Inventory Management with Two Demand Streams: A Maintenance Application

Rob J.I. Basten
Jennifer K. Ryan

Beta Working Paper series 480

BETA publicatie	WP 480 (working paper)
ISBN	
ISSN	
NUR	804
Eindhoven	August 2015

Inventory Management with Two Demand Streams: A Maintenance Application

Rob J.I. Basten*
Eindhoven University of Technology,
P.O. Box 513, 5600 MB, Eindhoven, The Netherlands

Jennifer K. Ryan†
University of Nebraska-Lincoln,
1400 R Street, Lincoln, NE 68588

August 18, 2015

Abstract

We consider an inventory system in which there are two types of demands with different priorities. The low priority demand is observed before the ordering decision, and thus exhibits perfect advance demand information, while the high priority demand is observed after the ordering decision. This type of inventory system arises in the context of spare parts inventory management, where the parts are used in the maintenance of capital systems, such as trains or airplanes. In such a system, parts are required to perform maintenance, which may be preventive, i.e., planned maintenance in which parts are replaced in order to prevent future failures, or corrective, i.e., unplanned maintenance in which parts that have failed unexpectedly are replaced. The former type of maintenance is planned in advance and thus the spare parts required can be ordered to arrive just-in-time. The latter type of maintenance cannot be planned in advance and thus, in order to prevent excessive downtime due to unplanned failures, safety stock must be held to meet these demands. While it is possible to manage the spare parts inventory for these two demand streams separately, we propose to jointly control their inventory. Specifically, we study a periodic review inventory system with a single stocking point used to meet deterministic, low priority demands, as well as stochastic, high priority demands, for a single product. We characterize the structure of the optimal inventory control policy and we demonstrate that the inventory requirements for the stochastic demand stream can be reduced by combining it with the deterministic demand stream. Finally, since the optimal policy is time-consuming to find, we propose a myopic heuristic policy and we demonstrate that this policy performs very well using a series of numerical experiments.

Keywords: Inventory · Two demand streams · Maintenance

*r.j.i.basten@tue.nl; Corresponding author: +31 40 247 2384
†jennifer.ryan@unl.edu

1 Introduction and Motivation

There are many practical settings in which an inventory system must satisfy multiple demand streams, with differing characteristics, for a single product. For example, one stream may come from high-priority customers, while the other comes from low-priority customers. These demand streams may also differ in terms of when they are observed, with advance demand information being available for some streams. As a result, there has been a considerable amount of literature considering inventory control with multiple demand streams or advance demand information (see Section 2). In this paper, we consider a problem setting which, to the best of our knowledge, has not been considered in the literature. In our setting, one demand stream is both deterministic and low priority, while the other demand stream is both stochastic and high priority. This situation arises in a variety of practical situations. For example, a warehouse ships goods to retailers on a regular basis. This demand is known well in advance and delaying a shipment a little is typically not a problem since the retailers hold inventory themselves. However, sometimes one of the retailers is out of stock and requests an emergency shipment. This type of request occurs randomly from the perspective of the warehouse. Because of the shortage at the retailer, any delay in the shipment could result in lost sales at the retailer. Therefore, such emergency shipment requests have high priority. This situation is also very common in two-echelon spare parts networks, where local stocking points are supplied from a central stocking point, which can also perform emergency shipments.

Our own interest in the problem studied in this paper stems from the maintenance of capital assets. Consider a user, such as a military organization or a rail operator, that maintains its own capital assets and holds the spare parts inventories required for that maintenance. Alternatively, consider an original equipment manufacturer (OEM), e.g., a manufacturer of jet engines or of high-tech manufacturing equipment, that is responsible for the maintenance of the assets that it has sold to customers. In these settings, the capital assets will require maintenance throughout their life cycle in order to ensure availability and functionality. Typically, some of the required maintenance is performed preventively, i.e., it is planned in advance. However, inevitably, some of the required maintenance will be corrective, i.e., performed in response to an unexpected failure, and thus is unplanned. As a result, the spare parts inventory manager will face two demand streams: one that is planned (deterministic) and one that is unplanned (stochastic). In addition, these two demand streams will generally have differing priorities. Since the planned demand is generated by preventive maintenance schedules, the consequence of a delay due to part unavailability is the need to reschedule that maintenance. However, because the delayed maintenance was preventive, the capital asset will be able to continue to function for some time. Thus, the cost associated with delaying a planned maintenance will be relatively low. In contrast, the unplanned demand is generated by unexpected failures, which may lead to unexpected downtime for the capital asset. In this case, any delay in performing the unplanned maintenance can result in lost equipment usage or lost production. Thus, the cost associated with delaying an unplanned maintenance may be significant. Therefore, in the context of spare parts inventory management for the maintenance of capital assets, we see that the deterministic demand stream will generally have a low priority, while the stochastic demand stream will have a high priority.

In such a setting, the inventory required to meet the planned and unplanned demand streams can be controlled separately. However, doing so ignores the benefits of inventory pooling (see, e.g., Basten and Van Houtum, 2014, p.50). Therefore, we propose joint inventory control, i.e., spare parts that are ordered for planned maintenance can be used for unplanned maintenance, if necessary. By managing the inventories jointly, we can reduce the amount of safety stock that must be held for unplanned maintenance. While there is a large amount of literature on inventory management when there are two demand streams (see Section 2 for a review), to the best of our knowledge, only a few papers consider a combination of deterministic and stochastic stream demand streams. In those papers, the priorities are reversed compared with our model. As a result, the system behavior is quite different.

Specifically, we consider a single stocking point for a single item that faces two streams of demands: a low-priority deterministic stream and a high-priority stochastic stream. We consider a periodic review setting with a finite horizon. However, we show that our results also hold for an infinite horizon setting. At the start of each period, the planned maintenance demand is observed. Spare parts can then be ordered and the ordered parts arrive instantaneously. Next, the unplanned maintenance demand is observed. The available spare parts can then be allocated to planned and unplanned maintenance actions, given the fact that delaying planned maintenance is less expensive than delaying unplanned maintenance. At the end of the period, delay costs are charged on all maintenance that is being delayed and holding costs are charged on all parts on-hand. We initially assume a replenishment lead time equal to zero. In this case, the difference between planned and unplanned maintenance may seem small. However, the difference is crucial: we can order parts for planned maintenance to arrive just-in-time, while we need to hold safety stock for the unplanned maintenance. In Section 8, we extend our results to consider a positive replenishment lead time.

For this model setting, we characterize the structure of the optimal inventory control policy for three cases: when planned maintenance may not be delayed, when planned maintenance may be delayed at most once, and when planned maintenance may be delayed an unlimited number of times. In the first case, inventory control for the two demand streams is separated. In the second case, we find that a myopic policy (Heyman and Sobel, 1984, Chapter 3) is optimal and that the optimal safety stock level for the unplanned maintenance is a function of the amount of planned maintenance scheduled for that period. In other words, information on previous and future periods can be ignored when ordering. In addition, we demonstrate that the optimal safety stock level is non-increasing in the amount of planned maintenance. In the third case, i.e., when maintenance can be delayed an unlimited number of times, we formulate the problem as a dynamic program and demonstrate that a base stock policy is optimal. However, the optimal policy is no longer myopic, i.e., we must consider the effect of decisions made in the current period on future periods. Fortunately, we can bound these effects and, using these bounds, we can derive lower and upper bounds on the optimal safety stock level in the current period. In a set of numerical experiments, we demonstrate that using the upper bound as a heuristic myopic policy leads to the optimal policy in most of the cases. Finally, we extend our model to incorporate a deterministic, positive lead time. We find that, for the first two cases, the results derived for the case of zero lead time can be readily extended to a setting with a positive lead time.

In Section 2, we discuss the related literature and highlight the key contributions of our work. We present our model and problem description in Section 3 and we present some key observations on the model in Section 4. In Section 5, we consider the setting in which planned maintenance may not be delayed, which leads to separate stock control for both types of maintenance. Then, in Section 6, we consider the case in which planned maintenance may be delayed at most once and we demonstrate that joint stock control is optimal. We then consider, in Section 7, the case in which planned maintenance may be delayed an unlimited number of times. In Section 8, we extend the model to allow for a positive supply lead time. We then perform a numerical experiment in Section 9 and conclude in Section 10. We defer all proofs to the appendix, but we will sometimes outline the idea behind the proof in the main text.

2 Literature Review

Our model considers two demand streams with different delay costs, i.e., different priorities. Thus, our work is related to the literature considering inventory control with multiple demand classes. Most of that work considers two (or more) *stochastic* streams of demand (see, e.g., Veinott, 1965; Topkis, 1968; Dekker et al., 2002; Deshpande et al., 2003; Kranenburg and Van Houtum, 2008). Often, Poisson distributed demand is assumed, and excess demand can be lost, backordered, or lost for one (or more) streams, but backordered for other streams. Both continuous review and periodic review settings have been considered. One advantage of fulfilling multiple streams of demand from one stocking point is risk pooling. Another advantage is that any fixed costs associated with ordering can be shared across the demand streams.

Within the literature considering multiple demand streams, a few authors have studied systems with one stochastic stream and one deterministic stream. To the best of our knowledge, the deterministic stream is always assumed to have the higher priority (see, e.g. Sobel and Zhang, 2001; Frank et al., 2003). One motivating example for these models is a manufacturer that has entered into a long term supply contract with certain customers, while other (smaller) customers place occasional (and less predictable) orders. For this problem setting, the authors formulate a periodic review inventory model, with a fixed cost associated with ordering, under the assumption that the higher priority deterministic demand must be fulfilled completely. In this problem setting, the advantage of combining the two streams of demand, i.e., making ordering decisions jointly for the two streams, is that the fixed cost for ordering can be shared. In addition, it may be optimal to ration inventory, i.e., to not fill some low priority demand even when inventory is available, because doing so may enable the inventory manager to delay ordering, thus avoiding the fixed order costs (Frank et al., 2003).

In this model setting, where the deterministic demand stream has higher priority, if there is no fixed order cost, then there is no need to manage the inventory jointly for the two demand streams, i.e., the problem is separable. In contrast, in our model setting, the deterministic demand stream is the lower priority stream. As a result, the advantage of jointly managing the inventory is that we can reduce the backorder (or delay) costs by exchanging a high cost backorder for a low cost backorder. As a result, in our model we do not require the assumption of a fixed order cost in order to derive insights into the value of jointly managing the inventory.

One other stream of literature that is closely related to our work is the literature that considers advance demand information (ADI) in inventory management. In some models, all demands are announced a certain time in advance, referred to as the demand lead time, e.g., the first model considered in Hariharan and Zipkin (1995). However, it is also possible to assume that different orders have different demand lead times, as in the other models considered in Hariharan and Zipkin (1995). While Hariharan and Zipkin (1995) consider ADI in a continuous review setting, Gallego and Özer (2001) consider ADI in a periodic review setting. Finally, there also exist models in which the ADI is imperfect, meaning that orders are announced with a certain probability of materializing. See, for example, Tan et al. (2007). Our model differs from this previous literature on ADI in that the backorder costs for the two types of demand in our model are different. In addition, the combination of one demand stream that has perfect ADI plus one demand stream that has no ADI has not been considered in the literature, to the best of our knowledge.

3 Model Description and Formulation

We next describe the problem we study and our model formulation. Throughout the paper, we will use the maintenance setting in our terminology. However, as discussed above, our problem is more broadly applicable. We consider the inventory management of spare parts that are used to maintain a large set of assets, such as a fleet of jet engines. We consider a finite horizon setting consisting of T periods, indexed by $t \in \{0, \ldots, T\}$. We will show that the key results from our analysis also hold for an infinite horizon setting. We initially assume a zero lead time for ordering parts and we assume that the fixed ordering costs are negligible. In Section 8, we discuss the setting with a positive supply lead time. While we do not include a variable ordering cost in our model, doing so would not change our results, as we show in Observation 4.5. Holding costs are charged on parts that remain on hand at the end of the period, at $C^{\mathrm{h}} > 0$ per part.

For the fleet of assets, planned maintenance actions are scheduled according to an exogenous process, resulting in P_t planned maintenance actions being scheduled in period t. P_t is a stochastic variable with a general discrete probability distribution with support $\mathbb{N}_0$ ($\mathbb{N}_0 = \mathbb{N} \cup 0$, with $\mathbb{N} = \{1, 2, \ldots\}$). The actual value of P_t is observed at the beginning of period t, before ordering. In addition to planned maintenance, unplanned maintenance may be required. The number of unplanned actions that arises in a certain period is a stochastic variable, U_t, which has a general discrete probability distribution with support $\mathbb{N}_0$, which is stationary over the horizon. (Notice, however, that we do not require stationarity of the planned maintenance, P_t.) The realized value of U_t is observed only after parts have been ordered, which means that there is a crucial difference between planned and unplanned maintenance: parts required for planned maintenance can be ordered to arrive in time with certainty (and we therefore consider this stream to be deterministic), while safety stock is required to deal with unplanned maintenance. We use O_t to denote the total number of parts to order in period t. The number of parts on hand just prior to ordering in period t is denoted by I_t, while the number of parts on hand after the order has arrived, but before demand is filled, is denoted by Y_t, so that $Y_t = I_t + O_t$.

It is possible that there may be insufficient inventory available to fill all maintenance events with no delay. Thus, it is necessary to state whether, and by how much, maintenance can be delayed. We consider three cases: planned maintenance may never be delayed, planned maintenance may be delayed at most once, and planned maintenance may be delayed an unlimited number of times. We make no assumptions regarding how long the unplanned maintenance may be delayed. We use P_t^{d} (U_t^{d}) to denote the number of planned (unplanned) maintenance actions that have already been delayed (and are waiting to be serviced) at the beginning of period t. Delaying a planned maintenance action for one additional period costs C^{p}, where $C^{\mathrm{p}} > 0$, while delaying an unplanned maintenance action for one additional period costs C^{u}, where $C^{\mathrm{u}} > C^{\mathrm{p}}$. As discussed above, in general, delaying unplanned maintenance will be more expensive than delaying planned maintenance. We assume that, at the end of the horizon, parts are ordered to fulfill any remaining delayed maintenance and that left over inventory is ignored, i.e., the salvage value is zero.

In any period t, we define the *known demand* to be equal to the known planned maintenance demands occurring in period t, plus all delayed maintenance still waiting to be serviced at the start of period t, i.e., the known demand in period t is $P_t + U_t^{\mathrm{d}} + P_t^{\mathrm{d}}$. We will show in Section 4 that, in any period t, the number of parts on hand after ordering, Y_t, will be at least equal to the known demand in period t. Any additional stock on hand, which can be used to cover the (unknown) unplanned maintenance demands, will be referred to as the *safety stock*, denoted by S_t. Thus, we have $Y_t = P_t + U_t^{\mathrm{d}} + P_t^{\mathrm{d}} + S_t$. Notice that our use of the term safety stock is slightly different than the conventional usage. Conventionally, safety stock refers to the number of units of inventory in excess of the mean demand. However, in this paper we will use the term safety stock to refer to the number of units of inventory in excess of the known demand.

Our goal is to order spare parts and, upon arrival, to allocate these parts to maintenance actions such that we achieve the minimum average cost per period over the T period planning horizon.

Our objective function can thus be written as:

$$\underset{S_t,\ t\in\{1,\ldots,T\}}{\text{minimize}} \quad \frac{1}{T}\sum_{t=1}^{T} C(S_t \mid P_t, P_t^{\mathrm{d}}),$$

where $C(S_t \mid P_t, P_t^{\mathrm{d}})$ is the expected cost in period t, which we define in Sections 5, 6 and 7, for each of the three cases that we study. The expected cost in period t does not depend on the delayed unplanned maintenance, U_t^{d}, since we show in Section 4 that there exists an optimal policy in which all delayed unplanned maintenance is fulfilled immediately.

4 Observations

We next provide some initial observations that will aid in our analysis. We provide the intuition for these observations, but no formal proofs, since the results are obvious. The first two observations consider the allocation policy, which assigns on hand inventory to maintenance actions. The optimal allocation policy can be defined completely without performing any formal analysis. Thus, in the remainder of this paper we focus on determining the optimal ordering policy.

Observation 4.1. Under any optimal allocation policy, the following rules are used to allocate on hand inventory to maintenance actions:

(i) To the extent to which it is allowed, given the limits on delaying planned maintenance, first satisfy all unplanned maintenance and then satisfy all planned maintenance.

(ii) Spare parts will never be rationed, i.e., a period will never end with both positive on hand stock and a positive number of delayed maintenance actions.

The assumption that $C^{\mathrm{u}} > C^{\mathrm{p}}$ implies that it is less expensive to delay planned maintenance ($P_t + P_t^{\mathrm{d}}$) than to delay unplanned maintenance ($U_t + U_t^{\mathrm{d}}$), if there is insufficient stock to fulfill both. However, if there is a limit on the number of times a planned maintenance can be delayed (e.g., never in Section 5 or once in Section 6), then to ensure feasibility, we may first need to fill the planned maintenance that has reached the specified limit. For example, if planned maintenance can never be delayed, then we must first fill any planned maintenance. Or, if planned maintenance may be delayed at most once, then we must first fill any planned maintenance that has already been delayed once. Otherwise, the policy will not be feasible. This concludes part (i). Part (ii) holds due to the assumption that ordered parts arrive instantaneously. Thus, rather than holding back one part and incurring a delay cost for an unfilled maintenance action, it would be better to use the available part to fill a maintenance action and then order one more part in the next period, since ordered parts arrive instantaneously.

Observation 4.2. There exists an optimal allocation policy:

(i) Under which first all planned (unplanned) maintenance that has already been delayed and then all new planned (unplanned) maintenance is fulfilled.

(ii) Which is completely determined by Observation 4.1 and part (i) of the current observation.

Notice that if planned maintenance may never be delayed, then part (i) is not relevant for planned maintenance. If planned maintenance may be delayed at most once, then part (i) is trivial for planned maintenance. When planned maintenance may be delayed an unlimited number of times, the order in which the planned maintenance actions are fulfilled is irrelevant, because any planned maintenance action incurs the same per period delay costs, regardless of whether or not it has been delayed before. A similar argument applies to unplanned maintenance. However, in practice it would typically make the most sense to satisfy the oldest requests first (i.e., to use first come first served allocation). Therefore, when we need to assume a specific allocation policy, this is what we will assume.

Part (i) of Observation 4.2, along with Observation 4.1, completely determines the optimal allocation policy for each of the three cases that we consider. This completes part (ii). Therefore, the allocation policies that we will assume throughout the paper are as specified in these observations.

The remaining decision in each period is the safety stock level, S_t, or, equivalently, the amount to order, O_t. We now focus our attention on characterizing properties of the optimal ordering policy.

Observation 4.3. Under any optimal policy:

(i) There will be no (planned or unplanned) maintenance actions that are deliberately delayed indefinitely.

(ii) After ordering, the amount of stock on hand in period t is at least as large as the known demands in period t: $Y_t \geq P_t + P_t^{\mathrm{d}} + U_t^{\mathrm{d}}$.

Part (i) follows immediately from the fact that we have assumed that all delayed maintenance actions are fulfilled at the end of the horizon. However, when an infinite horizon is assumed, then delaying a planned (unplanned) maintenance action indefinitely (by not buying a part for it) leads to infinite cost for that action, since C^{p} (C^{u}) is strictly positive. A feasible policy that results in finite cost is to buy an additional spare part after a planned (unplanned) maintenance action has been delayed for the first time, and allocate that spare part to the delayed maintenance action. Given that delaying maintenance actions indefinitely is never optimal, it is best to fulfill any known demands as soon as possible, leading to part (ii): deliberately delaying a known demand for some period of time does not bring any advantage, but it does lead to additional holding costs.

The following observation states that unplanned maintenance actions will be delayed at most once. This is useful both for our analysis and for the practicality of the optimal ordering policy.

Observation 4.4. There exists an optimal policy under which an unplanned maintenance action will be delayed at most once.

We know from Observation 4.3, part (ii), that the amount of stock on hand after ordering is larger than the amount of delayed unplanned maintenance, i.e., $Y_t \geq U_t^{\mathrm{d}}$. Since we further know that there exists an optimal policy under which first all delayed unplanned maintenance is fulfilled, from Observation 4.2, part (ii), it holds that an unplanned maintenance action will be delayed at most once.

Observation 4.5. Incorporating positive ordering costs that are linear in the number of parts ordered will not change the optimal stocking or allocation policy.

Let us denote these ordering costs by C^{o} per part. According to Observation 4.3, all maintenance actions eventually receive the spare parts that they require. Thus, incorporating the ordering cost implies that, under any given policy, the average cost per period is increased by $\frac{C^{\mathrm{o}}}{T}\sum_{t=1}^{T}\left[P_t + U_t\right]$.

5 Planned Maintenance: No Delay Allowed

We are now ready to consider the case in which planned maintenance may not be delayed. This assumption leads to the following observation.

Observation 5.1. Under an optimal policy, stock control for both types of maintenance is performed separately, i.e., the safety stock level, S_t, does not depend on the amount of planned maintenance, P_t.

Under any feasible policy, P_t parts are ordered in period t to cover the planned maintenance. After arrival of the order and observing the number of unplanned maintenance actions, all of these parts are allocated to the planned maintenance actions, since these may not be delayed. Therefore, when determining the optimal safety stock level to cover unplanned maintenance, the amount of planned maintenance is irrelevant.

Delaying an unplanned maintenance action leads to delay cost of C^{u} per period. By Observation 4.4, an unplanned maintenance will be delayed at most once. Given the safety stock level, S_t, the order quantity is $O_t = P_t + U_t^{\mathrm{d}} + S_t - I_t$. The total costs per period are:

$$C(S_t \mid P_t) = \sum_{k=0}^{S_t} \mathbb{P}\left\{U_t = k\right\} C^{\mathrm{h}}(S_t - k) + \sum_{k=1}^{\infty} \mathbb{P}\left\{U_t = S_t + k\right\} C^{\mathrm{u}}k. \tag{1}$$

Here we have suppressed the notation P_t^{d} in the cost function since it is always zero in this case.

Because the distribution of the number of unplanned maintenance actions is stationary over time, and because both the purchase cost and salvage value for inventory remaining at the end of the horizon are assumed to be zero, the optimal safety stock level is stationary and can be determined myopically (Heyman and Sobel, 1984, Chapter 3). Therefore, the optimal safety stock level, denoted by S^*, for each period, is as specified in the following theorem.

Theorem 1. *S^* is the smallest S_t for which $(C^h + C^u)\sum_{k=S_t+1}^{\infty} \mathbb{P}\{U_t = k\} \leq C^h$.*

This is a well-known result. This safety stock level can be achieved exactly in each period by ordering $O_t = P_t + U_{t-1}$, i.e., in each period we order the planned demand plus the unplanned demands that were filled (or delayed) in the previous period. If we were to include variable ordering costs, as discussed in Observation 4.5, these results would still hold, as long as these variable ordering costs are also incurred on the parts that are ordered at the end of the horizon (recall that we have assumed that all delayed maintenance actions are fulfilled at the end of the horizon) and if any excess inventory can be salvaged for the same variable ordering cost.

Remark 5.1. Given that a stationary myopic policy is optimal for the finite horizon setting, the result also holds in an infinite horizon setting.

6 Planned Maintenance: At Most One Delay Allowed

Next, we consider the case in which planned maintenance may be delayed at most once. The allocation policy remains clear (see Section 4). We further know from Observation 4.3, part (ii), that after ordering we have $Y_t \geq P_t + P_t^{\mathrm{d}} + U_t^{\mathrm{d}}$ on hand. We can suppress P_t^{d} (and U_t^{d}) in the cost function because these values do not influence the cost, i.e., all delayed planned (and unplanned) maintenance at the start of period t will be filled during period t. Thus, no further holding costs or delay costs are incurred for these demands. The difference between Y_t and the known demands, $P_t + P_t^{\mathrm{d}} + U_t^{\mathrm{d}}$, is the safety stock, S_t. In other words, $O_t = P_t + P_t^{\mathrm{d}} + U_t^{\mathrm{d}} + S_t - I_t$. Notice that since we must have that $O_t \geq 0$, there are situations in which the feasible region for S_t is not all non-negative integer values, i.e., S_t may be constrained by a lower bound. For example, if $P_t + P_t^{\mathrm{d}} + U_t^{\mathrm{d}} = 0$, then $S_t \geq I_t$. We will see in Theorem 4 that this bound is never binding, i.e., that the optimal unconstrained value of S_t is always feasible.

The expected costs in period t are:

$$C(S_t \mid P_t) = \sum_{k=0}^{S_t} \mathbb{P}\{U_t = k\} C^{\mathrm{h}}(S_t - k) + \sum_{k=1}^{P_t} \mathbb{P}\{U_t = S_t + k\} C^{\mathrm{p}} k$$
$$+ \sum_{k=1}^{\infty} \mathbb{P}\{U_t = S_t + P_t + k\} (C^{\mathrm{u}} k + C^{\mathrm{p}} P_t). \quad (2)$$

The first term in the cost function represents the case in which the amount of unplanned maintenance is below the safety stock level, so that no maintenance is delayed, while some on hand inventory remains at the end of the period, incurring holding costs. The second term represents the case in which only planned maintenance is delayed, while the third term represents the case in which all planned maintenance and some unplanned maintenance is delayed.

Lemma 2 provides some useful structural properties for $C(S_t \mid P_t)$.

Lemma 2. *The cost function $C(S_t \mid P_t)$ has the following properties:*

(i) $C(S_t \mid P_t)$ is strictly convex in S_t.

(ii) $C(S_t \mid P_t)$ is strictly decreasing and strictly convex in P_t.

(iii) $C(S_t \mid P_t)$ is supermodular in P_t and S_t.

The fact that the costs are decreasing in P_t, as stated in part (ii), may be counterintuitive. However, recall that we ignore the ordering costs, i.e., we focus on the costs incurred due to the uncertainty in unplanned maintenance demands. Because we may delay planned maintenance, the stock that we order for planned maintenance acts as a buffer, i.e., for a given safety stock level, when there happens to be a lot of unplanned maintenance, some planned maintenance can be delayed rather than delaying unplanned maintenance, which would be much more costly. The supermodularity result in part (iii) of the lemma is used to prove Lemma 3. This result is useful because it implies that if the planned maintenance scheduled in a period t increases, then the optimal safety stock level decreases or remains the same. This result is also due to the fact that the stock ordered to fill planned maintenance acts as a buffer, i.e., delaying planned maintenance is less costly than delaying unplanned maintenance. Thus, if there is more planned maintenance that can potentially be delayed, freeing up stock to satisfy unplanned maintenance, then less safety stock is required.

Let us now assume that there is a certain period t in which there is no stock on hand (i.e., $I_t = 0$) and no delayed demand at the beginning of the period (i.e., $P_t^{\mathrm{d}} + U_t^{\mathrm{d}} = 0$). Period 1 is such a period by definition. These assumptions imply that any $S_t \geq 0$ is feasible. We further ignore any information that we have on future periods, thus making the problem myopic. Lemma 3 provides a first order condition for determining the optimal safety stock level, S_t^*, in period t, under these assumptions. The lemma also indicates that this optimal safety stock level is decreasing in the amount of planned maintenance in that period, P_t. To be more precise, the lemma indicates that when the amount of planned maintenance is increased (decreased) by 1 unit, the optimal safety stock level will remain the same or decrease (increase) by 1 unit. Because the optimal safety stock level is a function of the amount of planned maintenance, we write it as $S_t^*(P_t)$.

Lemma 3. *If $I_t = 0$ and $P_t^d + U_t^d = 0$, the optimal safety stock level, $S_t^*(P_t)$, has the following properties:*

(i) $S_t^*(P_t)$ is the smallest S_t for which:

$$(C^p + C^h) \sum_{k=S_t+1}^{\infty} \mathbb{P}\{U_t = k\} + (C^u - C^p) \sum_{k=S_t+P_t+1}^{\infty} \mathbb{P}\{U_t = k\} \leq C^h.$$

(ii) $S_t^*(P_t)$ *satisfies* $S_t^*(P_t) - 1 \leq S_t^*(P_t + 1) \leq S_t^*(P_t)$.

Although the optimal safety stock level, S_t^*, can no longer be written using a standard newsvendor equation, it is still myopic and it is straightforward to find, since $C(S_t \mid P_t)$ is convex in S_t.

Theorem 4 states that, in any period t, it is optimal to use the myopically optimal safety stock level, S_t^*, as defined in Lemma 3. The key to the proof is to demonstrate that S_t^* is feasible in every period t, i.e., implementing S_t^* will not require $O_t < 0$, even without the conditions required for Lemma 3.

Theorem 4. *In any period* t*, if ordering in all previous periods has been performed using the myopically optimal policy specified in Lemma 3, then* $I_t \leq P_t + S_t^*(P_t)$*. Therefore, it is optimal and feasible to order such that the on hand inventory after ordering is* $Y_t = P_t^d + U_t^d + P_t + S_t^*(P_t)$*, with* $S_t^*(P_t)$ *as defined in Lemma 3.*

If we assume that period $t = 1$ starts with $I_1 = 0$ and $P_1^{\mathrm{d}} + U_1^{\mathrm{d}} = 0$, then by Lemma 3, the myopic policy is optimal for period $t = 1$. By induction and Theorem 4, we thus have that the myopic policy is optimal for each period in the planning horizon.

Given that the optimal safety stock level in any period t satisfies the first order condition specified in Lemma 3, part (i), we see that $S_t^*(P_t)$ depends on t only through P_t. In other words, since the unit costs (C^{h}, C^{u} and C^{p}) and the probability distribution of U_t ($\mathbb{P}\{U_t = k\}$ for $k \in \mathbb{N}_0$) are stationary, we can write $S_t^*(P_t) = S^*(P_t)$ and thus the safety stock function is stationary. Notice that this implies that it is straightforward to extend these results to an infinite horizon setting.

Finally, Lemma 5 states that the optimal cost in period t is decreasing in the amount of planned maintenance in period t. The intuition is the same as that provided for Lemma 2.

Lemma 5. $C(S_t^*(P_t) \mid P_t)$ *is strictly decreasing and convex in* P_t*.*

7 Planned Maintenance: Unlimited Delay

Next, we consider the case in which planned maintenance may be delayed an unlimited number of times. In this case, the ordering decision in period t may influence periods $t' > t$. Thus, it becomes necessary to formulate the problem as a discrete-time Markov decision problem (MDP). The state in period t consists of the planned maintenance in period t and the delayed planned maintenance carried over to period t, i.e., $\mathcal{S} = \{(P_t, P_t^{\mathrm{d}}) \in \mathbb{N}_0^2\}$. We do not need to include the delayed unplanned maintenance, U_t^{d}, in the state due to the fact that there exists an optimal policy under which U_t^{d} will be filled in period t (see Observation 4.4).

The allocation policy is as defined in Observation 4.2, part (ii). Thus, the action space in period t is simply $\mathcal{A} = \{S_t \in \mathbb{N}_0\}$. Given S_t, the order quantity is just $O_t = P_t + P_t^{\mathrm{d}} + U_t^{\mathrm{d}} + S_t - I_t$. The cost function in period t can thus be written as:

$$C(S_t \mid P_t, P_t^{\mathrm{d}}) = \sum_{k=0}^{S_t} \mathbb{P}\{U_t = k\} C^{\mathrm{h}}(S_t - k) + \sum_{k=1}^{P_t + P_t^{\mathrm{d}}} \mathbb{P}\{U_t = S_t + k\} C^{\mathrm{p}} k \\ + \sum_{k=1}^{\infty} \mathbb{P}\left\{U_t = S_t + P_t + P_t^{\mathrm{d}} + k\right\} \left(C^{\mathrm{u}} k + C^{\mathrm{p}}(P_t + P_t^{\mathrm{d}})\right). \tag{3}$$

This single period cost function is quite similar to Equation (2). However, since we now allow planned maintenance to be delayed indefinitely, the ranges over which we are summing are modified slightly. The dynamic programming recursion can now be written as follows:

$$V_t(P_t, P_t^{\mathrm{d}}) = \mathbb{E}\left\{\min_{S_t \in \mathbb{N}_0} G_t(S_t \mid P_t, P_t^{\mathrm{d}})\right\},$$

with the expectation being over the future maintenance actions, since they are not yet observed in period t, and:

$$G_t(S_t \mid P_t, P_t^{\mathrm{d}}) = \sum_{k=0}^{S_t} \mathbb{P}\{U_t = k\} \left(C^{\mathrm{h}}(S_t - k) + V_{t+1}(P_{t+1}, 0)\right) \\ + \sum_{k=1}^{P_t + P_t^{\mathrm{d}}} \mathbb{P}\{U_t = S_t + k\} (C^{\mathrm{p}} k + V_{t+1}(P_{t+1}, k)) \\ + \sum_{k=1}^{\infty} \mathbb{P}\left\{U_t = S_t + P_t + P_t^{\mathrm{d}} + k\right\} \left(C^{\mathrm{u}} k + C^{\mathrm{p}}(P_t + P_t^{\mathrm{d}}) + V_{t+1}(P_{t+1}, P_t + P_t^{\mathrm{d}})\right). \tag{4}$$

The three cases can be explained as follows:

- $U_t \leq S_t$: nothing is delayed and period $t + 1$ starts with positive on hand stock;
- $S_t < U_t \leq S_t + P_t + P_t^{\mathrm{d}}$: only some planned maintenance is delayed and period $t + 1$ starts with zero on hand stock; and
- $U_t > S_t + P_t + P_t^{\mathrm{d}}$: all planned maintenance is delayed, as well as some unplanned maintenance, and period $t + 1$ starts with zero on hand stock.

We define $V_{T+1}(P_{T+1}, P_{T+1}^{\mathrm{d}}) \equiv 0$, since, as discussed in Section 3, we assume that, at the end of the horizon, parts are ordered to fill any remaining delayed maintenance (resulting in no holding or delay costs after the horizon), and leftover inventory has no salvage value. This implies that:

$$G_T(S_T \mid P_T, P_T^{\mathrm{d}}) = C(S_T \mid P_T, P_T^{\mathrm{d}}). \tag{5}$$

The order in which we present the results is very similar to that in the previous section. In the appendix, we first prove results for period T and then for all other periods t by induction. Due to the similarity of $G_T(S_T \mid P_T, P_T^{\mathrm{d}})$ with the cost function from the previous section, Equation (2), many of the results for period T follow (almost) directly from the results in the previous section. Furthermore, in all periods t, the results for P_t^{d} are the same as those for P_t (e.g., convexity of the cost rate function), which is easily seen from Equation (4). Lemma 6 provides some useful properties of $G_t(S_t \mid P_t, P_t^{\mathrm{d}})$; the intuition for these results is similar to that provided for Lemma 2.

Lemma 6. *$G_t(S_t \mid P_t, P_t^d)$ has the following properties:*

(i) $G_t(S_t \mid P_t, P_t^d)$ is strictly convex in S_t.

(ii) $G_t(S_t \mid P_t, P_t^d)$ is strictly decreasing and strictly convex in P_t and P_t^d.

(iii) $G_t(S_t \mid P_t, P_t^d)$ is supermodular in each of the combinations of P_t, P_t^d and S_t.

In what follows we will use the notation $S_t^*(P_t)$ or $S_t^*(P_t^{\mathrm{d}})$ to explicitly show the dependence of S_t^* on P_t or P_t^{d} whenever that is useful. In addition, we define $\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, P_{t+1}^{\mathrm{d}}) = V_{t+1}(P_{t+1}, P_{t+1}^{\mathrm{d}} + 1) - V_{t+1}(P_{t+1}, P_{t+1}^{\mathrm{d}})$, which represents the increase in costs incurred when an additional unit of planned maintenance is delayed in period t. We can now state Lemma 7, which provides properties of the optimal safety stock level.

Lemma 7. *The optimal safety stock level in period t, S_t^*, has the following properties:*

(i) S_t^* is the smallest S_t for which:

$$(C^p + C^h) \sum_{k=S_t+1}^{\infty} \mathbb{P}\{U_t = k\} + (C^u - C^p) \sum_{k=S_t+P_t+P_t^d+1}^{\infty} \mathbb{P}\{U_t = k\}$$
$$+ \sum_{k=S_t+1}^{S_t+P_t+P_t^d} \mathbb{P}\{U_t = k\} \left[\Delta_{P_{t+1}^d} V_{t+1}(P_{t+1}, k-1)\right] \leq C^h.$$

(ii) $S_t^*(P_t)$ satisfies $S_t^*(P_t) - 1 \leq S_t^*(P_t + 1) \leq S_t^*(P_t)$ and $S_t^*(P_t^d)$ satisfies $S_t^*(P_t^d) - 1 \leq S_t^*(P_t^d + 1) \leq S_t^*(P_t^d)$.

Analogous to Lemma 3, part (ii), in Section 6, Lemma 7, part (ii), implies that if the amount of (new or delayed) planned maintenance in period t increases by one, then the optimal safety stock level remains the same or decreases by one. This has a further, counterintuitive implication: Suppose the safety stock level in period t is increased by one unit. This may result in the optimal safety stock level for period $t+1$ also increasing by one unit due to the fact that the amount of delayed planned maintenance in period $t+1$ may (depending on the amount of unplanned maintenance that arises in period t) be reduced by one. Recall that the stock ordered to fill planned maintenance (whether delayed or not) serves as a buffer, i.e., can be used to fill unplanned maintenance, if needed. Thus, if the amount of delayed planned maintenance is decreased by one unit, we will order one less unit to fill that delayed planned maintenance, which implies that there will be less of a buffer that can be used to meet unplanned maintenance. Thus, the safety stock may need to be increased by one unit.

Lemma 8 provides some key properties of the optimal cost in period t as a function of P_t or P_t^{d}. Part (ii) also holds when P_t^{d} is replaced by P_t, but in the remainder of the paper we only use the result for P_t^{d}.

Lemma 8. *$V_t(P_t, P_t^d)$ has the following properties:*

(i) $V_t(P_t, P_t^d)$ is strictly decreasing and convex in P_t and P_t^d.

(ii)

$$-C^h + C^p \sum_{k=S_t^*(P_t^d)+1}^{\infty} \mathbb{P}\{U_t = k\} \leq \Delta_{P_t^d} V_t(P_t, P_t^d) \leq -C^h + (C^p + C^h) \sum_{k=S_t^*(P_t^d)}^{\infty} \mathbb{P}\{U_t = k\}.$$

The lemma states that the optimal cost-to-go in period t decreases when the amount of planned maintenance increases. This result is similar to the results of the previous section. The stock ordered to fill planned maintenance acts as a buffer and can be used to fill unplanned maintenance, if needed. However, while P_t is exogenously given, P_{t+1}^{d} is dependent on the decisions made in period t, as well as the amount of unplanned maintenance in period t. Part (ii) of the lemma indicates that the cost-to-go will decrease at most by C^{h} when P_{t+1}^{d} is increased by one. Combining this result with Lemma 7, part (i), provides lower and upper bounds on the optimal safety stock level, S_t^*, which we denote by S_t^{lb} and S_t^{ub}. These bounds are defined in Corollary 9, which we provide without further proof.

Corollary 9. *S_t^* satisfies $S_t^{lb} \leq S_t^* \leq S_t^{ub}$, where S_t^{lb} is the smallest value of S_t for which:*

$$C^p \sum_{k=S_t+1}^{S_t+P_t+P_t^d} \mathbb{P}\{U_t = k\} + (C^u + C^h) \sum_{k=S_t+P_t+P_t^d+1}^{\infty} \mathbb{P}\{U_t = k\} \leq C^h,$$

and S_t^{ub} is the smallest value S_t for which:

$$(C^p + C^h) \sum_{k=S_t+1}^{S_t+P_t+P_t^d} \mathbb{P}\{U_t = k\} + (C^u + C^h) \sum_{k=S_t+P_t+P_t^d+1}^{\infty} \mathbb{P}\{U_t = k\} \leq C^h.$$

Notice the similarity between the upper bound, S_t^{ub}, and the optimal safety stock level, $S_t^*(P_t)$, for the case in which planned maintenance may be delayed at most once, as specified in Lemma 3, part (i). The only difference between these two values is that in the condition defining $S_t^*(P_t)$, in Lemma 3, part (i), the ranges in the summations do not include P_t^{d}.

In addition, note that if the lower and upper bounds provided in Corollary 9 coincide, i.e., if $S_t^{\mathrm{lb}} = S_t^{\mathrm{ub}}$, then this is the optimal safety stock level, S_t^*, and thus the optimal base stock level can be found myopically. In the numerical experiments presented in Section 9, we will show that these bounds often do coincide. In addition, we will develop a heuristic solution procedure for the case in which the bounds do not coincide. Finally, we note that if we know P_{t+1}, then it is possible to get even tighter bounds on S_t^* by first bounding S_{t+1}^* and then using those bounds to get tighter bounds on $\Delta_{P_t^{\mathrm{d}}} V_t(P_t, P_t^{\mathrm{d}})$. This can be done recursively for more future periods, and although the resulting policy is not myopic, it can be found quite efficiently.

Finally, we conclude this section with Lemma 10, in which we compare the costs and safety stock levels for the three cases considered in the paper: planned maintenance may not be delayed, may be delayed at most once, and may be delayed an unlimited number of times.

Lemma 10. *The following results hold:*

(i) For a given safety stock level:

(a) The per period cost for the case in which planned maintenance can be delayed at most once is less than or equal to the per period cost for the case in which planned maintenance can never be delayed, with the difference in cost equal to:

$$-\sum_{k=1}^{P_t} \mathbb{P}\left\{U_t = S_t + k\right\}(C^u - C^p)k - \sum_{k=1}^{\infty} \mathbb{P}\left\{U_t = S_t + P_t + k\right\}(C^u - C^p)P_t.$$

If $P_t > 0$, then the cost when planned maintenance can be delayed once is thus strictly lower than the cost if planned maintenance cannot be delayed.

(b) The per period cost for the case in which planned maintenance can be delayed an unlimited number of times is less than or equal to the per period cost for the case in which planned maintenance can be delayed at most once, with the difference in cost equal to:

$$-\sum_{k=1}^{P_t^d} \mathbb{P}\left\{U_t = S_t + P_t + k\right\}(C^u - C^p)k - \sum_{k=1}^{\infty} \mathbb{P}\left\{U_t = S_t + P_t + P_t^d + k\right\}(C^u - C^p)P_t^d.$$

If $P_t^d > 0$, then the cost when planned maintenance can be delayed an unlimited number of times is thus strictly lower than the cost if planned maintenance can be delayed once.

(ii) At the optimal safety stock level:

(a) The optimal per period cost for the case in which planned maintenance can be delayed at most once is less than or equal to the optimal per period cost for the case in which

planned maintenance can never be delayed.

(b) *The optimal per period cost for the case in which planned maintenance can be delayed an unlimited number of times is less than or equal to the optimal per period cost for the case in which planned maintenance can be delayed once.*

(ii) *The optimal safety stock levels satisfy:*

(a) *The optimal safety stock level for the case in which planned maintenance can be delayed at most once is less than or equal to the optimal safety stock level for the case in which planned maintenance can never be delayed.*

(b) *The optimal safety stock level for the case in which planned maintenance can be delayed an unlimited number of times is less than or equal to the optimal safety stock level for the case in which planned maintenance can be delayed once.*

8 Positive supply lead time

We have thus far considered the case in which parts that are ordered arrive instantaneously. We now introduce a supply lead time $L \in \mathbb{N}_0$. The order of events in every period remains the same, so that previously ordered parts arrive immediately after ordering new parts. However, in period t, the parts that arrive were ordered in period $t-L$. In addition, to enable the analysis, we now assume an infinite horizon, i.e., $T=\infty$, and we assume that P_t is stationary, i.e., $P_t = P_{t'}$ for all periods t and t'. In practice, this is a natural assumption since companies generally try to schedule a stable work load in the maintenance shop. We will consider only the cases in which planned maintenance may not be delayed and may be delayed at most once. We will demonstrate that the results derived for the zero lead time setting also hold, with some minor adaptation, when $L > 0$.

We first introduce some additional definitions and notation. The inventory position after ordering in period t, Y_t, is now defined as all parts on hand and on order after ordering in period t, $Y_t = I_t + \sum_{r=t-L}^{t} O_r$. We then define $\mathbf{P}_{t+L}$ to be the lead time planned maintenance, i.e., all planned maintenance that arises in periods t through $t+L$, $\mathbf{P}_{t+L} = \sum_{r=t}^{t+L} P_r = (L+1)P_t$. Similarly, we define the lead time unplanned maintenance, $\mathbf{U}_{t+L} = \sum_{r=t}^{t+L} U_r$. We next define lead time known demands as all demands over the lead time that are known at the beginning of period t, i.e., $P_t^{\mathrm{d}} + U_t^{\mathrm{d}} + \mathbf{P}_{t+L}$. Finally, we define $\mathbf{S}_{t+L}$ to be the safety stock level in period $t+L$, which we define to be the inventory position after ordering in period t minus the lead time known demands in period t. Thus, we have

$$\mathbf{S}_{t+L} = I_t + \sum_{r=t-L}^{t} O_r - P_t^{\mathrm{d}} - U_t^{\mathrm{d}} - \mathbf{P}_{t+L} = Y_t - P_t^{\mathrm{d}} - U_t^{\mathrm{d}} - \mathbf{P}_{t+L}. \tag{6}$$

The safety stock level, $\mathbf{S}_{t+L}$, is our decision variable in period t. Alternatively and equivalently, we could have chosen O_t to be our decision variable.

We next provide some observations without formal proof. Observations 4.1 and 4.2, which consider the allocation policy, and Observation 4.5, which considers ordering costs, remain unchanged. However, below we provide modified versions of the other two observations in Section 4.

Observation 8.1. Under any optimal policy:

(i) There will be no unplanned maintenance actions that are deliberately delayed indefinitely.

(ii) The inventory position after ordering in period t is at least as large as the lead time known demands at the beginning of period t: $Y_t = I_t + \sum_{r=t-L}^{t} O_r \geq P_t^{\mathrm{d}} + U_t^{\mathrm{d}} + \mathbf{P}_{t+L}$.

Part (i) follows immediately from part (i) of Observation 4.3. The reasoning for part (ii) is similar to that for part (ii) of Observation 4.3. Given that all demands will be fulfilled at some point, it is best to fill them as soon as possible. From part (ii) of Observation 8.1 it follows that $\mathbf{S}_{t+L} \geq 0$.

Observation 8.2. There exists an optimal policy under which an unplanned maintenance action will be delayed at most $L+1$ times.

This observation is a generalization of Observation 4.4 and the reasoning for the result is similar. The inventory position after ordering in period t, Y_t, is larger than the lead time known demands at the beginning of period t, $P_t^{\mathrm{d}} + U_t^{\mathrm{d}} + \mathbf{P}_{t+L}$. An unplanned maintenance action that has been delayed in period $t-1$ is a known demand in period t and thus can be filled in period $t+L$, at the latest.

We now consider the case in which planned maintenance may not be delayed, i.e., stock is controlled separately for planned and unplanned maintenance. We have the following cost function for period $t+L$:

$$C(\mathbf{S}_{t+L} \mid P_{t+L}) = \sum_{k=0}^{\mathbf{S}_{t+L}} \mathbb{P}\left\{\mathbf{U}_{t+L} = k\right\} C^{\mathrm{h}} \times (\mathbf{S}_{t+L} - k) + \sum_{k=1}^{\infty} \mathbb{P}\left\{\mathbf{U}_{t+L} = \mathbf{S}_{t+L} + k\right\} C^{\mathrm{u}} \times k. \qquad (7)$$

We suppress P_t^{d} in the cost function since it is always zero. This cost function is analogous to that for the case with zero lead time. Notice that it is not important to know when the unplanned maintenance has occurred exactly; this result is analogous to the observation that only the aggregate demand over the lead time matters in an inventory model with backordering and a positive supply lead time, as shown by Scarf (1959). By noticing the similarity between Equations (7) and (1), it is easily seen that the results in Theorem 1 also hold when $L > 0$, with the unplanned maintenance being replaced by its positive lead time equivalent and the ordering decisions now based on the inventory position. We refer to the appendix for the details on the initial periods.

We next consider the case in which planned maintenance may be delayed at most once. We have the following cost function for period $t+L$:

$$C(\mathbf{S}_{t+L} \mid P_{t+L}) = \sum_{k=0}^{\mathbf{S}_{t+L}} \mathbb{P}\left\{\mathbf{U}_{t+L} = k\right\} C^{\mathrm{h}} \times (\mathbf{S}_{t+L} - k) + \sum_{k=1}^{P_{t+L}} \mathbb{P}\left\{\mathbf{U}_{t+L} = \mathbf{S}_{t+L} + k\right\} C^{\mathrm{p}} \times k$$
$$+ \sum_{k=1}^{\infty} \mathbb{P}\left\{\mathbf{U}_{t+L} = \mathbf{S}_{t+L} + P_{t+L} + k\right\} \left(C^{\mathrm{u}} \times k + C^{\mathrm{p}} \times P_{t+L}\right). \tag{8}$$

We suppress P_t^{d} in the cost function since it is always fulfilled in period t and does not influence the costs or the optimal decision. By noticing the similarity between Equations (8) and (2), it can be seen that the results in Lemma 3, part (i) and Theorem 4 also hold when $L > 0$, with the inventory position, known demands and unplanned maintenance replaced by their positive lead time equivalents, and the ordering decisions now based on the inventory position. We refer to the appendix for a formal proof of this result.

9 Numerical results and heuristic performance

In this section, we report the results of a numerical experiment in which we compare the optimal costs in three cases:

- Planned maintenance may not be delayed;
- Planned maintenance may be delayed at most once; and
- Planned maintenance may be delayed an unlimited number of times.

By comparing the system performance across these three cases, we gain insights into the value of allowing planned maintenance to be delayed.

In addition, we investigate the performance of four heuristic solution approaches for the case with unlimited delay, for which finding the optimal solution can be computationally challenging. The first two heuristics are myopic policies based on the bounds on the optimal safety stock level derived in in Corollary 9. In the *LB myopic heuristic* we always use the lower bound, while in the *UB myopic heuristic* we always use the upper bound. As discussed in Section 7, when the lower and upper bounds coincide, we have (myopically) found the optimal safety stock level for the case with unlimited delay. The third and fourth heuristics apply the safety stock levels derived for the no delay and one delay cases, respectively, even though an unlimited number of delays is allowed. In other words, under this approach, we do not allow for unlimited delays at a tactical level, when setting the safety stock level, but operationally, when we actually allocate the safety stock to maintenance actions, we do allow for unlimited delays. We refer to these two heuristics as the *no delay heuristic* and *one delay heuristic*, respectively.

		Settings
Supply lead time	L	0
Number of periods	T	5; 25; ∞
Planned maintenance per period	P_t	5;25
Unplanned maintenance rate	λ^{u}	1;5
Holding costs	C^{h}	1
Delay costs for planned maintenance	C^{p}	1;5
Delay costs for unplanned maintenance	C^{u}	10;50

Table 1: Settings of the parameters in the numerical experiment

In summary, we consider seven different solution approaches: the optimal solutions for each of the three different cases, plus the four heuristics for the unlimited delay case. When we compare the performance of two solution approaches, it will be clear which approach provides the lowest cost. We use I to denote the lowest cost approach and II to denote the other approach. As outlined in Section 9.1, we consider 48 different problem instances. We will use $C_I(e)$ and $C_{II}(e)$ to denote the costs resulting from approach I and II, respectively, applied to problem instance e. Finally, we will calculate and report the average and maximum percentage increase in costs for approach II, relative to apporach I, as $\frac{1}{48}\sum_{e=1}^{48}\left(\frac{C_{II}(e)-C_I(e)}{C_I(e)}\right)$ and $\max_e\left\{\frac{C_{II}(e)-C_I(e)}{C_I(e)}\right\}$, respectively.

9.1 Experimental design

We use a full factorial experiment. We assume that $L = 0$, since we have shown in Section 8 that the structure of the results (for the first two cases) will be the same if $L > 0$. We further assume, for simplicity and because in practice companies try to schedule a stable workload in the maintenance shop, that the amount of planned maintenance that is scheduled in each period, P_t, is stationary over time. We assume that the unplanned maintenance, U_t, has a truncated Poisson distribution with rate λ^{u}, truncated at $10\lambda^u$, so that only values with a very small probability of occurring are ignored. (We do this truncation so that the amount of delayed maintenance is capped, which helps to avoid computational issues.) Assuming a Poisson distribution is realistic in practice and common in the literature on spare parts inventory management, see, e.g., Sherbrooke (2004, p.21) or Basten and Van Houtum (2014, p.40).

Table 1 provides the exact parameter values that we consider. We perform $1\times3\times2\times2\times1\times2\times2 = 48$ different experiments. With $T = \infty$, we mean that we consider the infinite horizon setting, in which the optimal safety stock level is stationary in all three cases, including the case in which planned maintenance may be delayed an unlimited number of times. In the latter case, we stop the value iteration once we have achieved a precision of 0.0001, i.e., when $|V_2 - V_1|/V_1 < 0.0001$. This never requires T to become higher than 44.

	No delay	At most one delay
Average	136%	0.50%
Maximum	334%	4.92%

Table 2: % additional costs relative to the setting with an unlimited number of delays

	No delay	At most one delay
$T = 5$	136%	0.44%
$T = 10$	137%	0.54%
$T = \infty$	137%	0.53%
$P_t = 5$	130%	1.01%
$P_t = 25$	143%	0.00%
$\lambda^u = 1$	163%	0.02%
$\lambda^u = 5$	110%	0.99%
$C^{\mathrm{p}} = 1$	217%	0.96%
$C^{\mathrm{p}} = 5$	55.5%	0.05%
$C^{\mathrm{u}} = 10$	95.9%	0.38%
$C^{\mathrm{u}} = 50$	177%	0.63%

Table 3: Average % additional costs relative to the setting with an unlimited number of delays

9.2 Impact of delayed planned maintenance

We first quantify the value of allowing planned maintenance to be delayed. The results are presented in Tables 2 through 4. If we do not allow planned maintenance to be delayed, the increase in costs compared to the case with unlimited delay is 136% on average and 334% at maximum. If we allow planned maintenance to be delayed at most once, the increase in costs compared to the case with unlimited delay is 0.50% on average and 4.92% at maximum. These results indicate that most of the benefits obtained from allowing planned maintenance to be delayed are gained from allowing the maintenance to be delayed at most once. This holds because, in general, if we use the optimal safety stock levels, it is rare that planned maintenance is delayed more than one time. In other words, the probability that the unplanned demand is high enough over two periods to cause a planned demand to be delayed twice is very small.

Table 3 shows the average percentage increase in costs, relative to the case with unlimited delay, for the various parameter settings. We focus our discussion on the first column, which shows the percentage increase in costs for the case with no delay. From the table, we can see that λ^u, C^{p} and C^{u} have a significant impact on the value of allowing planned maintenance to be delayed. Unsurprisingly, allowing planned maintenance to be delayed adds the most value in low demand settings, when the cost of delaying planned maintenance is low, and when the cost of delaying

	No delay	At most one delay	Unlimited delay $P_t^{\text{d}} = 0$	Unlimited delay $P_t^{\text{d}} > 0$
Total	5.75	3.81	3.81	3.75
$T = 5$	5.75	3.81	3.81	3.75
$T = 10$	5.75	3.81	3.81	3.75
$T = \infty$	5.75	3.81	3.81	3.75
$P_t = 5$	5.75	3.88	3.88	3.75
$P_t = 25$	5.75	3.75	3.75	3.75
$\lambda^u = 1$	2.50	1.50	1.50	1.50
$\lambda^u = 5$	9.00	6.13	6.13	6.00
$C^{\text{p}} = 1$	5.75	3.13	3.13	3.00
$C^{\text{p}} = 5$	5.75	4.50	4.50	4.50
$C^{\text{u}} = 10$	5.00	3.75	3.75	3.75
$C^{\text{u}} = 50$	6.50	3.88	3.88	3.75

Table 4: Average safety stock levels

unplanned maintenance is high.

Table 4 shows the optimal safety stock levels for the three cases. For the case with unlimited delays, there are three problem instances in which the stock levels differ for different amounts of already delayed planned maintenance, P_t^{d}; we come back to these problem instances in Section 9.3. As stated in Lemma 10, the safety stock levels are lowest in the unlimited delay case, and highest in the no delay case. In addition, we see that the safety stock levels for the one delay case are almost identical to those for the unlimited delay case. In fact, the stock levels are identical when $P_t^{\text{d}} = 0$.

These results again suggest that allowing planned maintenance to be delayed at most once is sufficient to achieve very close to the optimal performance for the unlimited delay case. It is interesting to note that, in the one delay and unlimited delay cases, the safety stock levels are relatively insensitive to P_t. In fact, although we have shown that S_t^* is decreasing in P_t, the numerical results indicate that this effect is quite small. However, the optimal safety stock levels in the one delay case are dependent on the cost of delaying planned maintenance, C^{p}. As one would expect, when this cost increases, the optimal safety stock level increases.

9.3 Heuristic approaches for setting with unlimited delays

Recall that the optimal solution for the case in which planned maintenance can be delayed an unlimited number of times can be found by solving the MDP formulated in Section 7. Since solving the MDP may be computationally intensive, and requires mathematical sophistication on the part of the inventory manager, we next evaluate the performance of the four proposed heuristic

	LB myopic	UB myopic	No delay	One delay
Average	35%	0.01%	69%	0.09%
Maximum	185%	0.21%	187%	1.64%

Table 5: % optimality gaps of the four heuristics for the case with an unlimited number of delays

approaches for setting the safety stock level.

The key results can be found in Table 5. We first consider the LB myopic and UB myopic heuristics, which use the bounds provided in Corollary 9. We find that the UB myopic heuristic performs very well, with an optimality gap of only 0.01% on average and 0.21% at maximum, while the LB myopic heuristic gives optimality gaps of 35% on average and 185% at maximum. To understand these results, we compare Lemma 7, part (i), which gives the optimal solution for the unlimited delay case, with Corollary 9, which gives the bounds that we use in our heuristics. In doing so, we notice the following: The LB myopic heuristic assigns the difference in cost in period $t+1$ due to a change in the number of delayed planned maintenance actions, i.e., $\Delta_{P^{\mathrm{d}}_{t+1}} V_{t+1}(P_{t+1}, P^{\mathrm{d}}_{t+1})$, a value of $-C^{\mathrm{h}}$, while the UB myopic heuristic assigns it a value of 0, i.e., the UB myopic heuristic ignores this term. The numerical results indicate that the effect of $\Delta_{P^{\mathrm{d}}_{t+1}} V_{t+1}(P_{t+1}, P^{\mathrm{d}}_{t+1})$ is small and thus it is better to ignore it than to assign it too much value.

It is also interesting to note that in half of the experiments, i.e., all of the experiments with $C^{\mathrm{p}} = 5$, both heuristics give the same safety stock level, which implies that they provide the optimal safety stock level. This observation is particularly important given that $C^{\mathrm{p}} = 5$ is the more realistic setting that we consider. It ensures that $C^{\mathrm{p}} > C^{\mathrm{h}}$, i.e., the cost of delaying planned maintenance is greater than the cost of holding inventory, which does not hold when $C^{\mathrm{p}} = 1$.

We next consider the no delay and one delay heuristics, which apply the safety stock levels derived for the no delay and one delay cases, respectively, even though an unlimited number of delays is actually allowed. Applying the no delay heuristic leads to an optimality gap of 69% on average and 187% at maximum, while the one delay heuristic leads to an optimality gap of 0.09% on average and 1.64% at maximum. The results for the no delay heuristic further show that if delaying planned maintenance is allowable at an operational level, i.e., when allocating inventory to maintenance actions, then it is important to incorporate those potential delays, at least to some extent, when determining safety stock levels.

Table 5 indicates that the UB myopic heuristic outperforms the LB myopic heuristic and that the one delay heuristic outperforms the no delay heuristic. The natural next question is whether the UB myopic or the one delay heuristic is best overall. Both heuristics have the advantage that they incorporate the reality that planned maintenance may be delayed, while still ensuring that the base stock level is easy to determine, which makes it more credible and implementable for the inventory planner.

In 45 of the 48 problem instances, the UB myopic heuristic and the one delay heuristic both find the optimal solution. In this optimal solution, the optimal safety stock level does not vary with P_t^{d}, i.e., the optimal safety stock level does not depend on the amount of delayed planned maintenance that is still waiting to be filled. For the three remaining problem instances, the optimal solution is to set $S_t^* = 6$ if $P_t^{\mathrm{d}} = 0$ and $S_t^* = 5$ if $P_t^{\mathrm{d}} > 0$ (this explains the difference in the final two columns of Table 4). The UB myopic heuristic sets the safety stock level equal to 6 if $P_t^{\mathrm{d}} \in \{0, 1\}$ and 5 if $P_t^{\mathrm{d}} > 1$, while the one delay myopic heuristic sets the safety stock level to 6. By its design, the safety stock levels under the one delay myopic heuristic do not depend on P_t^{d}. Thus, the UB myopic heuristic performs better than the one delay heuristic, due to the fact that the UB myopic heuristic can take into account P_t^{d} when setting the safety stock level, while the one delay myopic heuristic cannot.

Overall, the above results demonstrate that a myopic heuristic policy can very well be used in practice when planned maintenance may be delayed an unlimited number of times. The results also indicate that the best myopic heuristic policy will consider as much information as possible. As a result, the UB myopic heuristic policy, which was developed based on our understanding of the true optimal policy, derived through an analysis of the MDP formulation for the unlimited delay case, outperforms the one delay heuristic.

10 Conclusions

The problem of managing an inventory system with multiple classes of demand is one that arises in many practical settings. This problem can be particularly challenging when the demand classes vary in the level of demand uncertainty, the urgency of fulfillment, and the costs associated with unfilled demands. Unfortunately, the impact of these factors on the optimal inventory control policy is highly dependent on the specific details of the practical setting. In this paper, we consider a setting motivated by the inventory management of spare parts used to support the maintenance of capital systems, such as trains or airplanes. In this setting, there are generally two classes of demand, representing parts needed for planned and unplanned maintenance. Planned maintenance is highly scheduled and hence creates a deterministic demand for spare parts. In contrast, unplanned maintenance, by definition, is unpredictable and thus creates a random demand for spare parts. In addition, while planned maintenance is generally preventive in nature, and thus can be temporarily delayed without significant costs, unplanned maintenance is corrective in nature, and thus requires an immediate response in order to prevent costly system downtime. Therefore, the demands associated with unplanned maintenance are of higher priority, incurring a higher delay cost, than those associated with planned maintenance.

For this problem setting, we have developed a periodic review inventory model, which we use to characterize the optimal inventory control policy, which includes both the inventory ordering decision and the inventory allocation decision (i.e., how to allocate available inventory to planned and unplanned demands). We analyze this model for three cases, characterized by the amount of delay that can be imposed on planned maintenance. In the first case, planned maintenance can never be delayed. In this case, because the known planned maintenance demands must always be immediately fulfilled, the problem is separable, i.e., it is optimal to keep separate pools of inventory for planned and unplanned maintenance demands. In addition, the problem of managing inventory to meet the unplanned demands reduces to a standard single product stochastic inventory model, for which a myopic policy is optimal (under certain conditions).

The second and third cases consider settings in which planned maintenance may be delayed, at most once in the second case, and indefinitely in the third case. In these settings, the problem is no longer separable, i.e., it is optimal to stock a single pool of inventory used to fill both planned and unplanned demands. For this problem, we define the decision variable in the inventory system to be the safety stock level, i.e., the optimal amount of inventory to hold in excess of the known demands (which include planned demands, as well as any planned or unplanned demands that have been delayed), and we find that the safety stock level in a given period is a function of the amount of planned maintenance in that period. Thus, the optimal policy explicitly captures the dependence between inventory decisions and maintenance decisions. In other words, our model enables inventory planners and maintenance managers to better understand the impact that decisions regarding planned maintenance activities have on inventory requirements, thus enabling a better joint management of these two critical business functions.

Our model also captures one of the key benefits of allowing planned maintenance activities to be delayed: risk pooling. Specifically, when planned maintenance can be delayed and the inventory used to satisfy planned and unplanned demands is jointly managed, stock that was to be allocated to a planned demand can be used to meet a higher priority unplanned demand, if needed. As a result, the optimal system costs and the optimal safety stock levels are decreasing in the amount of delay allowed for planned maintenance. In addition, the optimal safety stock level in a given period is decreasing in the amount of planned maintenance in that period, i.e., when more inventory is stocked to meet planned maintenance activities, fewer units can be stocked to meet any unplanned maintenance events.

We used numerical experiments to quantify the benefits associated with delaying planned maintenance. Specifically, through a comparison of the system performance in each of the three cases (no delay, one delay, and unlimited delay), we found that most of the benefits associated with allowing planned maintenance to be delayed are gained from allowing just a single delay. In other words, while allowing an unlimited delay for planned maintenance activities can further reduce inventory costs, the savings, relative to allowing just one delay, are minimal. This observation is important because, in many practical settings, maintenance managers are hesitant or unable to delay planned maintenance for extended periods of time, often due to firm-level policies or government regulations developed in the interest of safety (e.g., aircraft engines have very specific planned maintenance schedules that must be closely followed).

Finally, our analysis provides inventory planners with easy-to-compute inventory control policies that capture the dependence between maintenance activities and inventory requirements. Specifically, when planned demands can be delayed at most once, we find that an easy-to-compute myopic inventory control policy is optimal. However, when planned demands may be delayed indefinitely, finding the optimal safety stock levels is more challenging. Therefore, in addition to characterizing the structure of the optimal inventory policy, we developed several heuristic solution approaches, all of which reduce to solving a single product myopic stochastic inventory control problem. Using a series of numerical experiments, we evaluated the proposed heuristics and concluded that some of them can provide a system performance very close to the optimal performance. In particular, the UB myopic policy, which was developed through an analysis of the MDP formulation of the problem setting with unlimited delay, was found to perform very well relative to the optimal policy, finding the optimal solution in 94% of problem instances.

Acknowledgements

The research of the second author was supported by National Science Foundation grant CMMI-1300968.

References

Basten, R. J. I. and Van Houtum, G. J. (2014). System-oriented inventory models for spare parts. *Surveys in Operations Research and Management Science*, 19(1):34–55.

Dekker, R., Hill, R. M., Kleijn, M. J., and Teunter, R. (2002). On the $(s-1, s)$ lost sales inventory model with priority demand classes. *Naval Research Logistics*, 49(6):593–610.

Deshpande, V., Cohen, M. A., and Donohue, K. (2003). A threshold inventory rationing policy for service-differentiated demand classes. *Management Science*, 49(6):683–703.

Frank, K. C., Zhang, R. Q., and Duenyas, I. (2003). Optimal policies for inventory systems with priority demand classes. *Operations Research*, 51(6):993–1002.

Gallego, G. and Özer, Ö. (2001). Integrating replenishment decisions with advance demand information. *Management Science*, 47(10):1344–1360.

Hariharan, R. and Zipkin, P. H. (1995). Customer-order information, leadtimes, and inventories. *Management Science*, 41(10):1599–1607.

Heyman, D. P. and Sobel, M. J. (1984). *Stochastic Models in Operations Research: Stochastic Optimization*, volume 2. McGraw Hill, New York (NY).

Kranenburg, A. A. and Van Houtum, G. J. (2008). Service differentiation in spare parts inventory management. *Journal of the Operational Research Society*, 59:946–955.

Scarf, H. (1959). The optimality of (S,s) policies in the dynamic inventory problem. In Arrow, K. J., Karlin, S., and Suppes, P., editors, *Mathematical Methods in the Social Science*, chapter 13, pages 196–202. Stanford University Press, Palo Alto (CA).

Sherbrooke, C. C. (2004). *Optimal inventory modelling of systems. Multi-echelon techniques.* Kluwer, Dordrecht (The Netherlands), second edition.

Sobel, M. J. and Zhang, R. Q. (2001). Inventory policies for systems with stochastic and deterministic demand. *Operations Research*, 49(1):157–162.

Tan, T., Güllü, R., and Erkip, N. K. (2007). Modelling imperfect advance demand information and analysis of optimal inventory policies. *European Journal of Operational Research*, 177:897–923.

Topkis, D. M. (1968). Optimal ordering and rationing policies in a nonstationary dynamic inventory model with n demand classes. *Management Science*, 15(3):160–176.

Veinott, A. F. (1965). Optimal policy in a dynamic, single product, nonstationary inventory model with several demand classes. *Operations Research*, 13(5):761–778.

Appendix

Throughout the appendix, we use $\Delta_x f(x,y) = f(x+1,y) - f(x,y)$ and $\Delta_x\Delta_y f(x,y) = \Delta_y f(x+1,y) - \Delta_y f(x,y)$. Thus, $\Delta_x^2 f(x,y) = \Delta_x f(x+1,y) - \Delta_x f(x,y)$.

A Proofs for Section 5

In Section 5, planned maintenance may not be delayed. The cost function is defined in Equation (1), where we suppress P_t^{d} since it is always zero. In the proof of Theorem 1, we use the Lemma A.1, which is stated and proven below.

Lemma A.1. *For the cost function $C(S_t \mid P_t)$, as defined in Equation* (1) *it holds that $\Delta_{S_t} C(S_t \mid P_t) = C^h - (C^h + C^u)\sum_{k=S_t+1}^{\infty} \mathbb{P}\{U_t = k\}$.*

Proof.

$$\Delta_{S_t} C(S_t \mid P_t) = C(S_t+1 \mid P_t) - C(S_t \mid P_t) = C^{\mathrm{h}} - (C^{\mathrm{h}} + C^{\mathrm{u}}) \sum_{k=S_t+1}^{\infty} \mathbb{P}\{U_t = k\}.$$

□

Proof of Theorem 1. Using Lemma A.1, we can calculate $\Delta_{S_t}^2 C(S_t \mid P_t)$ as follows:

$$\Delta_{S_t}^2 C(S_t \mid P_t) = \Delta_{S_t} C(S_t+1 \mid P_t) - \Delta_{S_t} C(S_t \mid P_t) = (C^{\mathrm{h}} + C^{\mathrm{u}})\mathbb{P}\{U_t = S_t+1\}.$$

This shows that $\Delta_{S_t}^2 C(S_t \mid P_t) > 0$, which implies that $C(S_t \mid P_t)$ is convex in S_t. Again using Lemma A.1, we see that S_t^* is the smallest S_t for which $(C^{\mathrm{h}} + C^{\mathrm{u}})\sum_{k=S+1}^{\infty} \mathbb{P}\{U_t = k\} \leq C^{\mathrm{h}}$. □

B Proofs for Section 6

In Section 6, planned maintenance may be delayed at most once. The cost function is defined in Equation (2), where we suppress P_t^{d} since it does not influence the cost, i.e., since all delayed planned (and unplanned) maintenance is fulfilled in period t, no holding costs or delay costs are incurred. In some of the proofs, we use the results in Lemma B.1, which is stated and proven below.

Lemma B.1. *The cost function $C(S_t \mid P_t)$, as defined in Equation* (2) *has the following properties:*

(i) $\Delta_{S_t} C(S_t \mid P_t) = C^h - (C^h + C^p)\sum_{k=S_t+1}^{\infty} \mathbb{P}\{U_t = k\} - (C^u - C^p)\sum_{k=S_t+P_t+1}^{\infty} \mathbb{P}\{U_t = k\}$.

(ii) $\Delta_{P_t} C(S_t \mid P_t) = -(C^u - C^p)\sum_{k=S_t+P_t+1}^{\infty} \mathbb{P}\{U_t = k\}$.

Proof. (i) To derive $\Delta_{S_t} C(S_t \mid P_t)$, we first rewrite $C(S_t \mid P_t)$ as follows:

$$\begin{aligned} C(S_t \mid P_t) = &\sum_{k=0}^{S_t} \mathbb{P}\{U_t = k\} C^{\mathrm{h}}(S_t - k) + \sum_{k=1}^{\infty} \mathbb{P}\{U_t = S_t + k\} C^{\mathrm{p}} k \\ &+ \sum_{k=1}^{\infty} \mathbb{P}\{U_t = S_t + P_t + k\} (C^{\mathrm{u}} - C^{\mathrm{p}}) k \end{aligned}$$

Then:

$$\begin{aligned} \Delta_{S_t} C(S_t \mid P_t) &= C(S_t + 1 \mid P_t) - C(S_t \mid P_t) \\ &= C^{\mathrm{h}} \sum_{k=0}^{S_t} \mathbb{P}\{U_t = k\} - C^{\mathrm{p}} \sum_{k=S_t+1}^{\infty} \mathbb{P}\{U_t = k\} - (C^{\mathrm{u}} - C^{\mathrm{p}}) \sum_{k=S_t+P_t+1}^{\infty} \mathbb{P}\{U_t = k\} \\ &= C^{\mathrm{h}} - (C^{\mathrm{h}} + C^{\mathrm{p}}) \sum_{k=S_t+1}^{\infty} \mathbb{P}\{U_t = k\} - (C^{\mathrm{u}} - C^{\mathrm{p}}) \sum_{k=S_t+P_t+1}^{\infty} \mathbb{P}\{U_t = k\}. \end{aligned}$$

(ii) Using the rewritten version of $C(S_t \mid P_t)$ in part (i), we see that:

$$\begin{aligned} \Delta_{P_t} C(S_t \mid P_t) &= C(S_t \mid P_t + 1) - C(S_t \mid P_t) \\ &= -(C^{\mathrm{u}} - C^{\mathrm{p}}) \sum_{k=S_t+P_t+1}^{\infty} \mathbb{P}\{U_t = k\}. \end{aligned}$$

□

Proof of Lemma 2. (i) Using Lemma B.1, part (i):

$$\begin{aligned} \Delta^2_{S_t} C(S_t \mid P_t) &= \Delta_{S_t} C(S_t + 1 \mid P_t) - \Delta_{S_t} C(S_t \mid P_t) \\ &= (C^{\mathrm{h}} + C^{\mathrm{p}}) \mathbb{P}\{U_t = S_t + 1\} + (C^{\mathrm{u}} - C^{\mathrm{p}}) \mathbb{P}\{U_t = S_t + P_t + 1\}. \end{aligned}$$

Since $C^{\mathrm{u}} > C^{\mathrm{p}}$, $\Delta^2_{S_t} C(S_t \mid P_t) > 0$ and $C(S_t \mid P_t)$ is strictly convex in S_t.

(ii) The fact that $C(S_t \mid P_t)$ is strictly decreasing follows from Lemma B.1, part (ii), by noticing that $C^{\mathrm{u}} > C^{\mathrm{p}}$. Then, using Lemma B.1, part (ii), we have:

$$\begin{aligned} \Delta^2_{P_t} C(S_t \mid P_t) &= \Delta_{P_t} C(S_t \mid P_t + 1) - \Delta_{P_t} C(S_t \mid P_t) \\ &= (C^{\mathrm{u}} - C^{\mathrm{p}}) \mathbb{P}\{U_t = S_t + P_t + 1\}. \end{aligned}$$

Since $C^{\mathrm{u}} > C^{\mathrm{p}}$, $\Delta^2_{P_t} C(S_t \mid P_t) > 0$ and thus $C(S_t \mid P_t)$ is strictly convex in P_t.

(iii) Using Lemma B.1, parts (i) and (ii), we have:

$$\begin{aligned}\Delta_{S_t}\Delta_{P_t}C(S_t \mid P_t) &= \Delta_{P_t}C(S_t+1 \mid P_t) - \Delta_{P_t}C(S_t \mid P_t)\\ &= (C^{\mathrm{u}} - C^{\mathrm{p}})\,\mathbb{P}\{U_t = S_t + P_t + 1\}.\end{aligned}$$

Since $C^{\mathrm{u}} > C^{\mathrm{p}}$, we have that $C(S_t \mid P_t)$ is supermodular in P_t and S_t.

□

Proof of Lemma 3. (i) This result follows immediately from the fact that $C(S_t \mid P_t)$ is convex in S_t (Lemma 2, part (i)) and the derivation of $\Delta_{S_t}C(S_t \mid P_t)$ (Lemma B.1, part (i)).

(ii) The second inequality holds since $C(S_t^*(P_t) \mid P_t)$ is supermodular in P_t and S_t (Lemma 2, part (iii)). The first inequality holds by definition if $S_t^*(P_t) \leq 1$. Otherwise, part (i) of the current lemma implies that $\Delta_{S_t}C(S_t^*(P_t)-1 \mid P_t) < 0$ and that if $\Delta_{S_t}C(S_t^*(P_t)-2 \mid P_t+1) < 0$, then $S_t^*(P_t)-2 < S_t^*(P_t+1)$. Therefore, showing that $\Delta_{S_t}C(S_t^*(P_t)-2 \mid P_t+1) < \Delta_{S_t}C(S_t^*(P_t)-1 \mid P_t)$ completes the proof:

$$\begin{aligned}&\Delta_{S_t}C(S_t^*(P_t)-2 \mid P_t+1) - \Delta_{S_t}C(S_t^*(P_t)-1 \mid P_t)\\ &= \Delta_{S_t}C(S_t^*(P_t)-2 \mid P_t+1) - \Delta_{S_t}C(S_t^*(P_t)-2 \mid P_t)\\ &\qquad + \Delta_{S_t}C(S_t^*(P_t)-2 \mid P_t) - \Delta_{S_t}C(S_t^*(P_t)-1 \mid P_t)\\ &= \Delta_{P_t}\Delta_{S_t}C(S_t^*(P_t)-2 \mid P_t) - \Delta_{S_t}^2C(S_t^*(P_t)-2 \mid P_t)\\ &= (C^{\mathrm{u}} - C^{\mathrm{p}})\mathbb{P}\{U_t = S_t^*(P_t) + P_t - 1\}\\ &\qquad - (C^{\mathrm{h}} + C^{\mathrm{p}})\mathbb{P}\{U_t = S_t^*(P_t) - 1\} - (C^{\mathrm{u}} - C^{\mathrm{p}})\mathbb{P}\{U_t = S_t^*(P_t) + P_t - 1\}\\ &= -(C^{\mathrm{h}} + C^{\mathrm{p}})\mathbb{P}\{U_t = S_t^*(P_t) - 1\}\\ &< 0.\end{aligned}$$

□

Proof of Theorem 4. Lemma 3 shows that it is optimal to order $P_{t-1} + S_{t-1}^*(P_{t-1})$ in period $t-1$ if (1) there are no spare parts left on hand at the beginning of the period, i.e., $I_{t-1} = 0$, (2) no delayed maintenance actions remain at the beginning of the period, i.e., $P_{t-1}^{\mathrm{d}} + U_{t-1}^{\mathrm{d}} = 0$, and (3) we ignore information on future periods. Our goal in this proof is to show that the policy specified in Lemma 3 is optimal in general for the case in which planned maintenance may be delayed at most once. To do so, we assume that the policy specified in the lemma is applied in period $t-1$ and we demonstrate that the policy specified in the lemma is feasible for period t, i.e., that the specified safety stock level is achievable in period t.

If we order $P_{t-1} + S_{t-1}^*(P_{t-1})$ in period $t-1$, i.e., if we use the policy specified in Lemma 3 in period $t-1$, then one of three cases occurs:

1. $U_{t-1} = S^*_{t-1}(P_{t-1})$. At the beginning of period t, no parts are left on stock, $I_t = 0$, and no unfulfilled maintenance actions remain, $P^{\mathrm{d}}_t + U^{\mathrm{d}}_t = 0$. We order $O_t = P_t + S^*_t(P_t) \geq 0$, so that $Y_t = P^{\mathrm{d}}_t + U^{\mathrm{d}}_t + P_t + S^*_t(P_t)$.

2. $U_{t-1} > S^*_{t-1}(P_{t-1})$. At the beginning of period t, no parts are left on stock, $I_t = 0$, but unfulfilled maintenance actions remain, $P^{\mathrm{d}}_t + U^{\mathrm{d}}_t = U_{t-1} - S^*_{t-1}(P_{t-1})$. We order $O_t = P_t + S^*_t(P_t) + U_{t-1} - S^*_{t-1}(P_{t-1}) > 0$ so that $Y_t = P^{\mathrm{d}}_t + U^{\mathrm{d}}_t + P_t + S^*_t(P_t)$.

3. $U_{t-1} < S^*_{t-1}(P_{t-1})$. At the beginning of period t, there are parts left on stock, $I_t = S^*_{t-1}(P_{t-1}) - U_{t-1}$, but no unfulfilled maintenance actions remain, $P^{\mathrm{d}}_t + U^{\mathrm{d}}_t = 0$. We order $O_t = P_t + S^*_t(P_t) - S^*_{t-1}(P_{t-1}) + U_{t-1}$ so that $Y_t = P^{\mathrm{d}}_t + U^{\mathrm{d}}_t + P_t + S^*_t(P_t)$. Ordering this amount would not be possible (i.e., we would have $O_t < 0$) if $S^*_{t-1}(P_{t-1}) - U_{t-1} > P_t + S^*_t(P_t)$. However, by considering three cases, we show that $O_t < 0$ is not possible, i.e., that achieving $S^*_t(P_t)$ is feasible:
   - $P_t = P_{t-1}$: $S^*_t(P_t) = S^*_{t-1}(P_{t-1})$ so that $S^*_{t-1}(P_{t-1}) - U_{t-1} \leq P_t + S^*_t(P_t)$.
   - $P_t < P_{t-1}$: Lemma 3, part (ii), shows that $S^*_{t-1}(P_{t-1}) \leq S^*_t(P_t)$, so that $S^*_{t-1}(P_{t-1}) - U_{t-1} \leq P_t + S^*_t(P_t)$.
   - $P_t > P_{t-1}$: Applying Lemma 3, part (ii), $(P_t - P_{t-1})$ times shows that $S^*_{t-1}(P_{t-1}) - P_t + P_{t-1} \leq S^*_t(P_t)$, so that $S^*_{t-1}(P_{t-1}) - U_{t-1} \leq P_t + S^*_t(P_t)$.

Thus, by induction, this demonstrates that as long as we followed the myopically optimal policy (specified in Lemma 3) in all previous periods t', i.e., for all $1 \leq t' < t$, then we can achieve the myopically optimal safety stock level in period t. In other words, as long as we have ordered myopically in periods $1, \ldots, t-1$, the myopically optimal safety stock level in period t is feasible. Thus, the myopic solution presented in Lemma 3 is optimal for the case in which planned maintenance may be delayed at most once. □

Proof of Lemma 5. The fact that $C(S^*_t(P_t) \mid P_t)$ is strictly decreasing follows from Lemma 2, part (ii). To show convexity, we first bound $\Delta^*_{P_t} C(S^*_t(P_t) \mid P_t) = C(S^*_t(P_t+1) \mid P_t+1) - C(S^*_t(P_t) \mid P_t)$ (notice the difference with $\Delta_{P_t} C(S^*_t(P_t) \mid P_t) = C(S^*_t(P_t) \mid P_t+1) - C(S^*_t(P_t) \mid P_t))$. We know from Lemma 3, part (ii), that we need to consider two cases only, i.e., if P_t is increased by one, then S^*_t either stays the same or decreases by one.

(1) If $S_t^*(P_t+1) = S_t^*(P_t)$, then:

$$\begin{aligned} C(S_t^*(P_t) \mid P_t+1) - C(S_t^*(P_t) \mid P_t) &= \Delta_{P_t} C(S_t^*(P_t) \mid P_t) \\ &= -(C^{\mathrm{u}} - C^{\mathrm{p}}) \sum_{k=S_t^*(P_t)+P_t+1}^{\infty} \mathbb{P}\left\{U_t = k\right\} \\ &\geq -C^{\mathrm{h}} + (C^{\mathrm{p}} + C^{\mathrm{h}}) \sum_{k=S_t^*(P_t)+1}^{\infty} \mathbb{P}\left\{U_t = k\right\}, \end{aligned}$$

where the second equation follows from Lemma B.1, part (ii), and the inequality follows from Lemma 3, part (i).

(2) If $S_t^*(P_t+1) = S_t^*(P_t) - 1$, then using Lemma B.1, parts (i) and (ii):

$$\begin{aligned} & C(S_t^*(P_t) - 1 \mid P_t+1) - C(S_t^*(P_t) \mid P_t) \\ &= C(S_t^*(P_t) - 1 \mid P_t+1) - C(S_t^*(P_t) \mid P_t+1) + C(S_t^*(P_t) \mid P_t+1) - C(S_t^*(P_t) \mid P_t) \\ &= -\Delta_{S_t} C(S_t^*(P_t) - 1 \mid P_t+1) + \Delta_{P_t} C(S_t^*(P_t) \mid P_t) \\ &= -C^{\mathrm{h}} + (C^{\mathrm{p}} + C^{\mathrm{h}}) \sum_{k=S_t^*(P_t)}^{\infty} \mathbb{P}\left\{U_t = k\right\} + (C^{\mathrm{u}} - C^{\mathrm{p}}) \sum_{k=S_t^*(P_t)+P_t+1}^{\infty} \mathbb{P}\left\{U_t = k\right\} \\ &\qquad - (C^{\mathrm{u}} - C^{\mathrm{p}}) \sum_{k=S_t^*(P_t)+P_t+1}^{\infty} \mathbb{P}\left\{U_t = k\right\} \\ &= -C^{\mathrm{h}} + (C^{\mathrm{p}} + C^{\mathrm{h}}) \sum_{k=S_t^*(P_t)}^{\infty} \mathbb{P}\left\{U_t = k\right\}. \end{aligned}$$

Combining both cases shows that:

$$-C^{\mathrm{h}} + \sum_{k=S_t^*(P_t)+1}^{\infty} \mathbb{P}\left\{U_t = k\right\} (C^{\mathrm{p}} + C^{\mathrm{h}}) \leq \Delta_{P_t}^* C(S_t^*(P_t) \mid P_t) \leq -C^{\mathrm{h}} + (C^{\mathrm{p}} + C^{\mathrm{h}}) \sum_{k=S_t^*(P_t)}^{\infty} \mathbb{P}\left\{U_t = k\right\}.$$

To show convexity of $C(S_t^*(P_t) \mid P_t)$ in P_t, we need to compare the change in the optimal costs when P_t is increased by one with when it is again increased by one. We know from Lemma 3, part (ii), that we need to consider four cases. if P_t is increased by one, then S_t^* stays the same or decreases by one. This holds also if P_t is again increased by one.

1. $S_t^*(P_t+2) = S_t^*(P_t+1) = S_t^*(P_t)$: Strict convexity follows from Lemma 2, part (ii).
2. $S_t^*(P_t+2) = S_t^*(P_t+1) - 1 = S_t^*(P_t) - 1$: From case (2) above and the fact that $S_t^*(P_t+1) =$

$S_t^*(P_t)$, we see that:

$$C(S_t^*(P_t+1)-1 \mid P_t+2) - C(S_t^*(P_t+1) \mid P_t+1) = -C^{\mathrm{h}} + (C^{\mathrm{p}}+C^{\mathrm{h}}) \sum_{k=S_t^*(P_t)}^{\infty} \mathbb{P}\left\{U_t = k\right\}.$$

This is equal to the upper bound on $\Delta^*_{P_t} C(S_t^*(P_t) \mid P_t)$, which we have found above. This means that $\Delta^*_{P_t} C(S_t^*(P_t) \mid P_t)$ is non-decreasing in P_t, i.e., $C(S_t^*(P_t) \mid P_t)$ is convex in P_t.

3. $S_t^*(P_t+2) = S_t^*(P_t+1) = S_t^*(P_t)-1$: From case (2) above, we see that:

$$C(S_t^*(P_t)-1 \mid P_t+1) - C(S_t^*(P_t) \mid P_t) = -C^{\mathrm{h}} + (C^{\mathrm{p}}+C^{\mathrm{h}}) \sum_{k=S_t^*(P_t)}^{\infty} \mathbb{P}\left\{U_t = k\right\}.$$

   This is equal to the lower bound on $\Delta^*_{P_t} C(S_t^*(P_t+1) \mid P_t+1)$, that we have found above, considering that $S_t^*(P_t+1) = S_t^*(P_t)-1$. This means that $\Delta^*_{P_t} C(S_t^*(P_t) \mid P_t)$ is non-decreasing in P_t.

4. $S_t^*(P_t+2) = S_t^*(P_t+1)-1 = S_t^*(P_t)-2$: From case (2) above, we see that:

$$C(S_t^*(P_t)-1 \mid P_t+1) - C(S_t^*(P_t) \mid P_t) = -C^{\mathrm{h}} + (C^{\mathrm{p}}+C^{\mathrm{h}}) \sum_{k=S_t^*(P_t)}^{\infty} \mathbb{P}\left\{U_t = k\right\}.$$

   Next, from case (2) above and the fact that $S_t^*(P_t+1) = S_t^*(P_t)-1$, we see that:

$$C(S_t^*(P_t+1)-1 \mid P_t+2) - C(S_t^*(P_t+1) \mid P_t+1) = -C^{\mathrm{h}} + (C^{\mathrm{p}}+C^{\mathrm{h}}) \sum_{k=S_t^*(P_t)-1}^{\infty} \mathbb{P}\left\{U_t = k\right\}.$$

   This means that $\Delta^*_{P_t} C(S_t^*(P_t) \mid P_t)$ is strictly increasing in P_t.

□

C Proofs for Section 7

In Section 7, planned maintenance may be delayed an unconstrained number of times. The cost function is defined in Equation (3). All results below are shown for any period t. For period T, the results simplify by noticing that $V_{T+1}(S_{T+1} \mid P_{T+1}, P^{\mathrm{d}}_{T+1}) \equiv 0$. Many of the results for period T also follow immediately from the analogous results in Section 6. The proofs for periods $t < T$ are by induction and we often use results on period $t+1$ which are proven later. The key results that we use are:

- $V_{t+1}(P_{t+1}, P^{\mathrm{d}}_{t+1})$ is convex in P^{d}_{t+1}. This follows from Lemma 8, part (i).

- $\Delta_{P^{\mathrm{d}}_{t+1}} V_{t+1}(P_{t+1}, P^{\mathrm{d}}_{t+1}) < 0$, i.e., $V_{t+1}(P_{t+1}, P^{\mathrm{d}}_{t+1})$ is strictly decreasing in P^{d}_{t+1}. This follows from Lemma 8, part (i).
- $\Delta_{P^{\mathrm{d}}_{t+1}} V_{t+1}(P_{t+1}, P^{\mathrm{d}}_{t+1}) > -C^{\mathrm{h}}$. This follows from Lemma 8, part (ii).

Finally, in some of the proofs we use the results in Lemma C.1, which is stated and proven below.

Lemma C.1. *$G_t(S_t \mid P_t, P_t^d)$ has the following properties:*

(i)

$$\Delta_{S_t} G_t(S_t \mid P_t, P_t^d) = C^h - (C^h + C^p) \sum_{k=S_t+1}^{\infty} \mathbb{P}\{U_t = k\} - (C^u - C^p) \sum_{k=S_t+P_t+P_t^d+1}^{\infty} \mathbb{P}\{U_t = k\}$$
$$- \sum_{k=1}^{P_t+P_t^d} \mathbb{P}\{U_t = S_t + k\} \left[\Delta_{P^d_{t+1}} V_{t+1}(P_{t+1}, k-1)\right].$$

(ii)

$$\Delta_{P_t} G_t(S_t \mid P_t, P_t^d) = \sum_{k=S_t+P_t+P_t^d+1}^{\infty} \mathbb{P}\{U_t = k\} \left[-(C^u - C^p) + \Delta_{P^d_{t+1}} V_{t+1}(P_{t+1}, P_t + P_t^d)\right].$$

The same holds for $\Delta_{P_t^d} G_t(S_t \mid P_t, P_t^d)$, i.e., $\Delta_{P_t^d} G_t(S_t \mid P_t, P_t^d) = \Delta_{P_t} G_t(S_t \mid P_t, P_t^d)$.

Proof. (i)

$$
\begin{aligned}
&\Delta_{S_t} G_t(S_t \mid P_t, P_t^{\mathrm{d}}) = G_t(S_t + 1 \mid P_t, P_t^{\mathrm{d}}) - G_t(S_t \mid P_t, P_t^{\mathrm{d}}) \\
&= C^{\mathrm{h}} - (C^{\mathrm{h}} + C^{\mathrm{p}}) \sum_{k=S_t+1}^{\infty} \mathbb{P}\{U_t = k\} - (C^{\mathrm{u}} - C^{\mathrm{p}}) \sum_{k=S_t+P_t+P_t^{\mathrm{d}}+1}^{\infty} \mathbb{P}\{U_t = k\} \\
&\quad + \sum_{k=0}^{S_t} \mathbb{P}\{U_t = k\} \left[V_{t+1}(P_{t+1}, 0) - V_{t+1}(P_{t+1}, 0)\right] \\
&\quad + \mathbb{P}\{U_t = S_t + 1\} V_{t+1}(P_{t+1}, 0) \\
&\quad + \sum_{k=2}^{P_t+P_t^{\mathrm{d}}} \mathbb{P}\{U_t = S_t + k\} \left[V_{t+1}(P_{t+1}, k-1) - V_{t+1}(P_{t+1}, k)\right] \\
&\quad + \mathbb{P}\left\{U_t = S_t + P_t + P_t^{\mathrm{d}} + 1\right\} V_{t+1}(P_{t+1}, P_t + P_t^{\mathrm{d}}) \\
&\quad - \mathbb{P}\{U_t = S_t + 1\} V_{t+1}(P_{t+1}, 1) \\
&\quad + \sum_{k=2}^{\infty} \mathbb{P}\left\{U_t = S_t + P_t + P_t^{\mathrm{d}} + k\right\} \left[V_{t+1}(P_{t+1}, P_t + P_t^{\mathrm{d}}) - V_{t+1}(P_{t+1}, P_t + P_t^{\mathrm{d}})\right] \\
&\quad - \mathbb{P}\left\{U_t = S_t + P_t + P_t^{\mathrm{d}} + 1\right\} V_{t+1}(P_{t+1}, P_t + P_t^{\mathrm{d}}) \\
&= C^{\mathrm{h}} - (C^{\mathrm{h}} + C^{\mathrm{p}}) \sum_{k=S_t+1}^{\infty} \mathbb{P}\{U_t = k\} - (C^{\mathrm{u}} - C^{\mathrm{p}}) \sum_{k=S_t+P_t+P_t^{\mathrm{d}}+1}^{\infty} \mathbb{P}\{U_t = k\} \\
&\quad - \sum_{k=1}^{P_t+P_t^{\mathrm{d}}} \mathbb{P}\{U_t = S_t + k\} \left[\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, k-1)\right],
\end{aligned}
$$

where the cost terms for the current period, after the first equation, are analogous to those in the proof of Lemma B.1, part (i), and the second equality follows by noticing that some terms cancel each other out.

(ii)

$$\Delta_{P_t} G_t(S_t \mid P_t, P_t^{\mathrm{d}}) = G_t(S_t \mid P_t + 1, P_t^{\mathrm{d}}) - G_t(S_t \mid P_t, P_t^{\mathrm{d}})$$
$$= -(C^{\mathrm{u}} - C^{\mathrm{p}}) \sum_{k=S_t+P_t+P_t^{\mathrm{d}}+1}^{\infty} \mathbb{P}\{U_t = k\}$$
$$+ \mathbb{P}\left\{U_t = S_t + P_t + P_t^{\mathrm{d}} + 1\right\} V_{t+1}(P_{t+1}, P_t + P_t^{\mathrm{d}} + 1)$$
$$- \mathbb{P}\left\{U_t = S_t + P_t + P_t^{\mathrm{d}} + 1\right\} V_{t+1}(P_{t+1}, P_t + P_t^{\mathrm{d}})$$
$$+ \sum_{k=2}^{\infty} \mathbb{P}\left\{U_t = S_t + P_t + P_t^{\mathrm{d}} + k\right\} \left[V_{t+1}(P_{t+1}, P_t + P_t^{\mathrm{d}} + 1) - V_{t+1}(P_{t+1}, P_t + P_t^{\mathrm{d}})\right]$$
$$= \sum_{k=S_t+P_t+P_t^{\mathrm{d}}+1}^{\infty} \mathbb{P}\{U_t = k\} \left[-(C^{\mathrm{u}} - C^{\mathrm{p}}) + \Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, P_t + P_t^{\mathrm{d}})\right],$$

where the cost terms for the current period, after the first equation, are analogous to those in the proof of Lemma B.1, part (ii).

□

Proof of Lemma 6. (i)

$$\Delta_{S_t}^2 G_t(S_t \mid P_t, P_t^{\mathrm{d}}) = \Delta_{S_t} G_t(S_t + 1 \mid P_t, P_t^{\mathrm{d}}) - \Delta_{S_t} G_t(S_t \mid P_t, P_t^{\mathrm{d}})$$
$$= (C^{\mathrm{h}} + C^{\mathrm{p}})\mathbb{P}\{U_t = S_t + 1\} + (C^{\mathrm{u}} - C^{\mathrm{p}})\mathbb{P}\left\{U_t = S_t + P_t + P_t^{\mathrm{d}} + 1\right\}$$
$$- \sum_{k=2}^{P_t+P_t^{\mathrm{d}}} \mathbb{P}\{U_t = S_t + k\} \left[\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, k-2) - \Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, k-1)\right]$$
$$- \mathbb{P}\left\{U_t = S + P_t + P_t^{\mathrm{d}} + 1\right\} \left[\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, P_t + P_t^{\mathrm{d}} - 1)\right]$$
$$+ \mathbb{P}\{U_t = S + 1\} \left[\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, 0)\right]$$

Since $\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, 0) > -C^{\mathrm{h}}$ (Lemma 8, part (ii)), the summation of the first plus last term (the two terms with $\mathbb{P}\{U_t = S_t + 1\}$) is positive. The second term is positive since $C^{\mathrm{u}} > C^{\mathrm{p}}$. The third term (with the summation) is positive due to the convexity of $V_{t+1}(P_{t+1}, P_{t+1}^{\mathrm{d}})$ in P_{t+1}^{d} (Lemma 8, part (i)). The fourth term is positive since $V_{t+1}(P_{t+1}, P_{t+1}^{\mathrm{d}})$ is decreasing in P_{t+1}^{d} (Lemma 8, part (i)). Therefore $\Delta_{S_t}^2 G_t(S_t \mid P_t, P_t^{\mathrm{d}}) > 0$, so that $G_t(S_t \mid P_t, P_t^{\mathrm{d}})$ is strictly convex in S_t.

(ii) The fact that $G_t(S_t \mid P_t, P_t^{\mathrm{d}})$ is strictly decreasing in P_t follows from Lemma C.1, part (ii), combined with the fact that $C^{\mathrm{u}} > C^{\mathrm{p}}$ and that $V_{t+1}(P_{t+1}, P_{t+1}^{\mathrm{d}})$ is decreasing in P_{t+1}^{d} (Lemma 8, part (i)). Then, to show strict convexity:

$$\begin{aligned}
&\Delta^2_{P_t} G_t(S_t \mid P_t, P^{\mathrm{d}}_t) \\
&= \Delta_{P_t} G_t(S_t \mid P_t + 1, P^{\mathrm{d}}_t) - \Delta_{P_t} G_t(S_t \mid P_t, P^{\mathrm{d}}_t) \\
&= \mathbb{P}\left\{U_t = S_t + P_t + P^{\mathrm{d}}_t + 1\right\}\left[(C^{\mathrm{u}} - C^{\mathrm{p}}) - \Delta_{P^{\mathrm{d}}_{t+1}} V_{t+1}(P_{t+1}, P_t + P^{\mathrm{d}}_t)\right] \\
&\quad + \sum_{k=2}^{\infty} \mathbb{P}\left\{U_t = S_t + P_t + P^{\mathrm{d}}_t + k\right\}\left[\Delta_{P^{\mathrm{d}}_{t+1}} V_{t+1}(P_{t+1}, P_t + P^{\mathrm{d}}_t + 1) - \Delta_{P^{\mathrm{d}}_{t+1}} V_{t+1}(P_{t+1}, P_t + P^{\mathrm{d}}_t)\right].
\end{aligned}$$

The first term is positive since $C^{\mathrm{u}} > C^{\mathrm{p}}$ and $V_{t+1}(P_{t+1}, P^{\mathrm{d}}_{t+1})$ is decreasing in P^{d}_{t+1} (Lemma 8, part (i)), and the second term is positive due to the convexity of $V_{t+1}(P_{t+1}, P^{\mathrm{d}}_{t+1})$ in P^{d}_{t+1} (Lemma 8, part (i)). Therefore $\Delta^2_{P_t} G_t(S_t \mid P_t, P^{\mathrm{d}}_t) > 0$, so that $G_t(S_t \mid P_t, P^{\mathrm{d}}_t)$ is strictly convex in P_t. Since $\Delta_{P^{\mathrm{d}}_t} G_t(S_t \mid P_t, P^{\mathrm{d}}_t) = \Delta_{P_t} G_t(S_t \mid P_t, P^{\mathrm{d}}_t)$, $G_t(S_t \mid P_t, P^{\mathrm{d}}_t)$ is also strictly decreasing and strictly convex in P^{d}_t.

(iii)

$$\begin{aligned}
&\Delta_{P_t}\Delta_{S_t} G_t(S_t \mid P_t, P^{\mathrm{d}}_t) \\
&= \Delta_{S_t} G_t(S_t \mid P_t + 1, P^{\mathrm{d}}_t) - \Delta_{S_t} G_t(S_t \mid P_t, P^{\mathrm{d}}_t) \\
&= \mathbb{P}\left\{U_t = S_t + P_t + P^{\mathrm{d}}_t + 1\right\}\left[(C^{\mathrm{u}} - C^{\mathrm{p}}) - \mathbb{E}\left\{\Delta_{P^{\mathrm{d}}_{t+1}} V_{t+1}(P_{t+1}, P_t + P^{\mathrm{d}}_t)\right\}\right].
\end{aligned}$$

This is positive since $C^{\mathrm{u}} > C^{\mathrm{p}}$ and $V_{t+1}(P_{t+1}, P^{\mathrm{d}}_{t+1})$ is decreasing in P^{d}_{t+1} (Lemma 8, part (i)). Since the expression is positive, $G_t(S_t \mid P_t, P^{\mathrm{d}}_t)$ is supermodular in P_t. Since $\Delta_{P^{\mathrm{d}}_t}\Delta_{S_t} G_t(S_t \mid P_t, P^{\mathrm{d}}_t) = \Delta_{P_t}\Delta_{S_t} G_t(S_t \mid P_t, P^{\mathrm{d}}_t)$ and $\Delta_{P^{\mathrm{d}}_t}\Delta_{P_t} G_t(S_t \mid P_t, P^{\mathrm{d}}_t) = \Delta^2_{P_t} G_t(S_t \mid P_t, P^{\mathrm{d}}_t)$, the other supermodularity properties are proven.

□

Proof of Lemma 7. (i) This result follows immediately from the fact that $G_t(S_t \mid P_t, P^{\mathrm{d}}_t)$ is convex in S_t (Lemma 6, part (i)) and the derivation of $\Delta_{S_t} G_t(S_t \mid P_t, P^{\mathrm{d}}_t)$ (Lemma C.1, part (i)).

(ii) We give the proof for P_t; that for P^{d}_t is identical.

The second inequality holds since $G_t(S_t \mid P_t, P^{\mathrm{d}}_t)$ is supermodular in P_t and S_t (Lemma 6, part (iii)). The first inequality holds by definition if $S^*_t(P_t) \le 1$. Otherwise, part (i) of the current lemma implies that $\Delta_{S_t} G_t(S^*_t(P_t) - 1 \mid P_t, P^{\mathrm{d}}_t) < 0$ and that if $\Delta_{S_t} G_t(S^*_t(P_t) - 2 \mid P_t + 1, P^{\mathrm{d}}_t) < 0$, then $S^*_t(P_t) - 2 < S^*_t(P_t + 1)$. Therefore, showing that $\Delta_{S_t} G_t(S^*_t(P_t) - 2 \mid$

$P_t + 1, P_t^{\mathrm{d}}) < \Delta_{S_t} G_t(S_t^*(P_t) - 1 \mid P_t, P_t^{\mathrm{d}})$ completes the proof:

$$
\begin{aligned}
&\Delta_{S_t} G_t(S_t^*(P_t) - 2 \mid P_t + 1, P_t^{\mathrm{d}}) - \Delta_{S_t} G_t(S_t^*(P_t) - 1 \mid P_t, P_t^{\mathrm{d}}) \\
&= \Delta_{S_t} G_t(S_t^*(P_t) - 2 \mid P_t + 1, P_t^{\mathrm{d}}) - \Delta_{S_t} G_t(S_t^*(P_t) - 2 \mid P_t, P_t^{\mathrm{d}}) \\
&\qquad + \Delta_{S_t} G_t(S_t^*(P_t) - 2 \mid P_t, P_t^{\mathrm{d}}) - \Delta_{S_t} G_t(S_t^*(P_t) - 1 \mid P_t, P_t^{\mathrm{d}}) \\
&= \Delta_{P_t} \Delta_{S_t} G_t(S_t^*(P_t) - 2 \mid P_t, P_t^{\mathrm{d}}) - \Delta_{S_t}^2 G_t(S_t^*(P_t) - 2 \mid P_t, P_t^{\mathrm{d}}) \\
&= (C^{\mathrm{u}} - C^{\mathrm{p}}) \mathbb{P}\left\{U_t = S_t^*(P_t) + P_t + P_t^{\mathrm{d}} - 1\right\} \\
&\qquad - \mathbb{P}\left\{U_t = S_t^*(P_t) + P_t + P_t^{\mathrm{d}} - 1\right\} \left[\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, P_t + P_t^{\mathrm{d}})\right] \\
&\qquad - (C^{\mathrm{h}} + C^{\mathrm{p}}) \mathbb{P}\left\{U_t = S_t^*(P_t) - 1\right\} \\
&\qquad - (C^{\mathrm{u}} - C^{\mathrm{p}}) \mathbb{P}\left\{U_t = S_t^*(P_t) + P_t + P_t^{\mathrm{d}} - 1\right\} \\
&\qquad + \sum_{k=2}^{P_t + P_t^{\mathrm{d}}} \mathbb{P}\left\{U_t = S_t^*(P_t) + k - 2\right\} \left[\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, k - 2) - \Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, k - 1)\right] \\
&\qquad + \mathbb{P}\left\{U_t = S_t^*(P_t) + P_t + P_t^{\mathrm{d}} - 1\right\} \left[\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, P_t + P_t^{\mathrm{d}} - 1)\right] \\
&\qquad - \mathbb{P}\left\{U_t = S_t^*(P_t) - 1\right\} \left[\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, 0)\right].
\end{aligned}
$$

The first and fourth terms cancel each other out, the second plus sixth term are together negative due to the convexity of $V_{t+1}(P_{t+1}, P_{t+1}^{\mathrm{d}})$ in P_{t+1}^{d} (Lemma 8, part (i)), and the same holds for the fifth term. The third plus seventh term are together negative, since $\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, 0) > -C^{\mathrm{h}}$ (Lemma 8, part (ii)).

□

Proof of Lemma 8. We only consider the properties for P_t^{d}, since we know that the same results hold for P_t.

For part (i), the fact that $V_t(P_t, P_t^{\mathrm{d}})$ is strictly decreasing in P_t^{d} can be seen as follows: If we assume a certain value for P_t^{d} and the corresponding optimal value for S_t, then we have $V_t(P_t, P_t^{\mathrm{d}})$. If we then increase P_t^{d} while we keep S_t constant, we know from Lemma 6, part (ii), that $G_t(S_t \mid P_t, P_t^{\mathrm{d}})$ will strictly decrease. $V_t(P_t, P_t^{\mathrm{d}})$ is the minimum of $G_t(S_t \mid P_t, P_t^{\mathrm{d}})$ over all possible values for S_t. Thus, it is necessarily lower at this higher value of P_t^{d}.

We will now bound $\Delta_{P_t^{\mathrm{d}}} V_t(P_t, P_t^{\mathrm{d}})$. We will use the bounds to prove part (ii) and the convexity in part (i). We know from Lemma 7, part (ii), that we need to consider two cases: if P_t^{d} is increased by one, then S_t^* either stays the same or decreases by one.

(1) If $S_t^*(P_t^{\mathrm{d}}+1) = S_t^*(P_t^{\mathrm{d}})$, then:

$$
\begin{aligned}
& G_t(S_t^*(P_t^{\mathrm{d}}) \mid P_t, P_t^{\mathrm{d}}+1) - G_t(S_t^*(P_t^{\mathrm{d}}) \mid P_t, P_t^{\mathrm{d}}) \\
& = \Delta_{P_t^{\mathrm{d}}} G_t(S_t^*(P_t^{\mathrm{d}}) \mid P_t, P_t^{\mathrm{d}}) \\
& = \sum_{k=S_t+P_t+P_t^{\mathrm{d}}+1}^{\infty} \mathbb{P}\left\{U_t = k\right\} \left[-(C^{\mathrm{u}} - C^{\mathrm{p}}) + \Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, P_t + P_t^{\mathrm{d}})\right] \\
& \geq -C^{\mathrm{h}} + (C^{\mathrm{p}} + C^{\mathrm{h}}) \sum_{k=S_t^*(P_t^{\mathrm{d}})+1}^{\infty} \mathbb{P}\left\{U_t = k\right\} \\
& \quad + \sum_{k=1}^{P_t+P_t^{\mathrm{d}}} \mathbb{P}\left\{U_t = S_t^*(P_t^{\mathrm{d}}) + k\right\} \left[\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, k-1)\right] \\
& \quad + \sum_{k=S_t^*(P_t^{\mathrm{d}})+P_t+P_t^{\mathrm{d}}+1}^{\infty} \mathbb{P}\left\{U_t = k\right\} \left[\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, P_t + P_t^{\mathrm{d}})\right],
\end{aligned}
$$

where the second equation follows from Lemma C.1, part (ii), and the inequality follows from Lemma 7, part (i).

(2) If $S_t^*(P_t^{\mathrm{d}}+1) = S_t^*(P_t^{\mathrm{d}}) - 1$, then:

$$
\begin{aligned}
&G_t(S_t^*(P_t^{\mathrm{d}}) - 1 \mid P_t, P_t^{\mathrm{d}}+1) - G_t(S_t^*(P_t^{\mathrm{d}}) \mid P_t, P_t^{\mathrm{d}}) \\
&= G_t(S_t^*(P_t^{\mathrm{d}}) - 1 \mid P_t, P_t^{\mathrm{d}}+1) - G_t(S_t^*(P_t^{\mathrm{d}}) \mid P_t, P_t^{\mathrm{d}}+1) \\
&\qquad + G_t(S_t^*(P_t^{\mathrm{d}}) \mid P_t, P_t^{\mathrm{d}}+1) - G_t(S_t^*(P_t^{\mathrm{d}}) \mid P_t, P_t^{\mathrm{d}}) \\
&= -\Delta_{S_t} G_t(S_t^*(P_t^{\mathrm{d}}) - 1 \mid P_t, P_t^{\mathrm{d}}+1) + \Delta_{P_t^{\mathrm{d}}} G_t(S_t^*(P_t^{\mathrm{d}}) \mid P_t, P_t^{\mathrm{d}}) \\
&= -C^{\mathrm{h}} + (C^{\mathrm{p}} + C^{\mathrm{h}}) \sum_{k=S_t^*(P_t^{\mathrm{d}})}^{\infty} \mathbb{P}\{U_t = k\} + (C^{\mathrm{u}} - C^{\mathrm{p}}) \sum_{k=S_t^*(P_t^{\mathrm{d}})+P_t+P_t^{\mathrm{d}}+1}^{\infty} \mathbb{P}\{U_t = k\} \\
&\qquad + \sum_{k=0}^{P_t+P_t^{\mathrm{d}}} \mathbb{P}\left\{U_t = S_t^*(P_t^{\mathrm{d}}) + k\right\} \left[\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, k)\right] \\
&\qquad + \sum_{k=S_t+P_t+P_t^{\mathrm{d}}+1}^{\infty} \mathbb{P}\{U_t = k\} \left[-(C^{\mathrm{u}} - C^{\mathrm{p}}) + \Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, P_t + P_t^{\mathrm{d}})\right] \\
&= -C^{\mathrm{h}} + (C^{\mathrm{p}} + C^{\mathrm{h}}) \sum_{k=S_t^*(P_t^{\mathrm{d}})}^{\infty} \mathbb{P}\{U_t = k\} \\
&\qquad + \sum_{k=0}^{P_t+P_t^{\mathrm{d}}} \mathbb{P}\left\{U_t = S_t^*(P_t^{\mathrm{d}}) + k\right\} \left[\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, k)\right] \\
&\qquad + \sum_{k=S_t^*(P_t^{\mathrm{d}})+P_t+P_t^{\mathrm{d}}+1}^{\infty} \mathbb{P}\{U_t = k\} \left[\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, P_t + P_t^{\mathrm{d}})\right].
\end{aligned}
$$

After the third equation, we use Lemma C.1, parts (i) and (ii); the first summation starts at $\big(S_t^*(P_t^{\mathrm{d}}) - 1\big) + 1$, the second summation at $\big(S_t^*(P_t^{\mathrm{d}}) - 1\big) + P_t + \big(P_t^{\mathrm{d}} + 1\big) + 1$, and the third summation ranges from $U_t = \big(S_t^*(P_t^{\mathrm{d}}) - 1\big) + 1$ to $U_t = \big(S_t^*(P_t^{\mathrm{d}}) - 1\big) + P_t + \big(P_t^{\mathrm{d}} + 1\big)$.

Case (1) gives the lower bound and case (2) gives the upper bound on $\Delta_{P_t^{\mathrm{d}}} V_t(P_t, P_t^{\mathrm{d}})$. Part (ii) follows from these bounds, by combining the lower bound with the fact that $-C^{\mathrm{h}} < \Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_t, P_t^{\mathrm{d}})$, and the upper bound with the fact that $\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_t, P_t^{\mathrm{d}}) < 0$. Notice that here, we use induction, i.e., for $t = T$, it is easily seen that these bounds hold.

To show convexity (part (i)), we need to compare the change in the optimal costs when P_t^{d} is increased by one with when it is again increased by one. We know from Lemma 7, part (ii), that we need to consider four cases only. If P_t^{d} is increased by one, then S_t^* stays the same or decreases by one. This holds also if P_t^{d} is again increased by one.

1. $S_t^*(P_t^{\mathrm{d}}+2) = S_t^*(P_t^{\mathrm{d}}+1) = S_t^*(P_t^{\mathrm{d}})$: Strict convexity follows from Lemma 6, part (ii).

2. $S_t^*(P_t^{\mathrm{d}}+2) = S_t^*(P_t^{\mathrm{d}}+1) - 1 = S_t^*(P_t^{\mathrm{d}}) - 1$: From case (2) above and the fact that $S_t^*(P_t^{\mathrm{d}}+1) =$

$S_t^*(P_t^{\mathrm{d}})$, we see that:

$$\begin{aligned}
&G_t(S_t^*(P_t^{\mathrm{d}}+1)-1 \mid P_t, P_t^{\mathrm{d}}+2) - G_t(S_t^*(P_t^{\mathrm{d}}+1) \mid P_t, P_t^{\mathrm{d}}+1)\\
&= -C^{\mathrm{h}} + (C^{\mathrm{p}}+C^{\mathrm{h}}) \sum_{k=S_t^*(P_t^{\mathrm{d}})}^{\infty} \mathbb{P}\{U_t = k\}\\
&\qquad + \sum_{k=0}^{P_t+P_t^{\mathrm{d}}+1} \mathbb{P}\left\{U_t = S_t^*(P_t^{\mathrm{d}})+k\right\} \left[\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, k)\right]\\
&\qquad + \sum_{k=S_t^*(P_t^{\mathrm{d}})+P_t+P_t^{\mathrm{d}}+2}^{\infty} \mathbb{P}\{U_t = k\} \left[\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, P_t+P_t^{\mathrm{d}}+1)\right]\\
&= -C^{\mathrm{h}} + (C^{\mathrm{p}}+C^{\mathrm{h}}) \sum_{k=S_t^*(P_t^{\mathrm{d}})}^{\infty} \mathbb{P}\{U_t = k\}\\
&\qquad + \sum_{k=0}^{P_t+P_t^{\mathrm{d}}} \mathbb{P}\left\{U_t = S_t^*(P_t^{\mathrm{d}})+k\right\} \left[\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, k)\right]\\
&\qquad + \sum_{k=S_t^*(P_t^{\mathrm{d}})+P_t+P_t^{\mathrm{d}}+1}^{\infty} \mathbb{P}\{U_t = k\} \left[\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, P_t+P_t^{\mathrm{d}}+1)\right].
\end{aligned}$$

(At the second equation, the term $\mathbb{P}\left\{U_t = S_t^*(P_t^{\mathrm{d}})+P_t+P_t^{\mathrm{d}}+1\right\}\left[\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, P_t+P_t^{\mathrm{d}}+1)\right]$ moves from the second summation to the third.) This is at least as high as the upper bound on $\Delta_{P_t^{\mathrm{d}}} V_t(P_t, P_t^{\mathrm{d}})$, which follows from case (2) above, by noticing that $\Delta_{P_{t+1}^{\mathrm{d}}} V_t(P_{t+1}, P_{t+1}^{\mathrm{d}})$ is convex in P_{t+1}^{d} (which follows from the current lemma, part (i)). This means that $\Delta_{P_t^{\mathrm{d}}} V_t(P_t, P_t^{\mathrm{d}})$ is non-decreasing in P_t^{d}, i.e., convex (except for $t = T$, all differences are strict).

3. $S_t^*(P_t^{\mathrm{d}}+2) = S_t^*(P_t^{\mathrm{d}}+1) = S_t^*(P_t^{\mathrm{d}})-1$: From case (2) above, we see that:

$$\begin{aligned}
&G_t(S_t^*(P_t^{\mathrm{d}})-1 \mid P_t, P_t^{\mathrm{d}}+1) - G_t(S_t^*(P_t) \mid P_t, P_t^{\mathrm{d}})\\
&= -C^{\mathrm{h}} + (C^{\mathrm{p}}+C^{\mathrm{h}}) \sum_{k=S_t^*(P_t^{\mathrm{d}})}^{\infty} \mathbb{P}\{U_t = k\}\\
&\qquad + \sum_{k=0}^{P_t+P_t^{\mathrm{d}}} \mathbb{P}\left\{U_t = S_t^*(P_t^{\mathrm{d}})+k\right\} \left[V_{t+1}(\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, k)\right]\\
&\qquad + \sum_{k=S_t^*(P_t^{\mathrm{d}})+P_t+P_t^{\mathrm{d}}+1}^{\infty} \mathbb{P}\{U_t = k\} \left[\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, P_t+P_t^{\mathrm{d}})\right].
\end{aligned}$$

This is at most as high as the lower bound on $\Delta_{P_t^{\mathrm{d}}} V_t(P_t, P_t^{\mathrm{d}}+1)$, which follows from case (1) above, the fact that $S_t^*(P_t^{\mathrm{d}}+1) = S_t^*(P_t^{\mathrm{d}})-1$, and noticing that $\Delta_{P_{t+1}^{\mathrm{d}}} V_t(P_{t+1}, P_{t+1}^{\mathrm{d}})$ is convex in P^{d}_{t+1} (which follows from the current lemma, part (i)). This means that $\Delta_{P^{\mathrm{d}}_t} V_t(P_t, P^{\mathrm{d}}_t)$ is non-decreasing in P^{d}_t, i.e., convex (except for $t = T$, all differences are strict).

4. $S^*_t(P^{\mathrm{d}}_t + 2) = S^*_t(P^{\mathrm{d}}_t + 1) - 1 = S^*_t(P^{\mathrm{d}}_t) - 2$: From case (2) above, we see that:

$$
\begin{aligned}
& G_t(S^*_t(P^{\mathrm{d}}_t) - 1 \mid P_t, P^{\mathrm{d}}_t + 1) - G_t(S^*_t(P^{\mathrm{d}}_t) \mid P_t, P^{\mathrm{d}}_t) \\
&= -C^{\mathrm{h}} + (C^{\mathrm{p}} + C^{\mathrm{h}}) \sum_{k=S^*_t(P^{\mathrm{d}}_t)}^{\infty} \mathbb{P}\{U_t = k\} \\
&\quad + \sum_{k=0}^{P_t + P^{\mathrm{d}}_t} \mathbb{P}\left\{U_t = S^*_t(P^{\mathrm{d}}_t) + k\right\} \left[\Delta_{P^{\mathrm{d}}_{t+1}} V_{t+1}(P_{t+1}, k)\right] \\
&\quad + \sum_{k=S^*_t(P^{\mathrm{d}}_t)+P_t+P^{\mathrm{d}}_t+1}^{\infty} \mathbb{P}\{U_t = k\} \left[\Delta_{P^{\mathrm{d}}_{t+1}} V_{t+1}(P_{t+1}, P_t + P^{\mathrm{d}}_t)\right].
\end{aligned}
$$

Next, from case (2) above and the fact that $S^*_t(P^{\mathrm{d}}_t + 1) = S^*_t(P^{\mathrm{d}}_t) - 1$, we see that:

$$
\begin{aligned}
& G_t(S^*_t(P^{\mathrm{d}}_t + 1) - 1 \mid P_t, P^{\mathrm{d}}_t + 2) - G_t(S^*_t(P^{\mathrm{d}}_t + 1) \mid P_t, P^{\mathrm{d}}_t + 1) \\
&= -C^{\mathrm{h}} + (C^{\mathrm{p}} + C^{\mathrm{h}}) \sum_{k=S^*_t(P^{\mathrm{d}}_t)-1}^{\infty} \mathbb{P}\{U_t = k\} \\
&\quad + \sum_{k=0}^{P_t + P^{\mathrm{d}}_t} \mathbb{P}\left\{U_t = S^*_t(P^{\mathrm{d}}_t) - 1 + k\right\} \left[V_{t+1}(\Delta_{P^{\mathrm{d}}_{t+1}} V_{t+1}(P_{t+1}, k)\right] \\
&\quad + \sum_{k=S^*_t(P^{\mathrm{d}}_t)+P_t+P^{\mathrm{d}}_t+1}^{\infty} \mathbb{P}\{U_t = k\} \left[\Delta_{P^{\mathrm{d}}_{t+1}} V_{t+1}(P_{t+1}, P_t + P^{\mathrm{d}}_t)\right].
\end{aligned}
$$

This means that $\Delta_{P^{\mathrm{d}}_t} V_t(P_t, P^{\mathrm{d}}_t)$ is strictly increasing in P^{d}_t, i.e., strictly convex.

□

Proof of Lemma 10. (i)

(a) For a given S_t, the difference in the per period costs (Equation (2)-Equation (1)) is:

$$
\begin{aligned}
&\sum_{k=0}^{S_t} \mathbb{P}\{U_t = k\} C^{\mathrm{h}}(S_t - k) + \sum_{k=1}^{P_t} \mathbb{P}\{U_t = S_t + k\} C^{\mathrm{p}} k \\
&\quad + \sum_{k=1}^{\infty} \mathbb{P}\{U_t = S_t + P_t + k\} (C^{\mathrm{u}} k + C^{\mathrm{p}} P_t) \\
&\quad - \sum_{k=0}^{S_t} \mathbb{P}\{U_t = k\} C^{\mathrm{h}}(S_t - k) - \sum_{k=1}^{\infty} \mathbb{P}\{U_t = S_t + k\} C^{\mathrm{u}} k \\
&= - \sum_{k=1}^{P_t} \mathbb{P}\{U_t = S_t + k\} (C^{\mathrm{u}} - C^{\mathrm{p}}) k - \sum_{k=1}^{\infty} \mathbb{P}\{U_t = S_t + P_t + k\} (C^{\mathrm{u}} - C^{\mathrm{p}}) P_t.
\end{aligned}
$$

Since $C^{\mathrm{u}} > C^{\mathrm{p}}$, this concludes the proof.

(b) For a given S_t, the difference in the per period costs (Equation (3)-Equation (2)) is:

$$
\begin{aligned}
&\sum_{k=0}^{S_t} \mathbb{P}\{U_t = k\} C^{\mathrm{h}}(S_t - k) + \sum_{k=1}^{P_t + P_t^{\mathrm{d}}} \mathbb{P}\{U_t = S_t + k\} (C^{\mathrm{h}} + C^{\mathrm{p}}) k \\
&\quad + \sum_{k=1}^{\infty} \mathbb{P}\left\{U_t = S_t + P_t + P_t^{\mathrm{d}} + k\right\} \left((C^{\mathrm{h}} + C^{\mathrm{u}}) k + (C^{\mathrm{h}} + C^{\mathrm{p}})(P_t + P_t^{\mathrm{d}})\right) \\
&\quad - \sum_{k=0}^{S_t} \mathbb{P}\{U_t = k\} C^{\mathrm{h}}(S_t - k) - \sum_{k=1}^{P_t} \mathbb{P}\{U_t = S_t + k\} (C^{\mathrm{h}} + C^{\mathrm{p}}) k \\
&\quad - \sum_{k=1}^{\infty} \mathbb{P}\{U_t = S_t + P_t + k\} \left((C^{\mathrm{h}} + C^{\mathrm{u}}) k + (C^{\mathrm{h}} + C^{\mathrm{p}}) P_t\right) \\
&= - \sum_{k=1}^{P_t^{\mathrm{d}}} \mathbb{P}\{U_t = S_t + P_t + k\} (C^{\mathrm{u}} - C^{\mathrm{p}}) k - \sum_{k=1}^{\infty} \mathbb{P}\left\{U_t = S_t + P_t + P_t^{\mathrm{d}} + k\right\} (C^{\mathrm{u}} - C^{\mathrm{p}}) P_t^{\mathrm{d}}.
\end{aligned}
$$

Since $C^{\mathrm{u}} > C^{\mathrm{p}}$, this concludes the proof.

(ii)

(a) If the per period costs are (non-strictly) lower for any given safety stock level, see part (i) (a), then that also holds for the safety stock level that is optimal when planned maintenance may not be delayed. The optimal costs when maintenance may be delayed at most once are at most equal to those costs.

(b) The reasoning is analogous to that at part (a).

(iii)

(a) To compare the optimal safety stock levels, we need to compare the first order differences with respect to the safety stock level (Lemma B.1, part (i) - Lemma A.1):

$$
\begin{aligned}
&C^{\mathrm{h}} - (C^{\mathrm{h}} + C^{\mathrm{p}}) \sum_{k=S_t+1}^{\infty} \mathbb{P}\{U_t = k\} - (C^{\mathrm{u}} - C^{\mathrm{p}}) \sum_{k=S_t+P_t+1}^{\infty} \mathbb{P}\{U_t = k\} \\
&\quad - C^{\mathrm{h}} + (C^{\mathrm{h}} + C^{\mathrm{u}}) \sum_{k=S_t+1}^{\infty} \mathbb{P}\{U_t = k\} \\
&= (C^{\mathrm{u}} - C^{\mathrm{p}}) \sum_{k=S_t+1}^{S_t+P_t} \mathbb{P}\{U_t = k\}.
\end{aligned}
$$

Since $C^{\mathrm{u}} > C^{\mathrm{p}}$, this shows that the first order difference with respect to the safety stock level is higher when delaying is allowed at most once. The optimal safety stock level is then non-strictly lower.

(b) To compare the optimal safety stock levels, we need to compare the first order differences with respect to the safety stock level (Lemma C.1, part (i) - Lemma B.1, part (i)):

$$
\begin{aligned}
&C^{\mathrm{h}} - (C^{\mathrm{h}} + C^{\mathrm{p}}) \sum_{k=S_t+1}^{\infty} \mathbb{P}\{U_t = k\} - (C^{\mathrm{u}} - C^{\mathrm{p}}) \sum_{k=S_t+P_t+P_t^{\mathrm{d}}+1}^{\infty} \mathbb{P}\{U_t = k\} \\
&- \sum_{k=1}^{P_t+P_t^{\mathrm{d}}} \mathbb{P}\{U_t = S_t + k\} \left[\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, k-1)\right] \\
&\quad - C^{\mathrm{h}} + (C^{\mathrm{h}} + C^{\mathrm{p}}) \sum_{k=S_t+1}^{\infty} \mathbb{P}\{U_t = k\} + (C^{\mathrm{u}} - C^{\mathrm{p}}) \sum_{k=S_t+P_t+1}^{\infty} \mathbb{P}\{U_t = k\} \\
&= (C^{\mathrm{u}} - C^{\mathrm{p}}) \sum_{k=S_t+P_t+1}^{S_t+P_t+P_t^{\mathrm{d}}} \mathbb{P}\{U_t = k\} - \sum_{k=1}^{P_t+P_t^{\mathrm{d}}} \mathbb{P}\{U_t = S_t + k\} \left[\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, k-1)\right].
\end{aligned}
$$

Since $C^{\mathrm{u}} > C^{\mathrm{p}}$ and $\Delta_{P_{t+1}^{\mathrm{d}}} V_{t+1}(P_{t+1}, k-1) < 0$ (according to Lemma 8, part (i)), this shows that the first order difference with respect to the safety stock level is higher when delaying is allowed an unconstrained number of times. The optimal safety stock level is then non-strictly lower.

□

D Proof for Section 8

Proof. This proof is relevant for both the case in which planned maintenance may never be delayed and the case in which planned maintenance may be delayed at most once. We next show, for the case in which planned maintenance may be delayed at most once, that the cost function can indeed be written as in Equation (8). From there, it is easily seen that the results in Lemma 3, part (i), and Theorem 4 also hold when $L > 0$, with the inventory position, known demands and unplanned maintenance replaced by their positive lead time equivalents, and with the ordering decision based on the inventory position. We first discuss the initial periods.

First, we consider the initial periods $t \in \{1, \dots, L\}$. We assume that, based on previous ordering decisions, P_t parts arrive in each of those periods. In other words, $\mathbf{S}_t = 0$, and the resulting expected costs in periods $t \in \{1, \dots, L\}$ are fixed, and thus they do not affect the optimal policy. In period 1, we will order enough to raise the inventory position, Y_1, to a level equal to the lead time known demands, $P_1^{\mathrm{d}} + U_1^{\mathrm{d}} + \mathbf{P}_{1+L}$, plus the safety stock level, $\mathbf{S}_{1+L}$.

We next focus our attention on the case in which planned maintenance may be delayed at most once. The cost function for periods $t + L \geq 1 + L$, with decisions being made in periods $t \geq 1$ can be written as:

$$C(\mathbf{S}_{t+L} \mid P_{t+L}) = \sum_{k=0}^{\mathbf{S}_{t+L}} \mathbb{P}\{\mathbf{U}_{t+L} = k\} C^{\mathrm{h}}(\mathbf{S}_{t+L} - k) + \sum_{k=1}^{\min\{P_{t+L}, O_{t+1}\}} \mathbb{P}\{\mathbf{U}_{t+L} = \mathbf{S}_{t+L} + k\} C^{\mathrm{p}} k$$
$$+ \sum_{k=1}^{\infty} \mathbb{P}\{\mathbf{U}_{t+L} = \mathbf{S}_{t+L} + \min\{P_{t+L}, O_{t+1}\} + k\} \left(C^{\mathrm{u}} k + C^{\mathrm{p}} \min\{P_{t+L}, O_{t+1}\}\right). \quad (9)$$

The term $\min\{P_{t+L}, O_{t+1}\}$, which appears three times in the equation, is required since in period $t + L$, we can potentially delay all of the planned maintenance, P_{t+L}. However, we can only do so if we know that, in the next period, enough new parts arrive to fulfill all of that delayed planned maintenance, since planned maintenance may be delayed at most once. The amount of stock that arrives in period $t + L + 1$ is O_{t+1}. Therefore, the maximum amount of planned maintenance that may be delayed in period $t+L$ is $\min\{P_{t+L}, O_{t+1}\}$. As a result, a myopic policy may not be feasible for this problem, since the cost function for period $t+L$ depends on our ordering decision in period t, O_t, (or, equivalently, $\mathbf{S}_{t+L}$), as well as the amount that is ordered in period $t+1$, O_{t+1}. However, we will show that under any optimal myopic policy, $O_{t+1} \geq P_{t+L}$, so that $\min\{P_{t+L}, O_{t+1}\} = P_{t+L}$, and the cost rate function in Equation (9) reduces to Equation (8). In other words, if we order myopically optimal in period t, ignoring how much we will order in period $t + 1$ (and thus ignoring how much planned maintenance we may be able to delay in period $t + L$), and we also order myopically optimal in period $t + 1$, then we will find that we can delay any amount of planned maintenance that we need to, and thus the myopically optimal order quantity is also optimal in general.

We now show that $O_{t+1} \geq P_{t+L}$. Since we have assumed that the amount of planned maintenance is stationary, showing that $O_{t+1} \geq P_{t+L}$ is equivalent to showing that $O_{t+1} \geq P_{t+L+1}$. This result should intuitively hold, since the inventory position is always sufficient to cover all known demands and there is no reason to order parts that are required in period $t + L + 1$ before period $t + 1$. We will next formally show that $O_{t+1} \geq P_{t+L}$ always holds.

We define $\hat{I}_t$ to be the stock on hand, or the shortage, at the beginning of period t, i.e.,

$\hat{I}_t = I_t - U_t^{\mathrm{d}} - P_t^{\mathrm{d}}$. From Equation (6), we know that the inventory position after ordering in period t, Y_t, will be equal to the lead time known demands plus the safety stock level:

$$Y_t = I_t + \sum_{r=t-L}^{t} O_r = P_t^{\mathrm{d}} + U_t^{\mathrm{d}} + \mathbf{P}_{t+L} + \mathbf{S}_{t+L},$$

so that:

$$\hat{I}_t + \sum_{r=t-L}^{t} O_r = \mathbf{P}_{t+L} + \mathbf{S}_{t+L} = \sum_{r=t}^{t+L} P_r + \mathbf{S}_{t+L}. \tag{10}$$

We can next write down a stock balance equation:

$$\begin{aligned}\hat{I}_{t+1} &= \hat{I}_t + O_{t-L} - P_t - U_t \\ &= \sum_{r=t+1}^{t+L} P_r - \sum_{r=t-L+1}^{t} O_r + \mathbf{S}_{t+L} - U_t.\end{aligned}$$

Since $\hat{I}_{t+1} + \sum_{r=t-L+1}^{t+1} O_r = \sum_{r=t+1}^{t+L+1} P_r + \mathbf{S}_{t+L+1}$ (by replacing t by $t+1$ in Equation (10)), after performing some algebra we have:

$$\mathbf{S}_{t+L} - U_t + O_{t+1} = P_{t+L+1} + \mathbf{S}_{t+L+1}.$$

Thus, since all costs and probability distributions for unplanned maintenance are stationary, $\mathbf{S}_{t+L} = \mathbf{S}_{t+L+1}$, so that $O_{t+1} \geq P_{t+L+1}$. $\square$