55	636.6	8.3	6.9
	4	1	7	30	956.6	7	11	15	19	23	31	39	46	965.3	2	5	9	8	18	30	37	44	843.5	11.8	12.6
		0.75	7	34	835.9	7	11	15	19	26	35	43	50	831.4	2	6	10	9	20	33	40	50	747.4	10.6	10.1
		0.5	8	40	677.1	7	11	15	19	30	40	49	57	680.5	5	6	10	11	23	42	48	55	622.8	8.0	8.5
	3	1	4	29	910.1	7	11	15	19	23	31	39	46	960.3	2	4	6	5	18	29	38	43	806.3	11.4	16.0
		0.75	4	35	797.6	7	11	15	19	26	35	43	50	827.1	2	5	8	7	19	32	40	49	716.8	10.1	13.3
		0.5	6	39	653.8	7	11	15	19	30	40	49	57	677.4	3	5	8	9	22	40	47	52	603.1	7.8	11.0
60	10	1	12	32	981.3	7	11	15	19	21	29	36	42	879.8	6	10	13	15	18	30	39	43	851.2	13.3	3.2
		0.75	13	35	842.3	7	11	15	19	23	31	39	46	758.8	7	10	14	15	21	34	42	47	739.7	12.2	2.5
		0.5	14	39	677.7	7	11	15	19	27	36	44	52	614.8	8	10	15	16	25	38	45	59	603.5	11.0	1.8
	5	1	9	29	888.7	7	11	15	19	21	29	36	42	853.3	3	7	10	11	16	28	37	41	774.8	12.8	9.2
		0.75	10	32	773.9	7	11	15	19	23	31	39	46	733.9	4	8	12	11	18	32	40	46	682.2	11.8	7.0
		0.5	9	40	629.2	7	11	15	19	27	36	44	52	594.3	6	8	13	13	23	36	43	56	565.7	10.1	4.8
	4	1	8	28	857.0	7	11	15	19	21	29	36	42	848.0	3	6	9	9	16	28	36	40	747.4	12.8	11.9
		0.75	9	31	750.2	7	11	15	19	23	31	39	46	729.0	3	7	11	10	18	31	39	45	661.4	11.8	9.3
		0.5	8	39	612.1	7	11	15	19	27	36	44	52	590.2	5	7	12	11	22	36	43	54	552.0	9.8	6.5
	3	1	6	27	813.1	7	11	15	19	21	29	36	42	842.7	2	5	8	7	15	27	33	40	711.1	12.5	15.6
		0.75	7	30	717.5	7	11	15	19	23	31	39	46	724.0	2	5	9	8	18	30	37	44	632.7	11.8	12.6
		0.5	7	38	589.4	7	11	15	19	27	36	44	52	586.1	3	7	10	9	21	34	42	52	532.8	9.6	9.1
40	10	1	13	30	863.0	7	11	15	19	19	25	32	38	752.7	5	10	14	16	16	27	33	40	725.1	16.0	3.7
		0.75	13	32	734.2	7	11	15	19	21	28	35	41	643.5	6	11	14	16	17	30	39	43	630.0	14.2	2.1
		0.5	14	36	585.7	7	11	15	19	24	32	40	47	522.0	8	11	15	16	22	35	43	48	513.5	12.3	1.6
	5	1	9	27	773.1	7	11	15	19	19	25	32	38	723.3	4	8	11	13	15	25	31	37	655.1	15.3	9.4
		0.75	10	30	668.8	7	11	15	19	20	28	34	41	615.0	4	9	12	13	16	28	37	41	576.7	13.8	6.2
		0.5	12	34	543.5	7	11	15	19	23	31	39	46	497.6	6	9	13	14	20	33	41	46	477.8	12.1	4.0
	4	1	9	27	742.8	7	11	15	19	19	25	32	38	717.5	3	7	10	12	14	25	31	37	630.9	15.1	12.1
		0.75	9	29	645.8	7	11	15	19	20	28	34	41	609.6	4	8	11	12	16	27	36	40	558.3	13.5	8.4
		0.5	11	33	528.6	7	11	15	19	23	31	39	46	492.6	5	8	12	12	19	32	40	46	465.2	12.0	5.6
	3	1	7	25	700.9	7	11	15	19	19	25	32	38	711.6	2	6	9	10	13	25	31	36	598.6	14.6	15.9
		0.75	8	28	615.2	7	11	15	19	20	28	34	41	604.1	3	7	9	10	15	27	35	38	533.4	13.3	11.7
		0.5	9	31	508.6	7	11	15	19	23	31	39	46	487.6	4	7	11	11	19	31	39	45	448.1	11.9	8.1
30	10	1	13	29	795.0	7	11	15	19	17	23	30	36	675.0	6	10	14	16	16	26	32	39	653.5	17.8	3.2
		0.75	14	31	672.9	7	11	15	19	19	26	32	39	576.1	6	11	15	17	17	28	34	41	564.3	16.1	2.1
		0.5	15	34	532.4	8	12	16	20	22	29	37	43	467.7	7	12	15	17	19	31	41	45	459.3	13.7	1.8
	5	1	10	26	706.9	7	11	15	19	17	23	30	35	642.6	4	8	12	13	14	24	30	36	584.9	17.3	9.0
		0.75	11	28	608.3	7	11	15	19	19	25	32	38	549.8	4	9	12	14	15	26	32	38	513.6	15.6	6.6
		0.5	12	32	490.7	7	11	15	19	21	29	36	42	439.9	6	10	13	15	18	30	39	43	425.6	13.3	3.2
	4	1	9	25	677.0	7	11	15	19	17	23	30	35	635.8	3	7	11	12	13	23	30	35	561.9	17.0	11.6
		0.75	10	28	586.0	7	11	15	19	19	25	32	38	544.0	4	8	12	13	15	25	32	37	496.4	15.3	8.8
		0.5	11	31	476.3	7	11	15	19	21	29	36	42	434.6	5	9	12	14	17	29	38	42	413.7	13.1	4.8
	3	1	7	24	637.4	7	11	15	19	17	23	30	35	629.0	2	6	10	11	12	22	29	35	531.6	16.6	15.5
		0.75	9	27	557.1	7	11	15	19	19	25	32	38	538.1	3	7	10	12	14	25	31	37	473.1	15.1	12.1
		0.5	10	30	457.1	7	11	15	19	21	29	36	42	429.3	4	8	11	12	16	28	37	41	397.8	13.0	7.3

(*continued on next page*)

Table 11 (*continued*)

ω	b	h	$\bar{s}^s$	$\bar{S}^s$	SH cost(S)	$\bar{s}^r$				$\bar{S}^r$				SA cost(S)	$\bar{s}^*$				$\bar{S}^*$				LS cost(S)	Savings (%) Static	Savings(%) SA
20	10	1	14	27	718.7	7	11	15	19	16	22	28	33	586.7	6	10	15	17	15	25	30	37	575.7	19.9	1.9
		0.75	14	29	603.2	8	12	16	20	17	23	30	36	507.1	7	11	15	17	16	26	32	39	491.7	18.5	3.0
		0.5	15	32	471.4	8	12	16	20	20	27	33	40	402.2	7	12	16	18	18	28	35	42	394.7	16.3	1.9
	5	1	11	25	633.4	7	11	15	19	15	21	27	32	549.2	4	8	12	14	13	23	28	35	509.1	19.6	7.3
		0.75	11	27	540.9	7	11	15	19	16	23	29	34	467.7	5	9	13	15	14	24	30	38	442.2	18.2	5.5
		0.5	13	30	431.5	7	11	15	19	19	25	32	38	376.3	5	10	14	16	16	27	33	40	362.5	16.0	3.7
	4	1	10	24	604.4	7	11	15	19	15	21	27	32	540.7	4	8	12	13	13	22	27	34	486.7	19.5	10.0
		0.75	10	26	519.3	7	11	15	19	16	23	29	34	460.4	4	8	12	14	14	24	30	37	425.2	18.1	7.6
		0.5	12	29	417.4	7	11	15	19	19	25	32	38	370.5	5	9	13	15	16	26	33	39	351.6	15.8	5.1
	3	1	8	23	565.7	7	11	15	19	15	21	27	32	532.3	3	7	10	11	12	21	27	33	457.4	19.1	14.1
		0.75	9	25	490.2	7	11	15	19	16	23	29	34	453.1	3	7	11	12	13	23	29	35	403.3	17.7	11.0
		0.5	10	28	398.6	7	11	15	19	19	25	32	38	364.6	4	9	12	14	15	26	32	38	337.0	15.5	7.6

Table 12
Numerical comparison of Models 1 and 2.

Case	s_1	s_2	s_3	s_4	S_1	S_2	S_3	S_4	Model 1 cost	Model 2 cost	Δ (%)
1	4	8	12	14	14	23	30	37	560.79	557.45	−0.60
2	4	8	12	14	14	24	30	37	571.57	568.16	−0.60
3	5	9	13	15	14	25	32	39	642.43	639.37	−0.48
4	5	10	14	15	14	26	33	40	666.53	663.97	−0.39
5	5	10	14	16	15	27	35	41	706.41	704.14	−0.32
6	6	11	16	18	15	28	37	45	781.9 6	783.16	0.15
7	6	13	19	21	16	30	39	48	839.57	851.78	1.43
8	6	13	19	22	16	30	40	49	851.59	868.11	1.90
9	7	14	21	25	16	32	41	52	870.51	907.69	4.10
10	7	15	22	26	16	33	42	54	869.75	927.72	6.25

Appendix C

C.1. Investigating the continuous ordering policy

The model in Section 4 assumes that an order is placed when a demand epoch is realized at time $k\delta$ and $\text{IP}(k\delta) \le s_k$, for $k = 1, \ldots, N-1$. Here we refer to the model of Section 4 by Model 1. We investigate altering the assumption by placing an order as soon as $\text{IP}(k\delta) \le s_k$ (Model 2).

The models behave differently when $s_{k-1} \le s_k$ and $s_{k-1} \le \text{IP}(k\delta) \le s_k$, for $k = 1, \ldots, N-1$. In the case where such a realization occurs, the time to place an order in Model 1 is delayed until the next demand arrival. Let X be the time for the next demand epoch from time

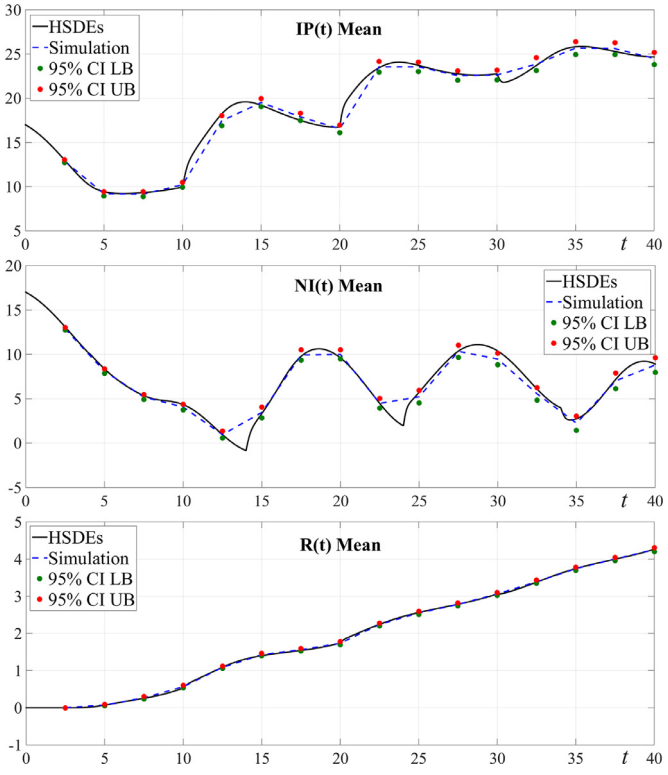

Fig. 17. Monte-Carlo simulation point estimates and CI.

$k\delta$. Model 1 results in an order delay of X time units (relative to Model 2) whenever $s_{k-1} \le \text{IP}(k\delta) \le s_k$. In the case where $\text{IP}(k\delta) > s_k$, then the ordering behavior of both models is identical (no order is placed).

Accordingly, the difference between both models depends on the duration of X. Notice that X represents the excess time of the duration of the inter-demand epoch (as defined by the Inspection Paradox, Ross, 2014). If Y is the duration of the inter-demand epoch then, $E[X] = \frac{1}{2}\frac{E[Y^2]}{E[Y]}$ (Ross, 2014). Since $E[X]$ is a function of the second moment of the inter-demand duration, we expect the difference in the performance of both models to be amplified as the variability in the demand process increases. We investigate this numerically by considering the Ph_t demand distributions utilized in Table 8 which exhibit an increasing trend in demand variability. We compare the results of Models 1 and 2 in Table 12. Notice that the cases in Table 12 are listed in ascending order of variability as illustrated in Section 6.4. The parameter Δ of Table 12 represents the percentage difference between the costs of Models 1 and 2. The numerical results illustrate that the performance of both Models is similar for Cases 1–6. Higher demand variability, Cases 7–10, results in a more significant difference as illustrated by Δ. The numerical results of Table 12 are consistent with our analytical analysis which indicates that behavior of Models 1 and 2 would vary as the variation in the demand process increases.

We also note that the behavior of Model 2 is captured by re-initializing the HSDE, Eqs. (5), at time $k\delta$ as follows. Re-initialize Eqs. (5) when $i = S_t$ to $P_{S_k,n}(k\delta) = \sum_{i=\ell}^{S_t} P_{i,n}$ for $n = 1 \ldots, m,\ k = 1, \ldots, N-1$ when $s_k > s_{k-1}$ and $\text{IP}(k\delta) \le s_k$.

C.2. Computational performance – Monte-Carlo simulation

We consider the base case demand process of Section 6.1 and calculate the time-dependent inventory measures, IP(t), NI(t) and R(t) (using the HSDEs of Sections 4.1, 4.2 and 4.4, respectively), and Monte-Carlo Simulation. We utilize 1000 replications to report the point estimates of the inventory measures at the points in time 2.5 to 40 in increments of 2.5, which results in 16 points in time (2.5, 5, 7.5, ..., 37.5, 40). The point estimates as well as the 95% Confidence Intervals (CI) for the selected points in time are plotted in Fig. 17 along with the time-continuous plot obtained from solving the HSDEs. The computational time of the HSDE for this numerical example is 0.734 seconds, compared to 99 seconds for the simulation run (not accounting for the time to sample the simulation run and generate the output report on the point estimates). The objective of this exercise is to provide a rough comparison of the computational performance of solving the HSDE vs Monte-Carlo simulation. It is also important to note that reporting on more systems statistics using the simulation approach would result in an even higher computational time.

Appendix D. Second moment of the counting process

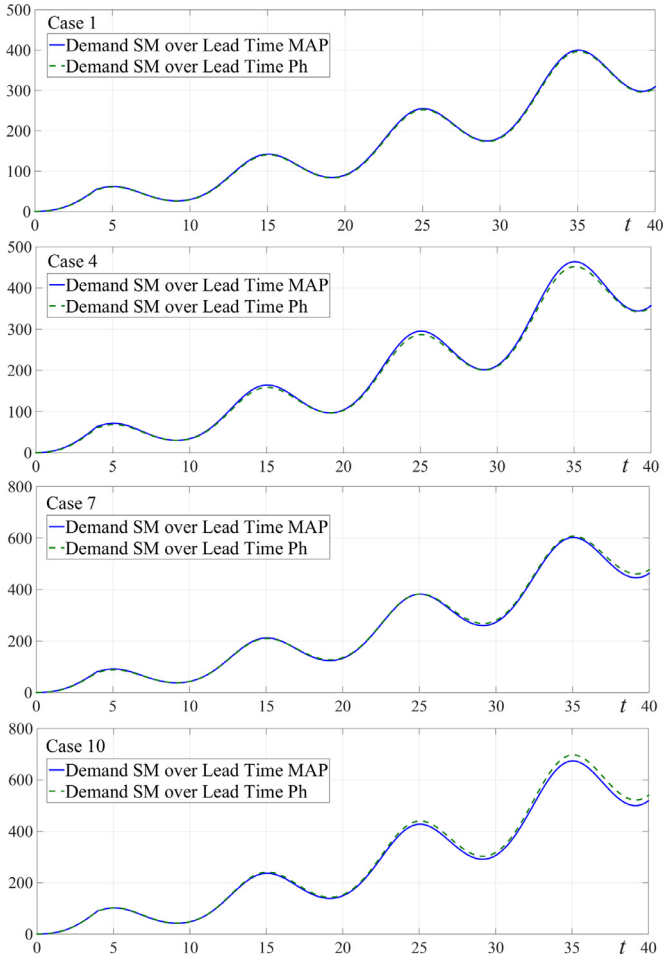

Fig. 18. Second moment over lead time.

References

Andersen, A. T., Neuts, M. F., & Friis Nielsen, B. (2004). On the time reversal of Markovian arrival processes. *Stochastic Models, 20*(3), 237–260.

Asmussen, S. (2008). *Applied probability and queues*: 51. Springer Science & Business Media.

Beyer, D., Cheng, F., Sethi, S. P., & Taksar, M. (2010). *Markovian demand inventory models*. New York: Springer.

Bollapragada, S., & Morton, T. E. (1999). A simple heuristic for computing non-stationary (*s, S*) policies. *Operations Research, 47*(4), 576–584.

Bookbinder, J. H., & Tan, J. Y. (1988). Strategies for the probabilistic lot-sizing problem with service-level constraints. *Management Science, 34*(9), 1096–1108.

Buchholz, P., Kriege, J., & Felko, I. (2014). *Input modeling with phase-type distributions and Markov models: Theory and applications*. Springer.

Dormand, J. R., & Prince, P. J. (1980). A family of embedded Runge–Kutta formulae. *Journal of computational and applied mathematics, 6*(1), 19–26.

Ehrenthal, J. C. F., Honhon, D. B. L. P., & Van Woensel, T. (2014). Demand seasonality in retail inventory management. *European Journal of Operational Research, 238*(2), 527–539.

Gerhardt, I., & Nelson, B. L. (2009a). Transforming renewal processes for simulation of nonstationary arrival processes. *INFORMS Journal on Computing, 21*(4), 630–640.

Gerhardt, I., & Nelson, B. L. (2009b). *On capturing dependence in point processes: Matching moments and other techniques* (pp. 1773–1808). Northwestern university. Technical report.

Gershwin, S. B., Dallery, Y., Papadopoulos, C. T., & Smith, J. M. (Eds.) (2012). Analysis and modeling of manufacturing systems, Vol. 60. Springer Science & Business Media.

Gusella, R. (1991). Characterizing the variability of arrival processes with indexes of dispersion. *IEEE Journal on Selected Areas in Communications, 9*(2), 203–211.

Johnson, M. A., & Luhman, J. A. (1994). Behaviour of queueing approximations based on sample moments. *Applied Stochastic Models in Business and Industry, 10*(4), 233–246.

Johnson, M. A., & Taaffe, M. R. (1989). Matching moments to phase distributions: mixtures of Erlang distributions of common order. *Stochastic Models, 5*(4), 711–743.

Johnson, M. A., & Taaffe, M. R. (1990). Matching moments to phase distributions: density function shapes. *Stochastic Models, 6*(2), 283–306.

Levi, R., Pl, M., Roundy, R. O., & Shmoys, D. B. (2007). Approximation algorithms for stochastic inventory control models. *Mathematics of Operations Research, 32*(2), 284–302.

Levi, R., Roundy, R. O., Shmoys, D. B., & Truong, V. A. (2008). Approximation algorithms for capacitated stochastic inventory control models. *Operations Research, 56*(5), 1184–1199.

Lim, Y. F., & Wang, C. (2017). Inventory management based on target-oriented robust optimization. *Management Science, 63*(12), 4409–4427.

Lucantoni, D. M., Meier-Hellstern, K. S., & Neuts, M. F. (1990). A single-server queue with server vacations and a class of non-renewal arrival processes. *Advances in Applied Probability, 22*(3), 676–705.

Ma, N., & Whitt, W. (2016). Efficient simulation of non-poisson non-stationary point processes to study queueing approximations. *Statistics & Probability Letters, 109*, 202–207.

Maddah, B., Nasr, W. W., & Charanek, A. (2017). A multi-station system for reducing congestion in high-variability queues. *European Journal of Operational Research, 2*(262), 602–619.

Moler, C., & Van Loan, C. (2003). Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later. *SIAM Review, 45*(1), 3–49.

Morton, T. E., & Pentico, D. W. (1995). The finite horizon nonstationary stochastic inventory problem: near-myopic bounds, heuristics, testing. *Management Science, 41*(2), 334–343.

Narayana, S., & Neuts, M. F. (1992). The first two moment matrices of the counts for the Markovian arrival process. *Communications in statistics. Stochastic models, 8*(3), 459–477.

Nasr, W. W., Charanek, A., & Maddah, B. (2017). MAP fitting by count and inter-arrival moment matching. (under review).

Nasr, W. W., & Maddah, B. (2015). Continuous (s, s) policy with MMPP correlated demand. *European Journal of Operational Research, 246*(3), 874–885.

Nasr, W. W., & Taaffe, M. R. (2013). Fitting the pht/mt/s/c time-dependent departure process for use in tandem queueing networks. *INFORMS Journal on Computing, 25*(4), 758–773.

Neale, J. J., & Willems, S. P. (2009). Managing inventory in supply chains with non-stationary demand. *Interfaces, 39*(5), 388–399.

Nelson, B. L., & Gerhardt, I. (2011). Modelling and simulating non-stationary arrival processes to facilitate analysis. *Journal of Simulation, 5*(1), 3–8.

Neuts, M. F. (1981). Matrix-geometric solutions in stochastic models: an algorithmic approach. *Courier Corporation*.

Nielsen, B. F., Nilsson, L. F., Thygesen, U. H. G., & Beyer, J. E. (2007). Higher order moments and conditional asymptotics of the batch Markovian arrival process. *Stochastic models, 23*(1), 1–26.

Qiu, R., Sun, M., & Lim, Y. F. (2017). Optimizing (s, s) policies for multi-period inventory models with demand distribution uncertainty: robust dynamic programing approaches. *European Journal of Operational Research, 261*(3), 880–892.

Ross, S. M. (2014). *Introduction to probability models*. Academic Press.

Sahin, I., & Perrakis, S. (1976). Moment inequalities for a class of single server queues. *INFOR: Information Systems and Operational Research, 14*(2), 144–152.

Scarf, H. (1960). The optimality of (s; s) policies in the dynamic inventory problem. In K. J. Arrow, S. Karlin, & P. Suppes (Eds.), *Mathematical methods in the social sciences*. chapter 13

Sethi, S. P., & Cheng, F. (1997). Optimality of (s, s) policies in inventory models with Markovian demand. *Operations Research, 45*(6), 931–939.

Silver, E. (1978). Inventory control under a probabilistic time-varying, demand pattern. *AIIE Transactions, 10*(4), 371–379.

Silver, E. A., Pyke, D. F., & Peterson, R. (1998). *Inventory management and production planning and scheduling: 3* p. 30. New York: Wiley.

Song, J. S., & Zipkin, P. (1993). Inventory control in a fluctuating demand environment. *Operations Research, 41*(2), 351–370.

Tarim, S. A., & Kingsman, B. G. (2006). Modelling and computing (rn, sn) policies for inventory systems with non-stationary stochastic demand. *European Journal of Operational Research, 174*(1), 581–599.

Tarim, S. A., & Smith, B. M. (2008). Constraint programming for computing non-stationary (*r, s*) inventory policies. *European Journal of Operational Research, 189*(3), 1004–1021.

Telek, M., & Horvath, G. (2007). A minimal representation of Markov arrival processes and a moments matching method. *Performance Evaluation, 64*(9), 1153–1168.

Tunc, H., Kilic, O. A., Tarim, S. A., & Eksioglu, B. (2011). The cost of using stationary inventory policies when demand is non-stationary. *Omega, 39*(4), 410–415.

Veinott Jr, F. A. (1966). On the opimality of (*s, S*) inventory policies: new conditions and a new proof. *SIAM Journal on Applied Mathematics, 14*(5), 1067–1083.

Whitt, W. (1982). Approximating a point process by a renewal process. I: Two basic methods. *Operations Research, 30*(1), 125–147.

Whitt, W. (1984). On approximations for queues. III: Mixtures of exponential distributions. *AT&T Bell Laboratories Technical Journal, 63*(1), 163–175.