Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Repairable Stocking and Expediting in a Fluctuating Demand Environment: Optimal Policy and Heuristics

Joachim Arts, Rob Basten, Geert-Jan Van Houtum

OPERATIONS RESEARCH
Articles in Advance, pp. 1–17
ISSN 0030-364X (print) | ISSN 1526-5463 (online)
http://dx.doi.org/10.1287/opre.2016.1498

Repairable Stocking and Expediting in a Fluctuating Demand Environment: Optimal Policy and Heuristics

Joachim Arts, Rob Basten, Geert-Jan Van Houtum
School of Industrial Engineering, Eindhoven University of Technology, Eindhoven, The Netherlands, PO Box 513, 5600MB
{j.j.arts@tue.nl, r.j.i.basten@tue.nl, g.j.v.houtum@tue.nl}

We consider a single stock-point for a repairable item facing Markov modulated Poisson demand. Repair of failed parts may be expedited at an additional cost to receive a shorter lead time. Demand that cannot be filled immediately is backordered and penalized. The manager decides on the number of spare repairables to purchase and on the expediting policy. We characterize the optimal expediting policy using a Markov decision process formulation and provide closed-form necessary and sufficient conditions that determine whether the optimal policy is a type of threshold policy or a no-expediting policy. We derive further asymptotic results as demand fluctuates arbitrarily slowly. In this regime, the cost of this system can be written as a weighted average of costs for systems facing Poisson demand. These asymptotics are leveraged to show that approximating Markov modulated Poisson demand by stationary Poisson demand can lead to arbitrarily poor results. We propose two heuristics based on our analytical results, and numerical tests show good performance with average optimality gaps of 0.11% and 0.33% respectively. Naive heuristics that ignore demand fluctuations have average optimality gaps of more than 11%. This shows that there is great value in leveraging knowledge about demand fluctuations in making repairable expediting and stocking decisions.

Keywords: expediting; Markov decision process; optimal policies; Markov modulated Poisson process; heuristics; inventory; dual-sourcing.
Subject classifications: dynamic programming: Markov; inventory/production: approximations/heuristics; reliability: maintenance/repairs.
Area of review: Operations and Supply Chains.
History: Received January 2014; revisions received May 2015, October 2015, February 2016; accepted March 2016. Published online in *Articles in Advance*.

1. Introduction

Both service and manufacturing industries depend on the availability of expensive equipment to deliver their products. Examples of such equipment include aircraft, rolling stock and manufacturing equipment. When this equipment is not working, the primary processes of their owners come to an immediate stop. To reduce the downtime of equipment, companies stock critical components such that the equipment can be returned to an operational state quickly by replacing a defective component with a ready-for-use component. Many components represent a significant financial investment and so they are repaired rather than discarded after a defect occurs. Consider for example, jet engines, bogies, or lens units for wafer steppers; these are components of aircraft, rolling stock, and integrated circuit manufacturing equipment, respectively, and their prices range from several hundreds of thousands up to millions of dollars. The best time for companies to buy these components is early in the lifecycle of the original equipment, because, at this time, it is possible to negotiate reasonable prices. In the literature, this is often referred to as the initial spare parts supply problem and it occurs in many different environments (e.g., Rustenburg et al. 2001, Pérès and Grenouilleau 2002). Also note that since such components are repairable, there is no natural inventory depletion that can reduce stock levels in case a company buys too many components. Later in the life cycle of a technical system, components often have to be custom made and prices are very steep, if the component can be purchased at all. An aggravating factor is that demand intensity for these components typically fluctuates over time, reflecting the fluctuating need for maintenance over time. Companies anticipate these demand fluctuations by leveraging the possibility of expediting the repair of defective components, rather than buying new components. Expediting a repair comes at a price, either because an external repair shop charges more for expedited repairs or because an internal repair shop can only handle a limited amount of expedited repairs. In the latter case, the cost of expediting can be thought of as a Lagrange multiplier that enforces a constraint on the number of expedited repairs that can be requested per time unit. Our model can then serve as a building block for a multi-item model with constraints on service and expediting frequencies.

Companies that operate in the environment described above face two major decisions related to their inventory control, one at the tactical level, and another at the operational level: (1) How many repairable spare parts should the firm buy? (tactical) (2) When should the firm request that the repair of a part be expedited? (operational). We refer to the first decision as the dimensioning decision and to the second as the expediting decision. The S spare repairables that are purchased early in the life cycle of a technical system are also called the *turn-around stock*. After this (initial) tactical decision, there is an operational recurring decision to either expedite or not expedite the repair of a spare part each time a demand/failure occurs. The latter decision should take demand fluctuations as well as current inventory levels into account. The model in this paper is intended to aid both the dimensioning and the expediting decision. For the dimensioning decision, it is important to consider the fact that expediting will occur later.

We study the decision problem described above via a stochastic inventory model. In this model, a defective item is replaced with a ready-for-use item and sent to a repair shop immediately after the defect occurs. At this point in time, the inventory manager is faced with the decision to either expedite or not expedite the repair of the part. Expediting a repair is more costly but has a shorter lead time. This expediting decision is informed by knowledge about the fluctuation of demand, which is modeled by a Markov modulated Poisson process. This demand model is quite rich in modeling fluctuations such as those that occur because of economic conditions, seasons of a year and the degradation of a fleet of equipment (Song and Zipkin 1993). It has also been observed empirically that this demand model fits well to practice since demand for repairable spare parts behaves as a nonstationary Poisson process that moves slowly relative to the replenishment lead time (Slay and Sherbrooke 1988).

We assume that inventory is replenished by an $(S-1, S)$-policy, meaning that each defective item is sent immediately to the repair shop. This replenishment policy is often used in practice and it is optimal when there are no economies of scale in replenishment. We model the expedited lead time as being deterministic and the regular lead time as being the sum of the expedited lead time and several exponential phases, the passing of which is monitored. This lead time model is a convenient device to investigate the value of tracking order progress information and the effect of different lead time distributions. (Gaukler et al. 2008, use a very similar model of order progress information.) Many lead time distributions can be modeled quite closely by this device. In practice, one often observes lead times that are close to deterministic and this device can approximate that arbitrarily closely by letting the number of exponential phases approach infinity.

The main contributions of this paper are the following: Firstly, we characterize the optimal repair expediting policy for the infinite horizon average and discounted cost criteria by formulating the problem as a Markov decision process (MDP). We find that the optimal policy may take two forms. The first form is simply to never expedite repair. The second form is a state dependent threshold policy, where the threshold depends on both the state of the Markov chain that modulates demand intensity, and the pipeline of repair orders. We also provide monotonicity results for the threshold as a function of the pipeline of repair orders. We give closed-form conditions that determine which of the two forms is optimal. In analyzing the optimal policy, we confirm a conjecture of Song and Zipkin (2009) that the expediting policy they propose is optimal for some special cases.

Secondly, we show that the joint problem of determining the turn-around stock and the expediting policy is not convex, but the cost function is submodular with respect to the turn-around stock and the expediting thresholds.

Thirdly, we show that when demand fluctuates arbitrarily slowly, the cost of any policy can be written as a weighted average of the cost for systems facing stationary Poisson demand. We employ this result to show that ignoring demand fluctuations by assuming stationary demand can lead to arbitrarily bad performance.

Finally, we propose two heuristics based on our analytical results for optimal policies in general and optimal policies for slowly fluctuating demand in particular. Numerical work shows that these policies perform very close to optimal and that naive heuristics that ignore demand fluctuations perform poorly.

This paper is organized as follows. In Section 2, we review relevant literature and position our contribution with respect to existing results. The model is described in Section 3. We study optimal expediting policies, optimal dimensioning and asymptotic result for slowly fluctuating environments in Sections 4–6, respectively. In Section 7, we propose heuristics based on our analytical results and test them numerically. Concluding remarks are provided in Section 8.

2. Literature Review

Our model is situated at the intersection of two streams of literature. The first one deals with sizing the turn-around stock of repairable item inventories and the second one with expediting, or inventory models with two (or more) supply modes.

An important characteristic of repairable item inventories is that inventory is replenished by repairing defective items. Repairable item inventory systems thus form a closed loop system that implicitly dictates the use of a base-stock policy. Often, the number of supported assets is large and the demand process is assumed to be independent of the number of outstanding orders. We make the same assumption which is in line with most of the repairable item inventory literature that was started with the METRIC model introduced by Sherbrooke (1968). Most of the important results in this stream of literature have been consolidated in the books by Sherbrooke (2004) and Muckstadt (2005), and the survey by Basten and Van Houtum (2014). This paper adds to the literature on repairable item inventories by studying what happens when it is possible to expedite the repair of a defective part, and in particular if this flexibility can be used to respond to a fluctuating demand environment. In doing this, we relax the commonly held assumption that demand is a stationary Poisson process. Our assumption of a Markov modulated Poisson process is more in line with empirical findings (Slay and Sherbrooke 1988). Verrijdt et al. (1998) already studied simple heuristics for the case that demand is a stationary Poisson process and emergency and regular repair lead times are both exponentially distributed. We relax the assumptions that the demand process is stationary and consider a more general lead time structure. Furthermore, we study optimal solutions as well as new heuristics informed by the structure of the optimal solution. We also remark that expediting a repair is not the same as shipping a ready-for-use part from a different stocking location which is commonly known as an emergency shipment (e.g., Alfredsson and Verrijdt 1999).

Inventory models with multiple supply modes have been reviewed by Minner (2003). Here we review the important and more recent results. Most authors consider a *periodic review* setting where the regular and expedited lead time differ by a single period and find that a base-stock policy is optimal for both the regular and expedited supply modes (e.g., Fukuda 1964). When the lead time of the regular and expedited supply modes differ by more than a single period, optimal policies do not exhibit simple structure and depend on the entire vector of outstanding orders (e.g., Whittmore and Saunders 1977, Feng et al. 2006). As a result, recent research considers heuristic policies for the control of dual supply systems, the most notable of these being the dual-index policy and variations thereof (Veeraraghavan and Scheller-Wolf 2008, Sheopuri et al. 2010, Arts et al. 2011). Under the dual-index policy, a regular and emergency inventory position are tracked separately, and both are kept at or above their respective order-up-to levels.

As opposed to the above mentioned papers, Moinzadeh and Schmidt (1991) consider a system running in *continuous time* facing Poisson demand with deterministic emergency and regular replenishment lead times. They show how to evaluate a given dual-index policy, although the name was not coined at the time, and the structure was not recognized as such. Song and Zipkin (2009) reinterpret the model of Moinzadeh and Schmidt (1991) revealing the simple structure of the policy and equivalence to a special type of queueing network that has a product form solution. Verrijdt et al. (1998) consider a similar system in the context of repairable items. In their model, the regular and expedited supply/repair modes have independent exponentially distributed lead times. They consider a different policy where repair is expedited when the inventory on-hand drops below a certain critical level.

While two different heuristic expediting policies have been suggested in the literature, one by Moinzadeh and Schmidt (1991) and Song and Zipkin (2009), and the other by Verrijdt et al. (1998), the optimal expediting policy has not yet been investigated. Song and Zipkin (2009) conjecture that their policy is optimal in some special cases. In this paper, we analyze the optimal repair expediting policy in the case of deterministic expedited repair lead times and stochastic regular repair lead times. As it turns out, the form suggested by Moinzadeh and Schmidt (1991) and Song and Zipkin (2009) is optimal in the special case that the regular repair lead time has a shifted exponential distribution and demand is a Poisson process. We note that Song and Zipkin (2009) also consider Markov modulated Poisson demand, but they focus on the performance evaluation of a stationary heuristic expediting policy, whereas we study optimal expediting policies that are not stationary.

3. Model Formulation

Our model supports two decisions that are on different time scales: (i) How to dimension the turn-around stock S and (ii) what expediting policy to follow. For the analysis, we use a nested procedure that determines the optimal expediting policy for a given turn-around stock (in Section 4), and uses this to determine the optimal turn-around stock (in Section 5). Below we give an integrated description of the model. In Section 3.1, we discuss the main assumptions and their justifications.

We consider a repairable item stock-point operated in continuous time with an infinite planning horizon $[0, \infty)$. The stock-point faces Markov modulated Poisson demand, i.e., demand is a Poisson process whose intensity varies with the state of an exogenous Markov process $Y(t)$. The Markov process $Y(t)$ is irreducible and has a finite state space $\Theta = \{1, \dots, |\Theta|\}$ with generator matrix $\mathbf{Q}$ whose elements we denote by q_{ij} and for convenience $q_i = -q_{ii}$. When $Y(t) = y$, the intensity of Poisson demand is given by $\lambda_y \geqslant 0$; $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_{|\Theta|})$; $\lambda_y > 0$ for at least one $y \in \Theta$. We denote demand in the time interval $(t_1, t_2]$ given $Y(t_1) = y$ as $D^y_{t_1, t_2}$. We assume that $Y(t)$ can be observed by the decision maker and so it provides a form of aggregated advance demand information.

The size of the turn-around stock of the repairable, $S \in \mathbb{N}_0$, is determined at time $t = 0$ and cannot be adapted afterward. (We relax this assumption in Section 5.1.) We assume that defective parts are sent to the repair shop immediately, i.e., we use an $(S-1, S)$ replenishment policy.

There exists a regular and an expedited repair option. The expedited repair lead time, l_e, is deterministic. The expedited

Figure 1. Repairable item inventory system with the possibility to expedite repair.

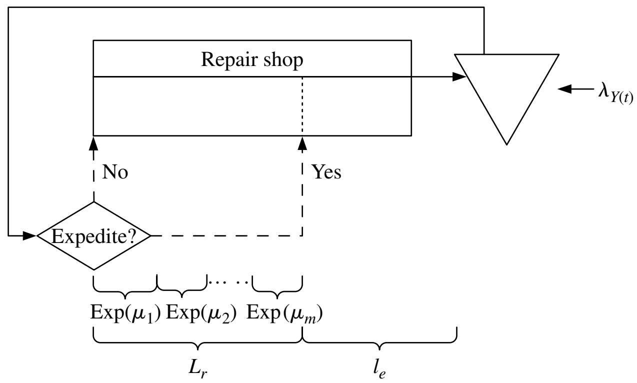

repair lead time can represent the minimal possible repair lead time or a lead time agreed upon with an external company that provides emergency repair service. We also refer to using the expedited repair mode as expediting a repair.

The regular repair lead time consists of the emergency repair lead time l_e, and a random component of length L_r, with mean $\mathbb{E}[L_r] < \infty$. We shall also refer to L_r as the *additional regular repair lead time*. We assume that this additional time is distributed as the sum of m exponential phases, with mean $1/\mu_i$ for the ith exponential phase. The inventory manager can observe the pipeline of outstanding orders and thus knows how many phases each part in the pipeline has completed. In particular, the inventory manager knows when the last phase (m) is completed and the remaining lead time of a regular order is l_e. A graphical representation of the system under study is given in Figure 1. Each failed part either enters regular or expedited repair when it is sent to the repair shop (and cannot change after that).

Turn-around stock holding (and depreciation) costs are incurred with a constant rate $h > 0$ for all repairable spare parts, regardless of where they are in the supply chain. Repair expediting costs per item are $c_e \geqslant 0$, i.e., c_e represents the (expected) cost difference between using the regular and emergency repair modes. A penalty cost rate $p > 0$ per item short per time unit for the repairable item inventory is also charged (backordering).

We let $X_i(t)$ denote the number of items in regular repair at time t that are in the ith phase of their additional repair lead time ($i = 1, \ldots, m$), and let $\mathbf{X}(t) = (X_1(t), \ldots, X_m(t))$. The following observation shows that $\mathbf{X}(t)$ and $Y(t)$ contain all the information needed to make expediting decisions. Let $c_p(x, y)$ denote the expected penalty cost rate at time $t + l_e$ conditional on $\sum_{i=1}^m X_i(t) = \mathbf{X}(t)\mathbf{e}^{\mathrm{T}} = x$ and $Y(t) = y$; c_p: $\mathbb{N}_0 \times \boldsymbol{\Theta} \to \mathbb{R}$ ($\mathbb{N}_0 = \mathbb{N} \cup \{0\}$ and $\mathbf{e} = (1, \ldots, 1)$). To find $c_p(x, y)$, note that $S - \mathbf{X}(t)\mathbf{e}^{\mathrm{T}} = S - x$ represents the net inventory at time t plus any parts that will arrive at the stock-point before time $t + l_e$. Thus the expected number of backorders at time $t + l_e$ given $\mathbf{X}(t)\mathbf{e}^{\mathrm{T}} = x$ and $Y(t) = y$ is $\mathbb{E}[(D^{Y(t)}_{t,\,t+l_e} - (S - \mathbf{X}(t)\mathbf{e}^{\mathrm{T}}))^+ \mid \mathbf{X}(t)\mathbf{e}^{\mathrm{T}} = x, Y(t) = y]$. From this it is easily verified that

$$
\begin{aligned}
c_p(x, y) &= p\,\mathbb{E}\big[(D^{Y(t)}_{t,\,t+l_e} - (S - \mathbf{X}(t)\mathbf{e}^{\mathrm{T}}))^+ \mid \mathbf{X}(t)\mathbf{e}^{\mathrm{T}} = x, Y(t) = y\big] \\
&= p \sum_{k=S-x}^{\infty} (k - (S - x))\,\mathbb{P}\{D^{y}_{t,\,t+l_e} = k\}. \qquad (1)
\end{aligned}
$$

When convenient, we use the notation $c_p(x, y \mid S)$ to explicitly show the dependence on S. To use (1), one must be able to evaluate $\mathbb{P}\{D^{y}_{t,\,t+l_e} = k\}$. This can be done numerically by inverting the generating function of $\mathbb{P}\{D^{y}_{t,\,t+l_e} = k \mid Y(t + l_e) = y'\}$ which is given in the form of a matrix exponential (e.g., Fischer and Meier-Hellstern 1992) and then unconditioning on the event $Y(t + l_e) = y'$. We relegate further details of this to Section EC.1 (available as supplemental material at http://dx.doi.org/10.1287/opre.2016.1498). Next, we note that whenever an item fails at time t, and its repair is not expedited, $X_1(t)$ increases by one. Thus, $\mathbf{X}(t)$ and $Y(t)$ contain all information needed to do cost accounting, and, in particular, to make optimal expediting decisions.

We are interested in minimizing the long run average cost rate $C(S) = hS + g^*(S)$ where $g^*(S)$ is the average expediting and penalty cost rate induced by an expediting policy that is optimal under turn-around stock level S:

$$
g^*(S) = \inf_{\pi \in \Pi} \limsup_{T \to \infty} \frac{1}{T} \mathbb{E}^{\pi}\Big[\int_0^T c_p(\mathbf{X}(t)\mathbf{e}^{\mathrm{T}}, Y(t) \mid S) + \lambda_{Y(t)} c_e \mathbf{1}^{\pi}(\mathbf{X}(t), Y(t))\, dt\Big]. \qquad (2)
$$

Here Π is the set of all Markovian policies for our system, $\mathbf{1}^{\pi}(\mathbf{x}, y)$ indicates whether policy π expedites demands in state $(\mathbf{x}, y)$, and $\mathbb{E}^{\pi}$ is the conditional expectation given policy π. Our decisions are (i) the turn-around stock, S, and (ii) a repair expediting policy $\pi \in \Pi$. (As an extension, we shall also consider minimizing total discounted costs for the expediting policy in Section 4.4, and the case where the holding costs for an item depend on its position in the pipeline in Section 4.5.)

3.1. Main Assumptions and Justifications

Some assumptions in our model require either a practical or analytical justification. Here we list the main assumptions and their justifications.

- The turn-around stock S is fixed: Although the turn-around stock may be increased by buying additional components, it never decreases since parts are repairable. Therefore, there is a fixed turn-around stock after a decision has been made to buy no more components. This situation usually occurs early in the life cycle of a technical system. We model this by assuming that all components are purchased at time $t = 0$. In some industries, it is not unusual that parts are not available from the original equipment manufacturer anymore a few years after the purchase of equipment. For example, this is the case for NedTrain, the company that motivated this study. NedTrain is a Dutch company that maintains rolling stock for the Dutch railways. This assumption may be relaxed under the infinite horizon discounted cost criterion; see Section 5.1.

• We consider an infinite planning horizon. The lifetime of repairables considered in the model is as long as the life cycle of the assets they support, which is typically several decades. This is long compared to other time characteristics in the problem such as lead times which are typically measured in weeks, and justifies using an infinite horizon model. Results for the finite horizon are also included in Theorem 1.

• Demand is a Markov modulated Poisson process: In spare parts literature, the Poisson demand model is perhaps the most common (e.g., Sherbrooke 2004, Muckstadt 2005). For relatively short periods of time, this demand model is often sufficiently accurate, and Markov modulated Poisson demand can handle Poisson demand as a special case. For longer periods of time, the demand intensity for repairables may be affected by things such as weather conditions (increased wear) and periodic inspections. Slay and Sherbrooke (1988) observe that demand for aircraft components behaves as a Poisson process for which the rate varies slowly over time. There are many reasons for this behavior such as weather, asset loading, and the fact that many capital assets undergo one or more major revisions during their lifetime. Demand for repairables peaks during these revision periods, as inspections reveal latent failures. Often, the exact timing of revision periods is uncertain when the asset is acquired. The Markov modulated Poisson process offers the flexibility to model these and many other demand scenarios while retaining tractability. It can also approximate a nonstationary Poisson process whose intensity over time is known using techniques as summarized in Artalejo et al. (2010). This model also enables us to asses the value of demand fluctuation information.

• The additional regular repair lead time, L_r, is a sum of exponential phases, and phase completions can be observed: Many inventory planners and IT systems work with deterministic planned lead times. To model all lead times in our model as deterministic is not only intractable computationally but also prohibits gaining structural results. The reason for this is that the state space becomes infinite dimensional: For each possible outstanding order (which can be infinitely many) we need to keep track of the remaining lead time. A pragmatic approach when the real L_r is (close to) deterministic is to set m large and interpret each stage as a deterministic time bucket. If the true distribution has $m < \infty$ phases, then the true state $\mathbf{X}(t)$ at time t cannot be inferred with certainty only from past order release times, but certain states are more likely than others.

We think of the m exponential phases of L_r primarily as a device to model order progress information, similar to Gaukler et al. (2008). The value of order progress information can be assessed using this model as will be shown in Section 7.4. Each time a phase passes, the remaining lead time of an item decreases so that inventory managers have an idea of the remaining lead time of outstanding orders. This model also allows for fitting lead times on the first two moments as long as $c_{L_r}^2 = \mathrm{Var}[L_r]/\mathbb{E}^2[L_r] \leqslant 1$ (Aldous and Shepp 1987) and can approximate deterministic lead time and perfect pipeline information as $m \to \infty$. Thus, this model allows us to gain insight on the added value of tracking repairable order progress carefully. In Section EC.3, we test our model via simulation for robustness against different lead time distribution assumptions and find that it is nearly insensitive.

• There is a penalty cost per backorder per time unit. Each backorder leads to the downtime of a technical system. Therefore, p may be interpreted as the cost of downtime per system and unit of time. If there is a requirement on the availability of the technical system, p can also be interpreted as a Lagrange multiplier enforcing this constraint.

4. Optimal Expediting Policies

In this section, we consider the problem of finding optimal repair expediting policies for fixed S, and call this problem $\mathfrak{M}(S)$. (In Section 5, we consider the joint problem of sizing the turn-around stock and expediting policy.) Since the holding costs depend linearly on S only, we need not consider holding cost in finding an optimal expediting policy for a fixed S.

We make several steps in our analysis. First in Section 4.1, we give closed-form conditions under which the state space can be truncated to yield a finite state space for the purpose of finding average optimal policies. This is not only technically convenient, but also necessary to ensure that the transition rates of the resulting MDP are bounded so that the Bellman equations can be used to analyze the structure of the optimal policy. We also show that when these conditions do not hold, the policy that never expedites repair is optimal. After that in Section 4.2, we formulate a finite horizon finite state space Markov decision process. The average optimal expediting policy is characterized in Section 4.3 and the infinite horizon discounted version in Section 4.4. A generalization of the holding cost structure is considered in Section 4.5.

4.1. State Space Truncation

Let Δ denote the difference operator with respect to the first argument of a function, i.e., $\Delta c_p(x,y) = c_p(x+1,y) - c_p(x,y)$. The following lemma establishes some useful properties of $c_p(x,y)$. The proof of Lemma 1, in Section EC.2.1, is similar to the proof of these same properties for the cost function of a news-vendor problem.

Lemma 1. *$c_p(x,y)$ has the following properties:*

(i) *$c_p(x+1,y) \geqslant c_p(x,y)$ for all $x \in \mathbb{N}_0$ and $y \in \Theta$, i.e., $c_p(x,y)$ is nondecreasing in x.*

(ii) $\Delta c_p(x+1, y) \geqslant \Delta c_p(x, y)$ *for all* $x \in \mathbb{N}_0$ *and* $y \in \Theta$, *i.e.,* $c_p(x, y)$ *is convex in* x.

(iii) $\Delta c_p(x, y) \leqslant p$ *for all* x *and* $y \in \Theta$ *and* $\Delta c_p(x, y) = p$ *for all* $x \geqslant S$ *and* $y \in \Theta$.

(iv) $\Delta c_p(x, y \mid S) \geqslant \Delta c_p(x, y \mid S+1)$ *for all* $x \in \mathbb{N}_0$, $y \in \Theta$ *and* $S \in \mathbb{N}_0$, *i.e.,* $c_p(x, y \mid S)$ *is submodular with respect to* x *and* S.

Proposition 1 below allows us to truncate the relevant state space if $c_e < p\mathbb{E}[L_r]$, and fully characterizes an optimal expediting policy if $c_e \geqslant p\mathbb{E}[L_r]$. It can be understood intuitively by making the following observation: Whenever a repair is expedited, this may avert a backorder at most for the additional regular repair lead time L_r. Thus, when the cost of expediting a repair is more than or equal to the expected backorder cost over the additional regular repair lead time, expediting is never beneficial (in expectation). Conversely, if expediting is cheaper than the cost of a backorder over the expected additional regular lead time, then expediting is beneficial in expectation if the number of parts already in repair that will not arrive within the expedited lead time is sufficiently large.

Proposition 1. *For the infinite horizon, average cost criterion, the following statements hold*

(i) *If* $c_e > p\mathbb{E}[L_r]$, *then the policy to never expedite is the* only *optimal policy.*

(ii) *If* $c_e = p\mathbb{E}[L_r]$, *then the policy to never expedite is* an *optimal policy.*

(iii) *If* $c_e < p\mathbb{E}[L_r]$, *then there is an* $M \in \mathbb{N}$ *such that any optimal policy will expedite repair whenever* $\mathbf{X}(t)\mathbf{e}^T \geqslant M$.

Proof. Let $\mathbf{e}_i \in \mathbb{R}^m$ be the unit vector in direction i. For parts (i) and (ii), the proof is based on showing that any policy that expedites in some state when $c_e \geqslant p\mathbb{E}[L_r]$ can be improved by a policy that is identical except that it does not expedite in that state. Let π denote an arbitrary policy that expedites for some state $(\mathbf{x}, y)$. Suppose now that at time t', the process is in state $(\mathbf{x}, y)$ and a demand occurs. Let $(\mathbf{X}(t), Y(t))$ denote the process under policy π at any time t. Next we construct a coupled process, $(\mathbf{X}'(t), Y(t))$, that is identical to $(\mathbf{X}(t), Y(t))$ except that the failed part arriving at time t' is *not* expedited. Let $\tilde{\mathbf{X}}(t)$ denote the evolution through the pipeline of the part arriving at time t' if it enter regular repair, i.e., $\tilde{\mathbf{X}}(t) = \mathbf{e}_i$ if the part sent to regular repair at time t' has completed its first $i-1$ phases of the additional regular repair lead time at time t, and $\tilde{\mathbf{X}}(t) = \mathbf{0}$ if the part has completed its additional regular repair lead time.

With this notation, we can write $\mathbf{X}'(t) = \mathbf{X}(t) + \tilde{\mathbf{X}}(t)$. Now let $T_r = \inf\{t - t' \mid \tilde{\mathbf{X}}(t) = \mathbf{0}, t \geqslant t'\}$ and note that $T_r \stackrel{d}{=} L_r$, where $\stackrel{d}{=}$ denotes equality in distribution. By construction, any cost difference between the processes $(\mathbf{X}'(t), Y(t))$ and $(\mathbf{X}(t), Y(t))$ must occur in the interval $[t', t' + T_r)$, because these processes are identical outside that interval. In $[t', t' + T_r)$, $\mathbf{X}(t)$ incurs exactly c_e more emergency repair costs because of the part expedited at time t', and $\mathbf{X}'(t)$ incurs more penalty costs because $\mathbf{X}'\mathbf{e}^{\mathrm{T}} = \mathbf{X}(t)\mathbf{e}^{\mathrm{T}} + 1$ for $t \in [t', t' + T_r)$. The expected cost difference between the processes $(\mathbf{X}(t), Y(t))$ and $(\mathbf{X}'(t), Y(t))$ thus satisfies

$$c_e - \mathbb{E}_{T_r}\left\{\mathbb{E}_{(\mathbf{X}(t), Y(t))}\left[\int_{t=t'}^{t'+T_r} \Delta c_p(\mathbf{X}(t)\mathbf{e}^{\mathrm{T}}, Y(t))\, dt \,\middle|\, T_r\right]\right\}$$
$$\geqslant c_e - \mathbb{E}_{T_r}[pT_r] = c_e - p\mathbb{E}[L_r] \geqslant 0,$$

where the first inequality holds by Lemma 1(iii). Notice that the latter inequality is strict when $c_e > p\mathbb{E}[L_r]$. This proves parts (i) and (ii).

Now for part (iii), we use the same coupling construction. The proof and intuition for part (iii) coincide nicely (although the proof has many technical details): As $\mathbf{X}\mathbf{e}^{\mathrm{T}}$ grows, the probability that expediting a repair avoids a backorder for L_r time units approaches unity. Therefore, the number of outstanding orders can be so large that the expected cost of expediting will almost surely be smaller than the expected additional backordering costs incurred by failing to expedite the repair of a part.

Let $\varepsilon = (p\mathbb{E}[L_r] - c_e)/3 > 0$. We denote the probability density function of L_r as f_{L_r} and fix $\alpha < \infty$ to verify

$$\int_{t=\alpha}^{\infty} t f_{L_r}(t_r)\, dt \leqslant \varepsilon / p. \tag{3}$$

Such an α exists because $tf_{L_r}(t) > 0$ for $t \in (0, \infty)$ so that $\int_{t=\alpha}^{\infty} tf_{L_r}(t)\, dt$ is strictly decreasing in α and furthermore $\lim_{\alpha\to\infty} \int_{t=\alpha}^{\infty} tf_{L_r}(t)\, dt = 0$. Let E_μ denote an exponential random variable with mean μ^{-1}. We fix an integer M' to verify

$$\mathbb{P}\{E_{\mu_m} < \alpha\}^{M'}\, \mathbb{E}[L_r] \leqslant \varepsilon. \tag{4}$$

Such an $M' \in \mathbb{N}$ exists because $\alpha < \infty$ and so $\mathbb{P}\{E_{\mu_m} < \alpha\} < 1$.

Now we consider an arbitrary policy π that does *not* expedite when $\mathbf{x}\mathbf{e}^{\mathrm{T}} \geqslant S + M' = M$ for some $(\mathbf{x}, y) \in \mathcal{S}$. Consider an arbitrary moment in time, t', when a failed part arrives to the system and $\sum_{i=1}^{m} X_i(t') \geqslant S + M' = M$ and policy π stipulates that the part should *not* be expedited. Denote this process $\mathbf{X}^\pi(t)$. We let $\tilde{\mathbf{X}}(t)$ denote the evolution of the part sent to regular repair at time t' by policy π, so $\tilde{\mathbf{X}}(t) = \mathbf{e}_i$ if the part sent to repair at time t' has completed its first $i-1$ phases of the additional regular repair, and $\tilde{\mathbf{X}}(t) = \mathbf{0}$ if the part has completed its additional regular repair lead time. Next, we consider an alternate process which is identical to $\mathbf{X}^\pi(t)$ except that it does expedite the unit arriving at t'. We denote this process $\mathbf{X}^e(t)$, and formally define it as $\mathbf{X}^e(t) = \mathbf{X}^\pi(t) - \tilde{\mathbf{X}}(t)$. We let $T_r = \inf\{t - t' \mid \tilde{\mathbf{X}}(t) = \mathbf{0}, t \geqslant t'\}$ and note that $T_r \stackrel{d}{=} L_r$.

Analogous to the proof of parts (i) and (ii), $\mathbf{X}^\pi(t)\mathbf{e}^{\mathrm{T}} = \mathbf{X}^e(t)\mathbf{e}^{\mathrm{T}} + 1$ for $t \in [t', t' + T_r)$, and $\mathbf{X}^\pi(t) = \mathbf{X}^e(t)$ for $t \geqslant t' + T_r$. Also both processes make exactly the same expediting decisions for all $t > t'$. Thus any cost differences between $\mathbf{X}^e(t)$ and $\mathbf{X}^\pi(t)$ occur in the time interval $[t', t'+T_r)$. Denote the expectation of this cost difference Ξ. Then we have

$$\begin{aligned}
\Xi &= c_e - \mathbb{E}_{T_r}\left\{\mathbb{E}_{(\mathbf{X}^e(t),Y(t))}\left[\int_{t=t'}^{t'+T_r}\Delta c_p(\mathbf{X}^e(t)\mathbf{e}^{\mathrm{T}}, Y(t))\,dt \,\middle|\, T_r\right]\right\}\\
&= c_e - \int_{t_r=0}^{\infty}\mathbb{E}_{(\mathbf{X}^e(t),Y(t))}\left[\int_{t=t'}^{t'+T_r}\Delta c_p(\mathbf{X}^e(t)\mathbf{e}^{\mathrm{T}}, Y(t))\,dt \,\middle|\, T_r=t_r\right]\\
&\quad\cdot f_{L_r}(t_r)\,dt_r\\
&= c_e - \int_{t_r=0}^{\infty}\mathbb{E}_{(\mathbf{X}^e(t),Y(t))}\left[\int_{t=t'}^{t'+T_r}\Delta c_p(\mathbf{X}^e(t)\mathbf{e}^{\mathrm{T}}, Y(t))\,dt \,\middle|\, T_r=t_r,\ \mathbf{X}^e(t)\mathbf{e}^{\mathrm{T}} \geqslant S \text{ for all } t\in(t', t'+T_r)\right]\\
&\quad\cdot \mathbb{P}\left\{\mathbf{X}^e(t)\mathbf{e}^{\mathrm{T}} \geqslant S \text{ for all } t\in(t', t'+t_r)\right\} f_{L_r}(t_r)\,dt_r\\
&\quad - \int_{t_r=0}^{\infty}\mathbb{E}_{(\mathbf{X}^e(t),Y(t))}\left[\int_{t=t'}^{t'+T_r}\Delta c_p(\mathbf{X}^e(t)\mathbf{e}^{\mathrm{T}}, Y(t))\,dt \,\middle|\, T_r=t_r,\ \mathbf{X}^e(t)\mathbf{e}^{\mathrm{T}} < S \text{ for some } t\in(t', t'+T_r)\right]\\
&\quad\cdot \mathbb{P}\left\{\mathbf{X}_e(t)\mathbf{e}^{\mathrm{T}} < S \text{ for some } t\in(t', t'+t_r)\right\} f_{L_r}(t_r)\,dt_r\\
&\leqslant c_e - \int_{t_r=0}^{\infty}\mathbb{E}_{(\mathbf{X}^e(t),Y(t))}\left[\int_{t=t'}^{t'+T_r}\Delta c_p(\mathbf{X}^e(t)\mathbf{e}^{\mathrm{T}}, Y(t))\,dt \,\middle|\, T_r=t_r,\ \mathbf{X}^e(t)\mathbf{e}^{\mathrm{T}} \geqslant S \text{ for all } t\in(t', t'+T_r)\right]\\
&\quad\cdot \mathbb{P}\left\{\mathbf{X}^e(t)\mathbf{e}^{\mathrm{T}} \geqslant S \text{ for all } t\in(t', t'+t_r)\right\} f_{L_r}(t_r)\,dt_r\\
&= c_e - \int_{t_r=0}^{\infty} p t_r f_{L_r}(t_r)\\
&\quad\cdot \mathbb{P}\left\{\mathbf{X}^e(t)\mathbf{e}^{\mathrm{T}} \geqslant S \text{ for all } t\in(t', t'+t_r)\right\} dt_r \qquad (5)
\end{aligned}$$

The third equality is obtained by conditioning on whether or not $\mathbf{X}^e(t)\mathbf{e}^{\mathrm{T}}$ stays above S on the interval $[t', t'+T_r)$. The first inequality follows from dropping the last term and the last equality follows from part (iii) of Lemma 1.

Next we observe that $\mathbb{P}\{\mathbf{X}^e(t)\mathbf{e}^{\mathrm{T}} \geqslant S \text{ for all } t\in(t', t'+t_r)\}$ is bounded below by the probability that fewer than M' parts already in additional regular repair at time t', finish additional regular repair before $t'+t_r$. Since the remaining time in regular repair for any of these parts is at least an E_{μ_m} random variable (by the lack of memory property), we conclude that

$$\begin{aligned}
&\mathbb{P}\left\{\mathbf{X}^e(t)\mathbf{e}^{\mathrm{T}} \geqslant S \text{ for all } t\in(t', t'+t_r)\right\}\\
&\qquad \geqslant 1-\mathbb{P}\{E_{\mu_m} < t_r\}^{M'}. \qquad (6)
\end{aligned}$$

Now continuing from (5) and using (6) we obtain

$$\begin{aligned}
\Xi &\leqslant c_e - \int_{t_r=0}^{\infty} p t_r f_{L_r}(t_r)\left(1-\mathbb{P}\{E_{\mu_m} < t_r\}^{M'}\right) dt_r\\
&= c_e - p\,\mathbb{E}[L_r] + \int_{t_r=0}^{\infty} p t_r f_{L_r}(t_r)\,\mathbb{P}\{E_{\mu_m} < t_r\}^{M'}\,dt_r\\
&= c_e - p\,\mathbb{E}[L_r] + \int_{t_r=0}^{\alpha} p t_r f_{L_r}(t_r)\,\mathbb{P}\{E_{\mu_m} < t_r\}^{M'}\,dt_r\\
&\quad + \int_{t_r=\alpha}^{\infty} p t_r f_{L_r}(t_r)\,\mathbb{P}\{E_{\mu_m} < t_r\}^{M'}\,dt_r\\
&\leqslant c_e - p\,\mathbb{E}[L_r] + \mathbb{P}\{E_{\mu_m} < \alpha\}^{M'} \int_{t_r=0}^{\alpha} p t_r f_{L_r}(t_r)\,dt_r\\
&\quad + \int_{t_r=\alpha}^{\infty} p t_r f_{L_r}(t_r)\,dt_r \qquad (7)\\
&\leqslant c_e - p\,\mathbb{E}[L_r] + \mathbb{P}\{E_{\mu_m} < \alpha\}^{M'}\,\mathbb{E}[L_r]\\
&\quad + \int_{t_r=\alpha}^{\infty} p t_r f_{L_r}(t_r)\,dt_r\\
&\leqslant -3\varepsilon + \varepsilon + \varepsilon = -\varepsilon < 0. \qquad (8)
\end{aligned}$$

Inequality (7) follows because $\mathbb{P}\{E_{\mu_m} < t_r\}$ is increasing in t_r and the final inequalities follow from the choice of ε, α and M'. Since $\Xi < 0$, we conclude that the expected cost of process $\mathbf{X}^\pi(t)$ is greater than the cost of $\mathbf{X}^e(t)$. Thus, we have shown that any policy that does not expedite when $\mathbf{X}(t)\mathbf{e}^{\mathrm{T}} \geqslant M$ and $c_e < p\,\mathbb{E}[L_r]$ can be strictly improved by expediting whenever $\mathbf{X}(t)\mathbf{e}^{\mathrm{T}} \geqslant M$. That is, if $c_e < p\,\mathbb{E}[L_r]$, then there is a $M \in \mathbb{N}$ such that whenever $\mathbf{X}(t)\mathbf{e}^{\mathrm{T}} \geqslant M$ it is optimal to expedite. □

Proposition 1 characterizes the optimal policy when $c_e \geqslant p\,\mathbb{E}[L_r]$. Note however that part (iii) of Proposition 1 does not characterize the optimal policy when $c_e < p\,\mathbb{E}[L_r]$. It only allows us to restrict the state space of $(\mathbf{X}(t), Y(t))$ to the finite set $\mathcal{S} = \{(\mathbf{x}, y) \in \mathbb{N}_0^m \times \Theta \mid \mathbf{x}\mathbf{e}^{\mathrm{T}} \leqslant M\}$ for some $M \in \mathbb{N}$ as constructed in the proof of Proposition 1. Proposition 1 thus enables us to study the optimal policy when $c_e < p\,\mathbb{E}[L_r]$ within a Markov decision process framework because transition rates are bounded after truncation of the state-space.

4.2. MDP Formulation with Bounded Transition Rates

In this section, we consider the model $\mathfrak{M}(S)$ with $c_e < p\,\mathbb{E}[L_r]$, and state space $\mathcal{S} = \{(\mathbf{x}, y) \in \mathbb{N}_0^m \times \Theta \mid \mathbf{x}\mathbf{e}^{\mathrm{T}} \leqslant M\}$, where M is chosen, as outlined in the proof of Proposition 1, such that it is optimal to expedite whenever $\mathbf{X}\mathbf{e}^{\mathrm{T}} \geqslant M$. With a slight abuse of notation, we term the problem of finding an optimal policy for this model as $\mathfrak{M}(S, M)$. In this finite state space, transition rates are bounded and so we can apply the technique of uniformization to transform the problem of finding an optimal expediting policy to discrete time.

In each state $(\mathbf{x}, y)$, we take a decision as to whether we expedite the repair of a part if the next event happens to be the arrival of a defective part. We let 1 denote the decision to expedite if a part arrives and let 0 be the decision to not expedite if a part arrives. Thus the action space in state $(\mathbf{x}, y)$ is $\mathcal{A}(\mathbf{x}, y) = \{0, 1\}$ when $\mathbf{x}\mathbf{e}^{\mathrm{T}} < M$ and $\mathcal{A}(\mathbf{x}, y) = \{1\}$ otherwise. Observe that if we take a decision 1 in some state of the system, this does not necessarily imply we will expedite some part, because the next event in the systems may not be the arrival of a defective part.

As uniform transition rate for this MDP, we choose $\Lambda = \lambda_{\max} + M\sum_{i=1}^{m}\mu_i + q_{\max}$ with $\lambda_{\max} = \max_{y\in\Theta}\lambda_y$ and $q_{\max} = \max_{y\in\Theta} q_y$. Let $p((\mathbf{x}', y') \mid (\mathbf{x}, y), a)$ denote the transition probability from state $(\mathbf{x}, y) \in \mathcal{S}$ to $(\mathbf{x}', y') \in \mathcal{S}$ under action $a \in \mathcal{A}(\mathbf{x}, y)$. Without loss of generality, we rescale time such that $\Lambda = 1$. Then we have

$$p((\mathbf{x}', y') \mid (\mathbf{x}, y), a) = \begin{cases} \lambda_y, & \text{if } \mathbf{x}' = \mathbf{x} + \mathbf{e}_1,\ y' = y,\ a = 0; \\ x_m\mu_m, & \text{if } \mathbf{x}' = \mathbf{x} - \mathbf{e}_m,\ y' = y,\ a \in \{0, 1\}; \\ x_i\mu_i, & \text{if } \mathbf{x}' = \mathbf{x} - \mathbf{e}_i + \mathbf{e}_{i+1},\ y' = y, \\ & \quad a \in \{0, 1\},\ i = 1, \ldots, m-1; \\ q_{y,y'}, & \text{if } \mathbf{x}' = \mathbf{x},\ y' \neq y,\ a \in \{0, 1\}; \\ \sum_{i=1}^{m} (M - x_i)\mu_i + q_{\max} - q_y + \lambda_{\max} - \lambda_y, & \\ & \text{if } (\mathbf{x}', y') = (\mathbf{x}, y),\ a = 0; \\ \sum_{i=1}^{m} (M - x_i)\mu_i + q_{\max} - q_y + \lambda_{\max}, & \\ & \text{if } (\mathbf{x}', y') = (\mathbf{x}, y) \text{ and } a = 1; \\ 0, & \text{otherwise,} \end{cases} \tag{9}$$

where $\mathbf{e}_i$ is the ith unit vector in dimension m. Regardless of the decision taken, between transitions, an expected penalty cost of $c_p(\mathbf{x}\mathbf{e}^{\mathrm{T}}, y)$ is incurred. Additionally, a cost of c_e is incurred if an arriving defective part is expedited.

Now let $V_n(\mathbf{x}, y)$ denote the optimal total cost function when in state $(\mathbf{x}, y)$ and having n transitions to go and define $V_0(\mathbf{x}, y) \equiv 0$. The finite horizon dynamic programming recursion (Bellman equation) is given by

$$\begin{aligned} V_{n+1}(\mathbf{x}, y) &= c_p(\mathbf{x}\mathbf{e}^{\mathrm{T}}, y) + \lambda_y \mathbf{1}_{\{\mathbf{x}\mathbf{e}^{\mathrm{T}} < M\}} \min\{c_e + V_n(\mathbf{x}, y), V_n(\mathbf{x} + \mathbf{e}_1, y)\} \\ &\quad + \lambda_y \mathbf{1}_{\{\mathbf{x}\mathbf{e}^{\mathrm{T}} = M\}} (c_e + V_n(\mathbf{x}, y)) + \sum_{i=1}^{m-1} x_i\mu_i V_n(\mathbf{x} - \mathbf{e}_i + \mathbf{e}_{i+1}, y) \\ &\quad + x_m\mu_m V_n(\mathbf{x} - \mathbf{e}_m, y) + \sum_{i=1}^{m} (M - x_i)\mu_i V_n(\mathbf{x}, y) \\ &\quad + \sum_{y' \in \Theta \setminus \{y\}} q_{yy'} V_n(\mathbf{x}, y') \\ &\quad + (q_{\max} - q_y + \lambda_{\max} - \lambda_y) V_n(\mathbf{x}, y), \end{aligned} \tag{10}$$

where $\mathbf{1}_{\{\cdot\}}$ is the indicator function.

Remark 1. Note that the smallest possible uniformization constant is given by $\Lambda' = \lambda_{\max} + M\mu_{\max} + q_{\max}$, where $\mu_{\max} = \max_{i \in \{1, \ldots, m\}} \mu_i$. We use a formulation based on this smaller constant in our numerical work (Section 7.4), because it leads to quicker convergence of value iteration algorithms (e.g., Kulkarni 1999). Here we work with the formulation based on Λ so that we can reuse some results in the literature to prove structural properties of optimal policies. ◇

To analyze the value function $V_n(\mathbf{x}, y)$ in Section 4.3, we employ the event based dynamic programming approach introduced by Koole (1998, 2006). To this end, let $\mathcal{V}$ denote the set of all functions $v: \mathcal{S} \to \mathbb{R}$ and let $f, f_1, \ldots, f_{m+2} \in \mathcal{V}$. We define the following operators $\mathbb{T}_{\text{cost}}$, $\mathbb{T}_{\text{AC}(i)}$, $\mathbb{T}_{\text{TD}(i)}$, $\mathbb{T}_{\text{D}(i)}$, $\mathbb{T}_{\text{env}}$: $\mathcal{V} \to \mathcal{V}$, $\mathbb{T}_{\text{unif}}$: $\mathcal{V}^{m+2} \to \mathcal{V}$.

$$\mathbb{T}_{\text{cost}} f(\mathbf{x}, y) = c_p(\mathbf{x}\mathbf{e}^{\mathrm{T}}, y) + f(\mathbf{x}, y) \tag{11}$$

$$\begin{aligned} \mathbb{T}_{\text{AC}(i)} f(\mathbf{x}, y) &= \mathbf{1}_{\{\mathbf{x}\mathbf{e}^{\mathrm{T}} < M\}} \min\{c_e + f(\mathbf{x}, y), f(\mathbf{x} + \mathbf{e}_i, y)\} \\ &\quad + \mathbf{1}_{\{\mathbf{x}\mathbf{e}^{\mathrm{T}} = M\}} (c_e + f(\mathbf{x}, y)) \end{aligned} \tag{12}$$

$$\mathbb{T}_{\text{TD}(i)} f(\mathbf{x}, y) = \frac{x_i}{M} f(\mathbf{x} - \mathbf{e}_i + \mathbf{e}_{i+1}, y) + \frac{M - x_i}{M} f(\mathbf{x}, y) \tag{13}$$

$$\mathbb{T}_{\text{D}(i)} f(\mathbf{x}, y) = \frac{x_i}{M} f(\mathbf{x} - \mathbf{e}_i, y) + \frac{M - x_i}{M} f(\mathbf{x}, y) \tag{14}$$

$$\begin{aligned} \mathbb{T}_{\text{env}} f(\mathbf{x}, y) &= \sum_{y \in \Theta \setminus \{y\}} q_{yy'} f(\mathbf{x}, y') \\ &\quad + (q_{\max} - q_y + \lambda_{\max} - \lambda_y) f(\mathbf{x}, y) \end{aligned} \tag{15}$$

$$\begin{aligned} &\mathbb{T}_{\text{unif}}(f_1, \ldots, f_{m+2})(\mathbf{x}, y) \\ &\quad = \lambda_y f_1(\mathbf{x}, y) + \sum_{i=1}^{m} M\mu_i f_{i+1}(\mathbf{x}, y) + f_{m+2}(\mathbf{x}, y). \end{aligned} \tag{16}$$

These operators are variations to operators defined by Koole (1998, 2004, 2006) and are originally intended to model various common queueing mechanisms such as arrival control ($\mathbb{T}_{\text{AC}(i)}$), transfer departures from multiserver tandem queues ($\mathbb{T}_{\text{TD}(i)}$), and departures from multiserver queues ($\mathbb{T}_{\text{D}(i)}$), while the operators $\mathbb{T}_{\text{cost}}$, $\mathbb{T}_{\text{env}}$, and $\mathbb{T}_{\text{unif}}$ are mainly convenient for bookkeeping. The Bellman recursion for our MDP, (10), can now be written succinctly as

$$\begin{aligned} V_{n+1}(\mathbf{x}, y) = \mathbb{T}_{\text{cost}} \mathbb{T}_{\text{unif}} \big[&\mathbb{T}_{\text{AC}(1)} V_n(\mathbf{x}, y), \mathbb{T}_{\text{TD}(1)} V_n(\mathbf{x}, y), \ldots, \\ &\mathbb{T}_{\text{TD}(m-1)} V_n(\mathbf{x}, y), \mathbb{T}_{\text{D}(m)} V_n(\mathbf{x}, y), \\ &\mathbb{T}_{\text{env}} V_n(\mathbf{x}, y) \big]. \end{aligned} \tag{17}$$

This formulation of the MDP recursion is convenient because the propagation of value function properties over n can be analyzed through the propagation properties of operators, for which results are available in the literature.

We remark that the operators used to rewrite the MDP recursion reveal that the MDP we are dealing with is equivalent to an admission control problem for a tandem line of ample exponential server queues. A similar equivalence is exploited by Song and Zipkin (2009) in finding effective means to evaluate heuristic policies.

4.3. Average Optimal Expediting Policies

To characterize average optimal policies, we study properties of the value function and how these properties propagate through recursion (17). We define the first order difference operator with respect to x_i, Δ_i, as $\Delta_i f(\mathbf{x}, y) = f(\mathbf{x} + \mathbf{e}_i, y) - f(\mathbf{x}, y)$. We distinguish the following subsets of $\mathcal{V}$:

$$\mathcal{I}(i) = \{f \in \mathcal{V} \mid f(\mathbf{x}, y) \leqslant f(\mathbf{x} + \mathbf{e}_i, y)\} \tag{18}$$

$$\mathcal{C}(i) = \{f \in \mathcal{V} \mid \Delta_i f(\mathbf{x}, y) \leqslant \Delta_i f(\mathbf{x} + \mathbf{e}_i, y)\} \tag{19}$$

$$\mathcal{UI} = \{f \in \mathcal{V} \mid f(\mathbf{x}+\mathbf{e}_{i+1}, y) \leqslant f(\mathbf{x}+\mathbf{e}_i, y),\ i=1,\ldots,m-1\} \quad (20)$$

$$\mathcal{SM}(i,j) = \{f \in \mathcal{V} \mid \Delta_i f(\mathbf{x}, y) \leqslant \Delta_i f(\mathbf{x}+\mathbf{e}_j, y)\}. \quad (21)$$

In (18)–(21), it is understood that the inequalities that characterize each set must hold when the arguments on both sides of the inequality exist in $\mathcal{S}$. The sets $\mathcal{I}(i)$ and $\mathcal{C}(i)$ contain nondecreasing and convex function with respect to x_i respectively. The set $\mathcal{UI}$ contains upstream increasing functions as introduced in Koole (2004) and renamed in Koole (2006). The set $\mathcal{SM}(i,j)$ consists of functions with a specific supermodularity property. Finally, define $\mathcal{F}$ as

$$\mathcal{F} = \left(\bigcap_{i=1}^{m} \mathcal{I}(i)\right) \cap \mathcal{C}(1) \cap \mathcal{UI} \cap \left(\bigcap_{j=2}^{m} \mathcal{SM}(1,j)\right). \quad (22)$$

Lemma 2. *The following statements hold*

(i) *The function $c_p(\mathbf{x}\mathbf{e}^T, y) \in \mathcal{V}$ is a member of $\mathcal{F}$.*

(ii) *If $f \in \mathcal{F}$ then $\mathbb{T}_{\text{cost}} f(\mathbf{x}, y)$, $\mathbb{T}_{\text{AC}(1)} f(\mathbf{x}, y)$, $\mathbb{T}_{\text{D}(m)} f(\mathbf{x}, y)$, $\mathbb{T}_{\text{env}} f(\mathbf{x}, y) \in \mathcal{F}$ and $\mathbb{T}_{\text{TD}(i)} f(\mathbf{x}, y) \in \mathcal{F}$ for $i = 1, \ldots, m-1$.*

(iii) *If $f_j \in \mathcal{F}$ for $j = 1, \ldots, m+2$, then $\mathbb{T}_{\text{unif}}(f_1, \ldots, f_{m+2}) \cdot (\mathbf{x}, y) \in \mathcal{F}$.*

The proof of this lemma is in Section EC.2.2. The properties of functions in $\mathcal{V}$ that are shown to propagate through operators (11)–(16) in Lemma 2, imply structure on the optimal policy. To state the next lemma, we introduce some notation. Let $\mathbf{x}^{(-1)}$ denote the vector $\mathbf{x}$ with its first component set to 0, i.e., $\mathbf{x}^{(-1)} = (0, x_2, \ldots, x_m)$. The next theorem characterizes the optimal policy for a finite planning horizon problem. (The proof of this Theorem is in Section EC.2.3.)

Theorem 1. *For a finite planning horizon with n decision epochs to go, if $V_{n-1} \in \mathcal{F}$, then, at transition epoch n, there are state dependent thresholds $T_n(\mathbf{x}^{(-1)}, y)$ such that it is optimal to expedite the repair of an arriving part at transition epoch n if and only if $X_1(t_n) \geqslant T_n(\mathbf{X}^{(-1)}(t_n), Y(t_n))$, where t_n is the time corresponding to transition epoch n. Furthermore, the thresholds $T_n(\mathbf{x}^{(-1)}, y)$ satisfy the following monotonicity property: $\Delta_i T_n(\mathbf{x}^{(-1)}, y) \leqslant 0$, for $i = 2, \ldots, m$.*

An alternative interpretation of Theorem 1 is that the optimal policy at transition epoch n (under the stated condition) is a switching curve between expediting and not expediting repairs. This switching curve is decreasing in x_i for $i = 2, \ldots, m$. Figure 2 illustrates two such switching curves. In part (a) of the figure, for a given x_2, it is optimal to expedite repair if x_1 is on or above the shown line. In part (b) of the figure, for given (x_2, x_3) it is optimal to expedite repair if x_1 is on or above the shown surface.

The policy described in Theorem 1 can also be reinterpreted as a state dependent expedite-up-to policy. To see this, define $\text{IP}_e(t) = S - \mathbf{X}(t)\mathbf{e}^{\text{T}}$, and note that this can be interpreted as the expedited inventory position: on-hand inventory minus backorders plus outstanding orders arriving within the expedited lead time l_e. The optimal policy is now to expedite parts to retain $\text{IP}_e(t_n)$ at or above the level $S - T_n(\mathbf{X}^{(-1)}(t_n), Y(t_n))$. Thus the resulting policy is a state dependent version of the dual-index policy

Figure 2. Part (a) shows the state dependent threshold for $n = 593$ in the case where L_r has an Erlang(2) distribution ($m = 2$). Part (b) shows the state dependent thresholds for $n = 1513$ when L_r is Erlang(3) distributed ($m = 3$).

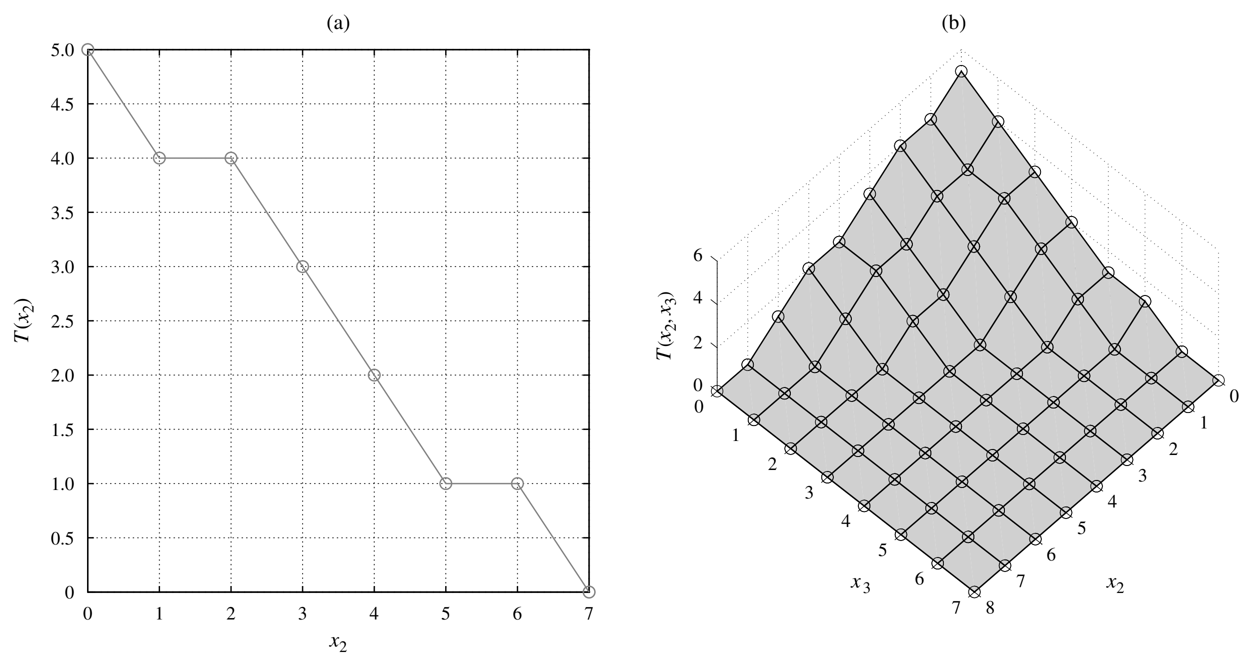

Notes. Both cases are based on the problem instance with $|\Theta| = 1$, $\lambda_1 = 1$, $\mathbb{E}[L_r] = 4$, $l_e = 2$, $c_e = 8$, $p = 10$, and $S = 8$. In both cases, n coincides with the iteration in which average optimal policies are found within some precision.

(Veeraraghavan and Scheller-Wolf 2008, Arts et al. 2011, consider state independent dual-index policies), where regular and emergency inventory positions are both kept at or above their order-up-to levels. Note however, that the regular order-up-to level was assumed to be S from the start as we are dealing with a closed loop system. Without this fixed base-stock level, a state dependent dual-index replenishment policy need not be optimal.

The main result of this section is that average optimal policies also have the structure described in Theorem 1 for the finite horizon problem. The proof of this result is in Section EC.2.4.

Theorem 2. *Consider the model* $\mathfrak{M}(S)$. *If* $c_e \geqslant p\,\mathbb{E}[L_r]$, *then it is average optimal to never expedite repair. If* $c_e < p\,\mathbb{E}[L_r]$, *then there are state dependent threshold levels* $T(\mathbf{x}^{(-1)}, y) \in \mathbb{N}_0$ *such that it is average optimal to expedite the repair of an arriving defective part at time t if and only if* $X_1(t) \geqslant T(\mathbf{X}^{(-1)}(t), Y(t))$. *Furthermore these threshold levels* $T(\mathbf{x}^{(-1)}, y)$ *satisfy the property in Lemma 1, i.e.,* $\Delta_i T(\mathbf{x}^{(-1)}, y) \leqslant 0$ *for* $i = 2, \ldots, m$.

Theorem 2 also answers a question and conjecture posed by Song and Zipkin (2009, p. 371): "Are there any systems for which some policy of the form above is in fact optimal?" The policy Song and Zipkin (2009) propose is exactly the policy described in Theorem 2 for the special case that $m = 1$. For $m \geqslant 2$ one obtains a generalized form of this policy.

4.4. Infinite Horizon Discounted Optimal Expediting Policies

The same policy structure results hold for the case where we are interested in the infinite horizon discounted cost criterion. In this case we wish to solve

$$\inf_{\pi\in\Pi} \mathbb{E}^{\pi}_{(\mathbf{x},y)}\left[\int_0^\infty e^{-\beta t}(c_p(\mathbf{X}\mathbf{e}^{\mathrm{T}}, Y(t)) + \lambda_{Y(t)} c_e \mathbf{1}^{\pi}(\mathbf{X}(t), Y(t)))\, dt\right],$$

where $\beta > 0$ is the discount rate and $\mathbb{E}^{\pi}_{(\mathbf{x},y)}$ is the conditional expectation given control policy π and initial state $(\mathbf{x}, y)$. Proposition 1 continues to hold with $p\,\mathbb{E}[L_r]$ replaced by the expected discounted penalty costs over an interval of length L_r:

$$\mathbb{E}_{L_r}\left[\int_0^{L_r} p e^{-\beta t}\, dt\right] = \frac{p}{\beta} - \frac{p}{\beta}\mathbb{E}[e^{-\beta L_r}] = \frac{p}{\beta}\left(1 - \prod_{i=1}^{m}\frac{\mu_i}{\mu_i + \beta}\right).$$

The last equality holds by observing that $\mathbb{E}[e^{-\beta L_r}]$ is the Laplace-Stieltjes transform of a sum of exponential random variables. The MDP recursion can be written in exactly the same manner as before except that $\mathbb{T}_{\text{cost}}$, needs to be changed to $\mathbb{T}^{\beta}_{\text{cost}}\colon \mathcal{V} \to \mathcal{V}$ with

$$\mathbb{T}^{\beta}_{\text{cost}} f(\mathbf{x}, y) = \frac{c_p(\mathbf{x}\mathbf{e}^{\mathrm{T}}, y)}{\Lambda + \beta} + \frac{\Lambda}{\Lambda + \beta} f(\mathbf{x}, y).$$

It is readily verified that $\mathbb{T}^{\beta}_{\text{cost}}$ propagates the same properties as $\mathbb{T}_{\text{cost}}$, that is, if $f \in \mathcal{F}$ then also $\mathbb{T}^{\beta}_{\text{cost}} f(\mathbf{x}, y) \in \mathcal{F}$. With this change, it is easy to verify that Theorem 2 still holds, again with $p\,\mathbb{E}[L_r]$ changed to $p/\beta - (p/\beta)\,\mathbb{E}[e^{-\beta L_r}]$. Therefore we omit the proof of the following theorem.

Theorem 3. *Consider the infinite horizon discounted cost criterion for model* $\mathfrak{M}(S)$ *with discount rate* β. *If* $c_e \geqslant p/\beta - (p/\beta)\,\mathbb{E}[e^{-\beta L_r}] = p/\beta - (p/\beta)\prod_{i=1}^m(\mu_i/(\mu_i+\beta))$, *then it is* β*-discounted optimal to never expedite repair. If* $c_e < p/\beta - (p/\beta)\,\mathbb{E}[e^{-\beta L_r}] = p/\beta - (p/\beta)\prod_{i=1}^m(\mu_i/(\mu_i+\beta))$, *then there are state dependent threshold levels* $T(\mathbf{x}^{(-1)}, y) \in \mathbb{N}_0$ *such that it is* β*-discounted optimal to expedite repair at time t if and only if* $X_1(t) \geqslant T(\mathbf{X}^{(-1)}(t), Y(t))$. *Furthermore these threshold levels* $T(\mathbf{x}^{(-1)}, y)$ *satisfy the property in Lemma 1, i.e.,* $\Delta_i T(\mathbf{x}^{(-1)}, y) \leqslant 0$ *for* $i = 2, \ldots, m$.

4.5. Stage Dependent Holding Costs

The holding cost rate for a part is assumed to be h regardless of the position of that part in the repair pipeline. In many settings, value is added to items as they progress through the repair system. We will model this by considering an alternate system that is identical to the original system in all respects except that holding cost is charged as follows: Parts in stage i of the additional regular repair time incur holding cost at a rate $\hat{h}_i$ ($\hat{h}_i \leqslant \hat{h}_{i+1}$). Parts in the pipeline following additional regular repair incur holding cost at rate $\hat{h}_e$, and parts on-hand incur holding cost at rate $\hat{h}_{\text{OH}}$. Let us further denote the penalty cost rate by $\hat{p}$ and expediting cost by $\hat{c}_e$ for the alternate system. Finally let $\hat{C}_\pi(S)$ denote the average cost rate of this system under expediting policy π and turn-around stock S, including holding cost rate. The average cost rate for our original system (also including holding cost rate) are then denoted by $C_\pi(S)$. Our model can be used to optimize $\hat{C}_\pi(S)$ because it is related to $C_\pi(S)$ by the following transformation (proof is in Section EC.2.5):

Proposition 2. *If the cost parameters of the original system are given by*

$$h = \hat{h}_{\text{OH}}, \quad c_e = \hat{c}_e + \sum_{i=1}^{m}\frac{\hat{h}_{\text{OH}} - \hat{h}_i}{\mu_i}, \quad \textit{and} \quad p = \hat{p} + \hat{h}_{\text{OH}},$$

then

$$\hat{C}_\pi(S) = C_\pi(S) - (\hat{h}_{\text{OH}} - \hat{h}_e)\bar{\lambda} l_e - \bar{\lambda}\sum_{i=1}^{m}\frac{\hat{h}_{\text{OH}} - \hat{h}_i}{\mu_i},$$

where $\bar{\lambda} = \lim_{t\to\infty}\sum_{y\in\Theta}\mathbb{P}(Y(t) = y)\lambda_y$ *is the average demand rate.*

5. Optimal Turn-Around Stock Levels

In this section, we focus on the joint optimization of the turn-around stock S and the expediting policy. We start with the average cost criterion and address the infinite horizon discounted cost criterion with multiple purchasing opportunities in Section 5.1.

Unfortunately, $g^*(S)$ is not convex in S so that optimization of $C(S)$ is hard and requires enumeration. The example in Figure 3 shows that $C(S)$ need not even be uni-modal. In Section 7.4, we present numerical work for instances as they are typically encountered in practice. For all these instances, $g^*(S)$ is convex and $C(S)$ is uni-modal. In fact, a cursory look at Figure 3(a) does not immediately reveal that $C(S)$ is not convex. This is typical for all counterexamples that we have found.

The optimal turn-around stock, S^*, can be found by enumeration and the next proposition gives a property that can be used to terminate an enumerative search (the proof can be found in EC.2.6).

Proposition 3. *If $C(S) \leqslant h$ for some $S \in \mathbb{N}_0$, then $S^* \leqslant S$.*

It is also assuring to note that $V_n(\mathbf{x}, y \mid S)$ is submodular with respect to x_1 and S so that the need for expediting decreases as the turn-around stock increases. In case $|\Theta| = 1$ and $m = 1$, this fact can be used to efficiently jointly optimize the turn around stock S and the expediting policy which is characterized by a single scalar threshold; see Song and Zipkin (2009). The following proposition makes this formal. Its proof is in Section EC.2.7.

Proposition 4. *Let $T_S(\mathbf{x}^{(-1)}, y)$ denote the expediting threshold that is average optimal under a turn-around stock level of S at $(\mathbf{x}, y) \in \mathcal{S}$. Then $T_S(\mathbf{x}^{(-1)}, y) \leqslant T_{S+1}(\mathbf{x}^{(-1)}, y)$ for all $(\mathbf{x}, y) \in \mathcal{S}$ and $S \in \mathbb{N}_0$, that is, when the turn-around stock increases, the need for expediting decreases.*

5.1. Multiple Purchasing Opportunities

So far, we have considered only a single purchasing opportunity at the beginning of the horizon. Under the infinite horizon β-discounted setting, this assumption can be relaxed to allow for multiple purchasing opportunities. In this section, we show how to do this for the setting with two purchasing opportunities; see Figure 4. The second purchasing opportunity is N decision epochs after the initial purchasing opportunity. (Since the length of a decision epoch can be tuned by the uniformization constant, N can be chosen so that the second purchasing opportunity is arbitrarily close to a specific time.) We let h_1 denote the (discounted) price of a component purchased initially, and we let h_2 denote the discounted price of purchasing a component at the second purchasing opportunity. We initially purchase S_1 and at the second purchasing opportunity, we increase the turn-around stock to $S_2 \geqslant S_1$. The total discounted cost of purchasing components is then given by $h_1 S_1 + h_2(S_2 - S_1)$.

The optimal expected discounted cost associated with backlogging and expediting can now be computed as follows. The discounted cost incurred from the second purchasing opportunity until the end of the infinite horizon can be computed as explained in Section 4.4. Let us denote the resulting value function by $V_\infty^\beta(\cdot, \cdot \mid S_2)$. Next we use this value function as the terminal reward function to compute the value functions with n decision epochs to go; $0 < n \leqslant N$. Let us denote the expected discounted cost from the initial purchasing opportunity by $V_N^\beta(\mathbf{x}, y \mid S_1, S_2)$. Now for a specific initial condition, the total expected discounted

Figure 3. The cost-rate $C(S)$ for the problem instance with Poisson demand with rate $\lambda = 3.43$, $l_e = 8$, $m = 1$, $\mathbb{E}[L_r] = 15$, $p = 37.2$ and $c_e = 116$.

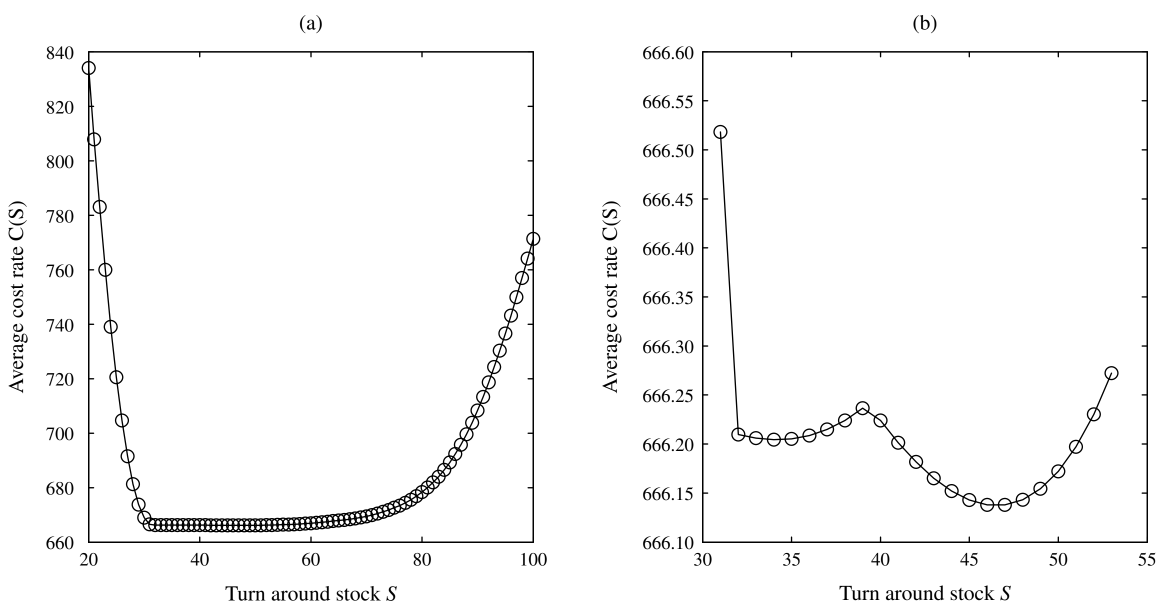

Notes. Although the fact that $C(S)$ is not unimodal is not immediately apparent from sub-figure (a), it is apparent from sub-figure (b). Any small deviation of any of the problem parameters will make $C(S)$ convex.

Figure 4. Situation with two purchasing opportunities.

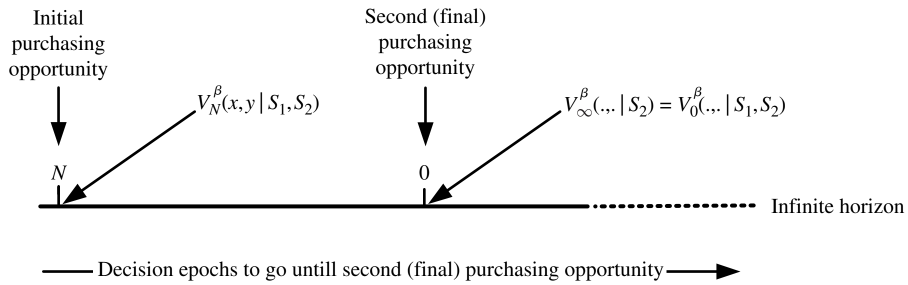

costs are given by $C^\beta(S_1, S_2) = h_1 S_1 + h_2(S_2 - S_1) + V_N^\beta(\mathbf{x}, y \mid S_1, S_2)$. Optimization of a single turn-around stock purchasing opportunity is already difficult, so the optimization of $C^\beta(S_1, S_2)$ will have to rely on enumeration. This idea can easily be extended to multiple purchasing opportunities.

6. Slowly Fluctuating Environments

Slay and Sherbrooke (1988) provide empirical evidence that demand for repairable parts is approximately a Poisson process for short periods of time, but that the demand intensity fluctuates slowly over time. In this section, we study the behavior of optimal expediting policies when demand fluctuations are slow. To this end, we will define a system denoted MMP(τ) for each scalar $\tau > 0$. Relative to our original model, the environment process $Y(t)$ is replaced in MMP(τ) with the scaled process $Y(\tau t)$. Note that if $Y(t)$ has generator $\mathbf{Q}$, then $Y(\tau t)$ has generator $\tau\mathbf{Q}$. Note further that the scalar τ scales the speed at which the environment fluctuates relative to the rest of the model: When $0 < \tau < 1$ the environment fluctuates slower than nominally. The average cost rate for MMP(τ) under policy π is denoted by $g^\pi_{\text{MMP}(\tau)}$, the optimal policy by $\pi^*_{\text{MMP}(\tau)}$, and $g^*_{\text{MMP}(\tau)} = g^{\pi^*_{\text{MMP}(\tau)}}_{\text{MMP}(\tau)}$. We also introduce the system P(y) for $y \in \Theta$, which is the original system except that demand is a stationary Poisson process with rate λ_y. As before, we let $g^\pi_{\text{P}(y)}$ denote the average cost rate of P(y) under policy π, $\pi^*_{\text{P}(y)}$ the optimal policy, and $g^*_{\text{P}(y)} = g^{\pi^*_{\text{P}(y)}}_{\text{P}(y)}$. Other than this, we use the same notation for all the parameters of models MMP(τ) and P(y). Furthermore we let $\lambda_{\min} = \min_{y\in\Theta} \lambda_y$, $y_{\min} = \arg\min_{y\in\Theta} \lambda_y$, $y_{\max} = \arg\max_{y\in\Theta} \lambda_y$, $\mathbb{P}(Y = y) = \lim_{t\to\infty} \mathbb{P}(Y(t) = y)$, and $\bar\lambda = \sum_{y\in\Theta} \mathbb{P}(Y = y)\lambda_y$. Note that $\bar\lambda$ is independent of τ because $Y(\tau t)$ has the same stationary distribution for any finite $\tau > 0$. One valid interpretation of a Markov modulated Poisson process is that there is a Poisson process for each $y \in \Theta$ and we only count an arrival of process $y \in \Theta$ at time t if $Y(t) = y$. Under this interpretation, we let $P^y_{t_1, t_2}$ denote the number of events generated by the Poisson process corresponding to state $y \in \Theta$ in the time interval $(t_1, t_2]$, regardless of $Y(t)$.

We will show in this section that when demand fluctuates arbitrarily slowly (i.e., as $\tau \downarrow 0$), the cost of any policy π for MMP(τ) converges to a weighted sum of costs for systems P(y) under the same policy for each $y \in \Theta$. We start our analysis with the following lemma whose proof is in Section EC.2.8.

Lemma 3. *Within model* MMP(τ), *for any* $\tau > 0$, $c_p(x, y)$ *can be bounded as*

$$\mathbb{E}[(P^y_{t,t+l_e} - (S-x))^+] - (1 - e^{-q_y l_e \tau})K_1 \leqslant \frac{c_p(x, y)}{p} \leqslant \mathbb{E}[(P^y_{t,t+l_e} - (S-x))^+] + (1 - e^{-q_y l_e \tau})K_2$$

with K_1 *and* K_2 *nonnegative constants given by*

$$K_1 = \max_{x\leqslant M}\left(\mathbb{E}[(P^y_{t,t+l_e} - (S-x))^+] - \mathbb{E}[(P^{y_{\min}}_{t,t+l_e} - (S-x))^+]\right)$$
$$K_2 = \max_{x\leqslant M}\left(\mathbb{E}[(P^{y_{\max}}_{t,t+l_e} - (S-x))^+] - \mathbb{E}[(P^y_{t,t+l_e} - (S-x))^+]\right).$$

Observe that $p\,\mathbb{E}[(P^y_{t,t+l_e} - (S-x))^+]$ is the value of $c_p(x, y)$ for system P(y). Lemma 3 thus bounds $c_p(x, y)$ with respect to the value of $c_p(x, y)$ in a pure Poisson demand system.

In preparation for the main results of this section, we will now give an alternative description of system MMP(τ) under a given expediting policy π in terms of a collection of renewal reward processes. To this end let us introduce the set $\mathscr{S}^{\mathbf{x},\pi}_{\text{rec}}$ of all states for $\mathbf{X}(t)$ that are recurrent under policy π. Observe that this set is unique, because MMP(τ) is unichain, as established in the proof of Theorem 2, and independent of $Y(\tau t)$, because if state $(\mathbf{x}, y)$ is recurrent, then so is $(\mathbf{x}, y')$ for any $y, y' \in \Theta$. We construct a renewal reward process for each state $(\mathbf{x}, y) \in \mathscr{S}^{\mathbf{x},\pi}_{\text{rec}} \times \Theta$ as follows: A renewal occurs for state $(\mathbf{x}, y)$ when the state changes to $(\mathbf{x}, y)$ because of a change in the modulating chain of demand $Y(\tau t)$. We divide the length of a renewal cycle in two parts. The first part, denoted $U^{(\mathbf{x},y)}$, is the time until $Y(\tau t)$ changes state for the first time (after the renewal point). Note that $U^{(\mathbf{x},y)}$ is exponentially distributed with mean $(q_y\tau)^{-1}$. The second part is denoted $W^{(\mathbf{x},y)}$ and consists of the time from the first state change in $Y(\tau t)$ until $(\mathbf{x}, y)$ is entered again because of a change in $Y(\tau t)$ (the next renewal point). Observe that $W^{(\mathbf{x},y)}$ has a phase type distribution with finite mean.

The expected reward $R^{(\mathbf{x},y)}$ associated with a renewal are the expected expediting and penalty costs incurred during the first part, $U^{(\mathbf{x},y)}$:

$$R^{(\mathbf{x},y)} = \mathbb{E}^{\pi}_{(\mathbf{x},y)}\left[\int_{t=0}^{\infty} e^{-q_y\tau t}(c_p(\mathbf{X}(t)\mathbf{e}^{\mathrm{T}}, y) + c_e\lambda_y \mathbf{1}^{\pi}(\mathbf{X}(t), y))\,dt\right]. \quad (23)$$

Note that $R^{(\mathbf{x},y)}$ can also be interpreted as the total discounted cost for a system facing Poisson demand at rate λ_y with discount rate $q_y\tau$, and a penalty cost rate function that still assumes the original demand from the MMP(τ) system. A crucial observation is that we now have

$$g^{\pi}_{\mathrm{MMP}(\tau)} = \sum_{y\in\Theta}\sum_{\mathbf{x}\in\mathcal{S}^{\mathbf{x},\pi}_{\mathrm{rec}}} \frac{R^{(\mathbf{x},y)}}{\mathbb{E}[U^{(\mathbf{x},y)}]+\mathbb{E}[W^{(\mathbf{x},y)}]}, \quad (24)$$

by the renewal reward theorem (e.g., Ross 1996, Theorem 3.6.1). With these preliminaries, we will state and prove the main results of this section.

Theorem 4. *In system* MMP(τ), *for any policy* π,

$$g^{\pi}_{\mathrm{MMP}(\tau)} \to \sum_{y\in\Theta}\mathbb{P}(Y=y)g^{\pi}_{\mathrm{P}(y)}, \quad \textit{as } \tau\downarrow 0,$$

where policy π *for system* MMP(τ) *is mapped to policy* π *for system* P(y) *by making decisions for state* $(\mathbf{x}, y)$ *in* MMP(τ) *coincide with those for state* $\mathbf{x}$ *in system* P(y).

Proof. Rewriting (24) we have

$$g^{\pi}_{\mathrm{MMP}(\tau)} = \sum_{y\in\Theta}\underbrace{\sum_{\mathbf{x}\in\mathcal{S}^{\mathbf{x},\pi}_{\mathrm{rec}}}\frac{\mathbb{E}[U^{(\mathbf{x},y)}]}{\mathbb{E}[U^{(\mathbf{x},y)}]+\mathbb{E}[W^{(\mathbf{x},y)}]}}_{\text{term 1}}\cdot\underbrace{\frac{R^{(\mathbf{x},y)}}{\mathbb{E}[U^{(\mathbf{x},y)}]}}_{\text{term 2}}. \quad (25)$$

We will proceed to analyze terms 1 and 2 above as $\tau\downarrow 0$. Since the occupancy distribution coincides with the limiting distribution for a Markov chain (e.g., Kulkarni 1999, Theorem 6.9) we have for term 1 that

$$\sum_{\mathbf{x}\in\mathcal{S}^{\mathbf{x},\pi}_{\mathrm{rec}}}\frac{\mathbb{E}[U^{(\mathbf{x},y)}]}{\mathbb{E}[U^{(\mathbf{x},y)}]+\mathbb{E}[W^{(\mathbf{x},y)}]} = \mathbb{P}(Y=y) \quad (26)$$

for any $\tau > 0$ and in particular as $\tau\downarrow 0$.

Now consider term 2 from (25). Using Lemma 3 we have

$$\begin{aligned}\frac{R^{(\mathbf{x},y)}}{\mathbb{E}[U^{(\mathbf{x},y)}]} &\leqslant q_y\tau\,\mathbb{E}^{\pi}_{(\mathbf{x},y)}\left[\int_0^{\infty} e^{-q_y\tau t}(p\,\mathbb{E}[(P^y_{t,t+l_e}-(S-\mathbf{X}(t)\mathbf{e}^{\mathrm{T}}))^+] + p(1-e^{-q_y l_e\tau})K_2 + c_e\lambda_y\mathbf{1}^{\pi}(\mathbf{X}(t),y))\,dt\right]\\ &= q_y\tau\,\mathbb{E}^{\pi}_{(\mathbf{x},y)}\left[\int_0^{\infty} e^{-q_y\tau t}\big(p\,\mathbb{E}[(P^y_{t,t+l_e}-(S-\mathbf{X}(t)\mathbf{e}^{\mathrm{T}}))^+] + c_e\lambda_y\mathbf{1}^{\pi}(\mathbf{X}(t),y)\big)\,dt\right]\\ &\quad + q_y\tau\int_0^{\infty} e^{-q_y\tau t}p(1-e^{-q_y l_e\tau})K_2\,dt\\ &= q_y\tau\,\mathbb{E}^{\pi}_{(\mathbf{x},y)}\left[\int_0^{\infty} e^{-q_y\tau t}\big(p\,\mathbb{E}[(P^y_{t,t+l_e}-(S-\mathbf{X}(t)\mathbf{e}^{\mathrm{T}}))^+] + c_e\lambda_y\mathbf{1}^{\pi}(\mathbf{X}(t),y)\big)\,dt\right] + p(1-e^{-q_y l_e\tau})K_2. \end{aligned} \quad (27)$$

In the same manner, one may find the following lower bound for term 2:

$$\frac{R^{(\mathbf{x},y)}}{\mathbb{E}[U^{(\mathbf{x},y)}]} \geqslant q_y\tau\,\mathbb{E}^{\pi}_{(\mathbf{x},y)}\left[\int_0^{\infty} e^{-q_y\tau t}\big(p\mathbb{E}[(P^y_{t,t+l_e}-(S-\mathbf{X}(t)\mathbf{e}^{\mathrm{T}}))^+] + c_e\lambda_y\mathbf{1}^{\pi}(\mathbf{X}(t),y)\big)\,dt\right] - p(1-e^{-q_y l_e\tau})K_1. \quad (28)$$

Now we have

$$\lim_{\tau\downarrow 0}\frac{R^{(\mathbf{x},y)}}{\mathbb{E}[U^{(\mathbf{x},y)}]} = \lim_{\tau\downarrow 0} q_y\tau\,\mathbb{E}^{\pi}_{(\mathbf{x},y)}\left[\int_0^{\infty} e^{-q_y\tau t}\big(p\mathbb{E}[(P^y_{t,t+l_e}-(S-\mathbf{X}(t)\mathbf{e}^{\mathrm{T}}))^+] + c_e\lambda_y\mathbf{1}^{\pi}(\mathbf{X}(t),y)\big)\,dt\right] = g^{\pi}_{\mathrm{P}(y)} \quad (29)$$

for every $(\mathbf{x}, y)\in\mathcal{S}$. The first equality follows from (27), (28), and the squeeze theorem, the second from Puterman (1994, Corollary 8.2.5). Combining (29), (26), (25), and the fact that $\mathbf{\Theta}$ and $\mathcal{S}^{\mathbf{x},\pi}_{\mathrm{rec}}$ have finite cardinality completes the proof. □

This result is important: Slow demand fluctuations can be accounted for by considering a time averaged sum of stationary problems. Also note that the proof of this result does not use anything specific about our model other than the finite state and action space established by Proposition 1 and the bound on the penalty cost function in Lemma 3. Therefore, this result also holds for other finite state MDPs subject to MMP input and cost functions that either do not depend on the dynamics of $Y(\tau t)$ or for which the cost functions satisfy a condition similar to that in Lemma 3. The main result of this section is a direct consequence of Theorem 4 and its proof. The proof is straightforward and therefore omitted.

Theorem 5. *The cost of an optimal policy*

$$g^{*}_{\mathrm{MMP}(\tau)} \to \sum_{y\in\Theta}\mathbb{P}(Y=y)g^{*}_{\mathrm{P}(y)}, \quad \textit{as } \tau\downarrow 0.$$

Furthermore, let π^*_{P} *be the policy for the* MMP(τ) *system that takes the same decision in* $(\mathbf{x}, y)$ *as an optimal policy for system* $P(y)$ *in state* $\mathbf{x}$. *Then* $\lim_{\tau\downarrow 0} g^{\pi^*_P}_{\mathrm{MMP}(\tau)} = \lim_{\tau\downarrow 0} g^{*}_{\mathrm{MMP}(\tau)}$.

Theorem 5 implies that π^*_{P} is a near optimal heuristic expediting policy for sytems where $\max_{y\in\Theta} q_y \ll \mathbb{E}[L_r]^{-1}$.

6.1. Value of Anticipating Demand Fluctuations

We will now use Theorems 4 and 5 to show that approximating demand by a Poisson process when it is in fact an Markov modulated Poisson process can lead to arbitrarily bad results. To do this we introduce system BE(τ, p, λ) (BE stands for bad example) that is parameterized as follows:

$$h=1, \quad c_e=0, \quad \mathbf{Q}=\tau\begin{pmatrix}-1 & 1\\ 1 & -1\end{pmatrix}, \quad \boldsymbol{\lambda}=\begin{pmatrix}0\\ 2\lambda\end{pmatrix}. \quad (30)$$

Note that the optimal expediting policy for system BE(τ, p, λ) is to always expedite because $c_e = 0$. (Therefore we can (and will) leave L_r unspecified.) We let $g^*_{\mathrm{BE}(\tau,p,\lambda)}$ denote the optimal cost rate associated with backlogging and expediting for system BE(τ, p, λ). Let P(p, λ) denote a system that is identical to E(τ, p, λ, S) except that demand is a stationary Poisson process with mean λ. (Note that the mean demand rate in system BE(τ, p, λ) is also λ.) We let $g^{\mathrm{P}}_{\mathrm{BE}(\tau,p,\lambda)}$ denote the backlogging and expediting cost rate of system BE(τ, p, λ) under a policy that is optimal for system P(p, λ); that is under the policy that approximates the MMP demand process by a Poisson process with the same mean. Finally let $C_{\mathrm{BE}(\tau,p,\lambda)}(S)$ be the total cost associated with system BE(τ, p, λ), i.e., $C_{\mathrm{BE}(\tau,p,\lambda)}(S) = hS + g^*_{\mathrm{BE}(\tau,p,\lambda)}$. Similarly let $C^{\mathrm{P}}_{\mathrm{BE}(\tau,p,\lambda)}(S)$ be the total cost associated by applying the expediting policy that is optimal for a system with stationary Poisson demand at rate λ to system BE(τ, p, λ), i.e., $C^{\mathrm{P}}_{\mathrm{BE}(\tau,p,\lambda)}(S) = hS + g^{\mathrm{P}}_{\mathrm{BE}(\tau,p,\lambda)}$.

Proposition 5. *Approximating a Markov modulated Poisson demand process by a Poisson demand process with the same mean can lead to arbitrarily bad performance. In particular,*

$$\lim_{p\to\infty}\lim_{\lambda\to\infty}\lim_{\tau\downarrow 0}\frac{\inf_{S\in\mathbb{N}_0} C^{\mathrm{P}}_{\mathrm{BE}(\tau,p,\lambda)}(S)}{\inf_{S\in\mathbb{N}_0} C_{\mathrm{BE}(\tau,p,\lambda)}(S)}=\infty.$$

7. Heuristics

In the previous sections, we have analyzed exact and asymptotic solutions to our problem. However, finding the optimal solution involves repeatedly solving a MDP that suffers from the curse of dimensionality. Furthermore, the optimal expediting policy is rather intricate, depending on the entire vector of repair jobs that will not arrive within the expedited lead time. In this section, we describe several heuristics for our model that involve a heuristic expediting policy that is much easier to interpret and that does not impose the same computational burden. We call this expediting heuristic the world driven threshold (WDT) heuristic for reasons that will become clear later. We discuss this heuristic expediting policy in Section 7.1 and the special case of stationary Poisson demand models in Section 7.2. Section 7.3 presents several heuristics for the joint optimization of expediting policy and turn-around stock. These heuristics use the heuristic WDT expediting policy.

7.1. World Driven Threshold Policies

Computing the state dependent optimal threshold levels quickly becomes computationally prohibitive as m increases. A plausible heuristic policy is to aggregate all orders in $\mathbf{X}(t)$ and to put a threshold expediting level, $T(y)$, on their sum, $\mathbf{X}(t)\mathbf{e}^{\mathrm{T}}$. This threshold will then only depend on $Y(t)$ and so, borrowing the terminology of Zipkin (2000), we call such a policy a WDT policy. The WDT policy satisfies the monotonicity property in Theorem 2 that $\Delta_i T(\mathbf{x}^{(-1)}, y) \leqslant 0$. Indeed, observe that the thresholds ($T^{\mathrm{WDT}}(\mathbf{x}, y)$) of a WDT policy satisfy

$$\Delta_i T^{\mathrm{WDT}}(\mathbf{x}, y) = \begin{cases}-1, & \text{if } T(\mathbf{x}, y) > 0;\\ 0, & \text{otherwise.}\end{cases}$$

This is shown graphically in Figure 5, where the optimal thresholds are shown with the best WDT thresholds. As before, the most convenient way to interpret Figure 5 is to think of it as a switching curve: If x_1 is on or above the shown line for some x_2, then expedite the repair, otherwise do not expedite repair.

For $m > 1$, finding the best WDT policy is about as difficult as finding an optimal policy since the stationary distribution of $\mathbf{X}(t)$ under such a policy still requires the evaluation of an $m+1$ dimensional Markov chain. A notable exception, that we discuss in Section 7.2, occurs when demand is a stationary Poisson process, i.e., $|\Theta| = 1$. In general, for $|\Theta| > 1$ and $c_e < p\mathbb{E}[L_r]$, we propose the following heuristic way of finding a good WDT policy. Instead of working with the $(m+1)$-dimensional space, move to two-dimensional space by approximating L_r by a single exponential phase with the same mean $\mu_1 = \mu = 1/\mathbb{E}[L_r]$. Then we are left with a two-dimensional space for which we can easily solve the resulting MDP to optimality using any common algorithm to solve finite state and action space MDPs.

The WDT policies that result from this procedure are not necessarily optimal within the class of WDT policies and the computed cost is not exact but an approximation. Since the system under study is equivalent to a type of ample server queue, we may expect this approximation to be quite accurate. In the next subsection, we show that this approach is exact for Poisson demand and in Section EC.3, we provide numerical evidence that WDT policies that are found in this manner perform exceptionally well compared to optimal policies under Markov modulated Poisson demand.

7.2. Special Case: Poisson Demand

Now we consider the evaluation of WDT policies for the special case where $|\Theta| = 1$, so that we are dealing with stationary Poisson demand. We termed this system P(y) in Section 6 when demand intensity is λ_y. With a slight abuse

Figure 5. Optimal and heuristic policies.

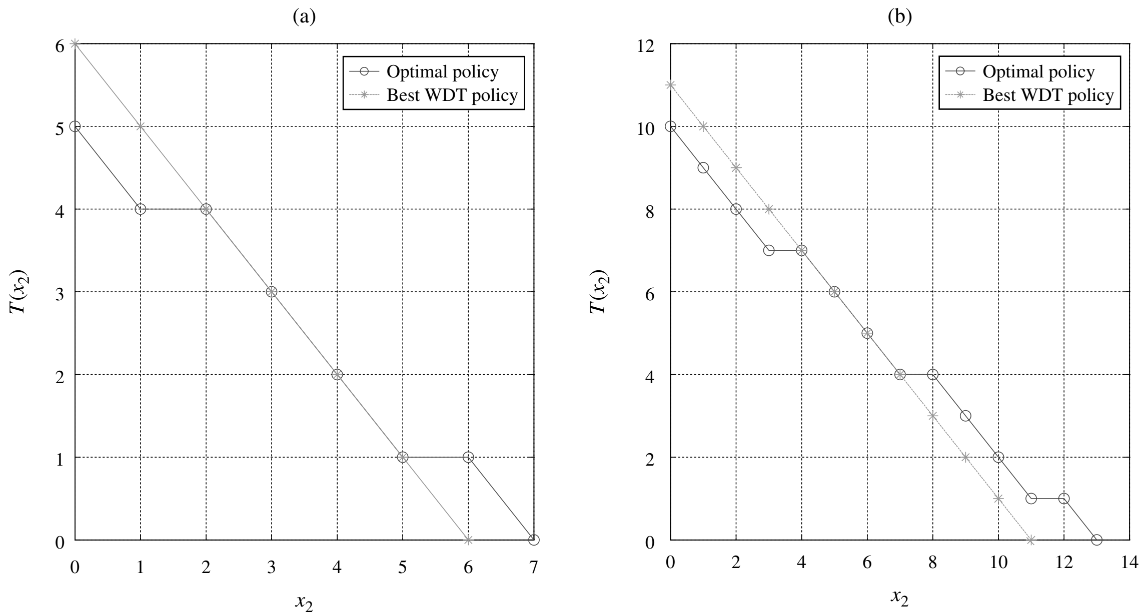

Notes. Part (a) shows the optimal state dependent threshold in conjunction with the best WDT policy for the case with $\lambda = 1$, $\mathbb{E}[L_r] = 4$, $l_e = 2$, $p = 10$, $c_e = 8$, $S = 8$ and $m = 2$. Part (b) shows the optimal state dependent threshold in conjunction with the best heuristic policy for the same case except $S = 12$.

of notation we now call such a system P(λ) when demand intensity is λ. In this case, the evaluation of a WDT policy can be done exactly for any distribution of L_r in closed-form using the results of Song and Zipkin (2009). (In this context, it might be appropriate to refer to a WDT policy simply as a threshold policy. For convenience, we use the name WDT policy also in this context.) Alternatively, one may simply observe that under such a policy, $\mathbf{X}(t)\mathbf{e}^{\mathrm{T}}$ has the same stationary distribution as the number of customers in an $M/G/c/c$ queue, where the number of servers c is set equal to the threshold T and the service time is distributed as L_r. In this equivalence, a customer being blocked from the queue because all T servers are busy corresponds to a repair being expedited because there are T or more parts that will not arrive within l_e. The average expediting and backorder penalty cost rate for such a policy with threshold level T and base-stock level S, $g_{\mathrm{P}(\lambda)}(S, T)$ is therefore given by

$$g_{\mathrm{P}(\lambda)}(S, T) = \lambda c_e B(T, \lambda \mathbb{E}[L_r]) + \sum_{x=0}^{T} c_p(x \mid S) \frac{(\lambda \mathbb{E}[L_r])^x / x!}{\sum_{k=0}^{T} (\lambda \mathbb{E}[L_r])^k / k!}, \quad (31)$$

where $B(c, \rho) = (\rho^c / c!)/(\sum_{k=0}^{c} \rho^k / k!)$ is the Erlang loss function with c servers and traffic intensity ρ, and $c_p(x \mid S) = c_p(x, 1 \mid S)$. Expression (31) also reveals that the performance of a WDT policy is insensitive to the distribution of L_r for the special case of Poisson demand. This insensitivity does not hold for Markov modulated Poisson demand. However, as a direct consequence of Theorem 4, this insensitivity does hold for system MMP(τ) as $\tau \downarrow 0$. In Section EC.3, we provide evidence that the performance evaluation of a WDT policy is nearly insensitive to the exact distribution of L_r for Markov modulated Poisson demand processes.

7.3. Joint Optimization Heuristics

We propose the following heuristics for the joint optimization of stocking level and expediting policy using the WDT expediting policy of Section 7.1.

- *E-WDT heuristic*: The E-WDT heuristic greedily optimizes $C(S)$ where we approximate L_r as being exponentially distributed by setting $m = 1$ and $\mu_1 = \mu = 1/\mathbb{E}[L_r]$. This simple heuristic ignores that L_r may not be exponential and that $C(S)$ may not be convex.
- *POIS-ASYMP heuristic*: Theorem 5 shows that that as fluctuations occur arbitrarily slowly, the optimal expediting policy when $Y(t) = y$ coincides with the expediting policy that is optimal for P(y). Therefore the POIS-ASYMP heuristic finds S and $T(y)$ for all $y \in \Theta$ by greedily minimizing

$$C_{\text{POIS-ASYMP}}(S, T(1), \ldots, T(|\Theta|)) = hS + \sum_{y \in \Theta} \mathbb{P}(Y = y) g_{\mathrm{P}(\lambda_y)}(S, T(y)),$$

where $g_{\mathrm{P}(\lambda_y)}(S \mid T(y))$ is the cost of applying a WDT expediting policy with threshold $T(y)$ in a system with Poisson demand with rate λ_y; see (31). Since $C_{\text{POIS-ASYMP}}(S, T(1), \ldots, T(|\Theta|))$ has a closed-form, this is easy to implement. The expediting policy is the WDT policy characterized by $T(y)$ for $y \in \Theta$.
- *POIS-AVG heuristic*: An easy, but naive, heuristic is to ignore entirely that demand fluctuates. The POIS-AVG heuristic does exactly this by greedily minimizing

Table 1. Average and maximum percentage optimality gaps for different heuristics.

Heuristic	E-WDT	POIS-ASYMP	POIS-AVG	POIS-MAX
AVG	0.11	0.33	12.81	11.66
MAX	0.76	2.37	68.1	26.35

$C_{\text{P}(\bar{\lambda})}(S, T)$ with $\bar{\lambda} = \sum_{y\in\Theta} P(Y = y)\lambda_y$. The WDT expediting policy then has a single threshold T independent of $Y(t)$. Proposition 5 shows that this heuristic can turn out arbitrarily bad. Its most important merit is its simplicity.

• *POIS-MAX heuristic*: Another easy, and seemingly prudent, policy is to account for demand fluctuations by optimizing for peak demand. The POIS-MAX heuristic does this by greedily minimizing $C_{\text{P}(\lambda_{\max})}(S, T)$. Here too, the expediting policy has a single threshold T independent of $Y(t)$.

7.4. Numerical Results

We report extended numerical result in Section EC.3, including a detailed description of our test-bed of 512 problem instances, and simulation results for different distributional assumptions on L_r. This section only reports aggregate statistics of percentage optimality gaps of the heuristics in Section 7.3 in Table 1. The E-WDT heuristic is the most sophisticated and performs very well with an average gap of 0.11% and maximum gap of only 0.76%. Nevertheless, the POIS-ASYMP heuristic performs almost as well and requires only the optimization of a closed-form expression. The POIS-AVG and POIS-MAX heuristics perform quite poorly. Proposition 5 already suggested that POIS-AVG can perform quite poorly and optimality gaps up to 68.1% confirm this. This shows that there is great value in leveraging knowledge about demand fluctuations when making expediting and stocking decisions.

Finally, the results show that although the optimal policy uses information about the progress of orders in the pipeline, the E-WDT and POIS-ASYMP heuristics that ignore this information perform quite well. This suggests that order progress information has limited value beyond knowing whether an order will arrive within l_e time units. At the same time, the value of information about demand fluctuations that is ignored by the POIS-AVG and POIS-MAX heuristics appears to be very important for good performance.

8. Conclusion

In this paper, we have considered the joint problem of finding the best turn-around stock and expediting policy for repairables that experience fluctuating demand. We have confirmed a conjecture by Song and Zipkin (2009) regarding the form of the optimal expediting policy and proposed the WDT expediting policies that perform well and are computationally tractable. We have shown that when demand fluctuates slowly, the performance of the system can be written as a weighted sum of the performance of systems facing stationary Poisson demand. Numerical and analytical results show that there is great value in leveraging knowledge about demand fluctuations when making expediting and stocking decisions.

Supplemental Material

Supplemental material to this paper is available at http://dx.doi.org/10.1287/opre.2016.1498.

Acknowledgments

The authors thank the referees and the associate and department editor for their constructive comments that improved the content and presentation of the paper. The authors thank Bob Huisman of NedTrain for useful discussions on the real-life situation that led to the model in this paper and Martijn van Aspert for finding the counterexample in Section 5. The second author gratefully acknowledges the support of the Lloyd's Register Foundation (LRF). The first and second author thank NedTrain and the Netherlands Foundation for Scientific Research for funding this research.

References

Aldous D, Shepp L (1987) The least variable phase type distribution is Erlang. *Comm. Statist.: Stochastic Models* 3(3):467–473.
Alfredsson P, Verrijdt J (1999) Modeling emergency supply flexibility in a two-echelon inventory system. *Management Sci.* 45(10):1416–1431.
Artalejo JR, Gómez-Corral A, He QM (2010) Markovian arrivals in stochastic modelling: A survey and some new results. *SORT: Statist. Oper. Res. Trans.* 34(2):101–156.
Arts J, Van Vuuren M, Kiesmüller GP (2011) Efficient optimization of the dual index policy using Markov chains. *IIE Trans.* 43(8):604–620.
Basten RJI, Van Houtum GJ (2014) System-oriented inventory models for spare parts. *Surveys Oper. Res. Management Sci.* 19(1):34–55.
Feng Q, Sethi SP, Yan H, Zhang H (2006) Are base-stock policies optimal in inventory problems with multiple delivery modes? *Oper. Res.* 54(4):801–807.
Fischer W, Meier-Hellstern K (1992) The Markov-modulated Poisson process (MMPP) cookbook. *Performance Eval.* 18(2):149–171.
Fukuda Y (1964) Optimal policies for the inventory problem with negotiable leadtime. *Management Sci.* 10(4):690–708.
Gaukler GM, Özer Ö, Hausman WH (2008) Order progress information: Improved dynamic emergency ordering policies. *Production Oper. Management* 17(6):599–613.
Koole G (1998) Structural results for the control of queueing systems using event-based dynamic programming. *Queueing Systems* 30(3):323–339.
Koole G (2004) Convexity in tandem queues. *Probab. Engrg. Informational Sci.* 18(1):13–31.
Koole G (2006) Monotonicity in Markov reward and decision chains: Theory and applications. *Foundations Trends in Stochastic Systems* 1(1):1–76.
Kulkarni VG (1999) *Modeling, Analysis, Design, and Control of Stochastic Systems* (Springer, New York).
Minner S (2003) Multiple supplier inventory models in supply chain management: A review. *Internat. J. Production Econom.* 81–82:265–279.
Moinzadeh K, Schmidt CP (1991) An $(S-1, S)$ inventory system with emergency orders. *Oper. Res.* 39(3):308–321.
Muckstadt JA (2005) *Analysis and Algorithms for Service Part Supply Chains* (Springer, Berlin).
Pérès F, Grenouilleau JC (2002) Initial spare parts supply of an orbital system. *Aircraft Engrg. Aerospace Tech.* 74(3):252–262.
Puterman ML (1994) *Markov Decision Processes: Discrete Stochastic Dynamic Programming* (John Wiley & Sons, New York).

Ross SM (1996) *Stochastic Processes*, 2nd ed. (John Wiley & sons, New York).

Rustenburg WD, van Houtum GJ, Zijm WHM (2001) Spare parts management at complex technology-based organizations: An agenda for research. *Internat. J. Production Econom.* 71(1—3):177–193.

Sheopuri A, Janakiraman G, Seshadri S (2010) New policies for the stochastic inventory control problem with two supply sources. *Oper. Res.* 58(3):734–745.

Sherbrooke CC (1968) METRIC: A multiechelon technique for recoverable item control. *Oper. Res.* 16(1):122–141.

Sherbrooke CC (2004) *Optimal Inventory Modeling of Systems: Multiechelon Techniques*, 2nd ed. (Wiley, Hoboken, NJ).

Slay FM, Sherbrooke C (1988) The nature of the aircraft component failure process. Technical Report IR701R1, Logistics Management Institute, Washington D.C.

Song JS, Zipkin P (1993) Inventory control in a fluctuating demand environment. *Oper. Res.* 41(2):351–370.

Song JS, Zipkin P (2009) Inventories with multiple supply sources and networks of queues with overflow bypasses. *Management Sci.* 55(3): 362–372.

Veeraraghavan S, Scheller-Wolf A (2008) Now or later: A simple policy for effective dual sourcing in capacitated systems. *Oper. Res.* 56(4): 850–864.

Verrijdt J, Adan I, de Kok T (1998) A trade off between emergency repair and inventory investment. *IIE Trans.* 30(2):119–132.

Whittmore AS, Saunders SC (1977) Optimal inventory under stochastic demand with two supply options. *SIAM J. Appl. Math.* 32(2): 293–305.

Zipkin PH (2000) *Foundations of Inventory Management* (McGraw-Hill, New York).

Joachim Arts is an assistant professor at Eindhoven University of Technology. His research interests are in the application of operations research to problems in maintenance, inventory, and logistics. He often collaborates with industrial partners such as NedTrain, ASML, and Philips. He was a visiting scholar at the MIT Sloan School of Management and the MIT Operations Research Center. He is a recipient of the EURO doctoral dissertation award and a Prins Bernhard Cultuurfonds Fellow.

Rob Basten is an assistant professor at Eindhoven University of Technology since 2014. His research focuses on after sales service for capital goods: maintenance policies and maintenance optimization, spare parts inventory control, and design of after sales service supply chains. The research is typically performed in cooperation with industry. He has been an assistant professor at University of Twente, a visiting scholar at the University of Texas at Austin in 2008 and the Rensselaer Polytechnic Institute in 2014.

Geert-Jan van Houtum is a professor of Maintenance and Reliability at the School of Industrial Engineering, Eindhoven University of Technology. He is the scientific director of the Beta Research School for Operations Management and Logistics. His research is focused on spare parts management, production-inventory systems, maintenance optimization, and the effect of design decisions on the total cost of ownership of capital goods. Together with Bram Kranenburg, he wrote a book entitled *Spare Parts Inventory Control under System Availability Constraints*.