Warranty Matching in a Consumer Electronics Closed-Loop Supply Chain

INSEAD The Business School for the World®

Working Paper

2020/03/TOM

(Revised version of 2019/36/TOM)

Warranty Matching in a Consumer Electronics Closed-Loop Supply Chain

Andre P. Calmon
INSEAD, andre.calmon@insead.edu

Stephen C. Graves
Massachusetts Institute of Technology, sgraves@mit.edu

Stef Lemmens
INSEAD, stef.lemmens@insead.edu

Problem definition: We examine a dynamic assignment problem faced by a large Wireless Service Provider (WSP) that is a Fortune 100 company. This company manages two warranties: (i) a customer warranty that WSP offers its customers, and (ii) an Original Equipment Manufacturer (OEM) warranty that OEMs offer the WSP. The WSP uses devices refurbished by the OEM as replacement devices and hence their warranty operation is a closed-loop supply chain. Depending on the assignment policy WSP uses, the customer and OEM warranties might become misaligned for customer-device pairs, potentially incurring a cost for WSP.
Academic/practical relevance: We identify, model, and analyze a new dynamic assignment problem that emerges in this setting called the warranty matching problem. We introduce a new class of policies, called farsighted policies, which can perform better than myopic policies. We also propose a new heuristic assignment policy, the sampling policy, which leads to a near-optimal assignment. Our model and results are motivated by a real-world problem, and our theory-guided assignment policies can be used in practice; we validate our results using data from our industrial partner.
Methodology: We formulate the problem of dynamically assigning devices to customers as a discrete-time stochastic dynamic programming problem. Because this problem suffers from the curse of dimensionality, we propose and analyze a set of reasonable classes of assignment policies.
Results: The performance metric that we use for a given assignment policy is the average time that a replacement device under a customer warranty is uncovered by an OEM warranty. We show that our assignment policies reduce the average uncovered time and the expected number of out-of-OEM-warranty returns by more than 75% in comparison to our industrial partner's current assignment policy. We also provide distribution-free bounds for the performance of a myopic assignment policy, and of random assignment, which is a proxy for the current WSP policy.
Managerial implications: Our results indicate that, in closed-loop supply chains, being completely far-sighted might be better than being completely myopic. Also, policies that are effective in balancing short-term and long-term costs can be simple and effective, as illustrated by our sampling policy. We describe how the performance of myopic and farsighted policies depend on the size and length of inventory build-up.

Key words : Closed-Loop Supply Chains, Reverse-Logistics, Inventory Management, Sustainability.
History : Working Paper - November 2019.

1. Introduction

We introduce and analyze an assignment problem that occurs in multi-player closed-loop supply chains. We consider a supply chain with three players: an Original Equipment Manufacturer (OEM) that produces electronic devices, a retailer that sells the devices manufactured by the OEM and offers warranties to customers, and customers that purchase and use these devices and file a warranty claim when the device fails. A real-world system inspires this setting: our research emerged from a collaboration with a large Wireless Service Provider (WSP), a Fortune 100 company that acts as a retailer and sells mobile service contracts and mobile devices to their customers.

To manage such a system, our industrial partner faces the following three operational challenges: (*i*) forecasting the failure time distribution of new products; (*ii*) deciding how many devices should be kept in inventory to serve the warranty claims; and (*iii*) assigning devices in inventory to incoming customer warranty requests.

Although the three challenges are intertwined, we uncouple them and study them separately. In Calmon (2015) we address challenge (*i*) and introduce two strategies for forecasting the failure time distribution of new products. In Calmon and Graves (2017) we address challenge (*ii*), namely the determination of an inventory control policy to manage the inventory level at the WSP's reverse logistics facility over a product life time. In this paper we address challenge (*iii*), which we term *the warranty matching problem*, for which we take the inventory management policy as a given. Uncoupling the inventory management and warranty matching problems implies some degree of sub-optimization. However, uncoupling them leads to practical policies for which we can derive both theoretical and managerial insights. Finally, we note that the numerical experiments in Calmon (2015) indicate that making these decisions independently does not lead to a significant additional cost.

As discussed in Calmon and Graves (2017), there are two warranties in place in this closed-loop supply chain: (a) the customer warranty, offered by the retailer to the customer, and (b) the OEM warranty, provided by the OEM to the retailer. These warranties can have very different characteristics. In our partner WSP's case, as depicted in Figure 1, the customer warranty guarantees an immediate replacement of the faulty device with a working (new or refurbished) one, while the OEM warranty requires WSP to wait for the device repair, such that a replacement device is not sent immediately to WSP. If a customer covered by the customer warranty (which is usually 12 months in length) has a device that fails, WSP's reverse logistics facility sends a replacement from its inventory, which is largely refurbished devices. The remainder of the customer warranty applies to the replacement device. Hence, the customer warranty is associated with the customer, not the device that the customer possesses. Once the customer receives a replacement device, the customer sends back the failed device to WSP's reverse logistics facility.

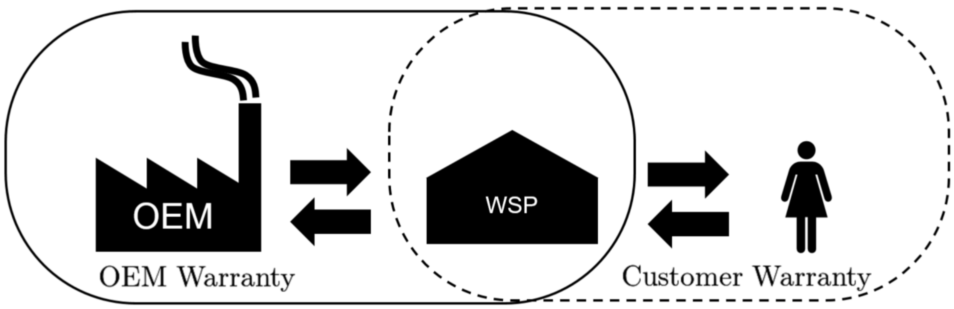

- Faulty device is shipped to OEM, fixed and only then returns to retailer
- Backlogging allowed
- Loss occurs

- Replacement device is sent immediately
- Replacement is usually a refurbished device
- No backlogging

Figure 1 Characteristics of the OEM warranty and of the customer warranty.

When the failed device arrives at the reverse logistics facility, WSP sends it to the OEM if it is still under OEM warranty (which is usually 12 months in length as well); and the device is then repaired by the OEM and returned to the centralized inventory held by WSP at its reverse logistics facility. If the device is no longer under OEM warranty, WSP either pays for the OEM to repair the device, or scraps the failed device. In the latter case, there is a unit loss in the inventory held by WSP for servicing warranty claims, and WSP may need to procure a replacement device for its inventory. Hence, in either case, there can be an additional cost to WSP, when the failed device is no longer under OEM warranty.

To minimize these costs, WSP wants all device-customer pairs to have matched warranties: that is, the remaining time on the customer warranty is (roughly) the same as (or less than) that for the OEM warranty on the device. When this is true, WSP can invoke the OEM warranty to repair the failed device from a customer warranty claim.

With new device sales the warranties are generally matched. At the time of sale, the customer has a 12-month warranty; even though the OEM warranty starts when WSP acquires the device, a device typically has close to 12 months remaining on the OEM warranty when it is sold to a customer.

However, these warranties can become unmatched when there is a warranty claim. When the customer makes a warranty claim, WSP must select a replacement device from the inventory at its reverse logistics facility to send to the customer. If the remaining OEM warranty on the replacement device is less than the time remaining on the customer warranty, the customer's device will be *uncovered* for some time, potentially leading to an *out-of-OEM-warranty return* for WSP. Furthermore, the OEM warranty does not pause during device repair or while the device is in inventory. Thus, a working device that is sitting in inventory continues to age and to consume its OEM warranty.

The practice at WSP has been to select the replacement device randomly, due to the large volume of transactions in its reverse logistics system. As such, the remaining OEM warranty on the replacement device has no relationship to the remaining time on the customer warranty. Since some devices can fail multiple times, randomly assigning replacement devices to customers can lead to a significant number of out-of-OEM-warranty returns. For the devices that we examine in this paper, WSP's data indicate that, between 15% and 25% of returned devices were out-of-OEM-warranty returns under WSP's current practice. The hundreds of out-of-OEM-warranty returns that WSP receives per day are a significant source of costs and operational overhead.

In this paper, we propose and analyze assignment strategies that mitigate the out-of-OEM-warranty returns in this system, i.e., minimize the number of replacement requests from customers still covered under the customer warranty, but whose failed device is not covered by the OEM warranty. In particular, we consider four classes of policies:

- *Random assignment policies*: we assign replacement devices to customers, ignoring the time remaining in both the customer and device warranties. This is representative of WSP's current policy for assigning devices to customers.
- *Myopic assignment policies*: we assign devices in inventory to customers so as to minimize the warranty mismatch in the current period. In particular, we examine the *Youngest-Out-First (YOF) assignment policy.*
- *Farsighted assignment policies*: we assign devices to customers so as to minimize the future expected cost (ignoring the current period cost). In WSP's case, we examine the *Oldest-Out-First (OOF) assignment policy.*
- *Heuristic policies*: we assign devices based on a policy that blends elements of the youngest-out-first policy and the oldest-out-first policy. We consider the *Sampling Policy*, where devices in inventory are sampled, sorted, and then matched to demand.

We display in Table 1 a summary of the performance of these policies on real data from WSP, in comparison to WSP's current practice for five device models. The performance metric is the *average uncovered time* of the customer-device pairs that are sent as replacements to customers that filed a warranty claim. If a customer has less time left in its customer warranty than the device has left in its OEM warranty, then the uncovered time is zero. Conversely, if the customer has more time left in its customer warranty than the device has left in the OEM warranty, then the uncovered time is the time difference between the end of the customer warranty and the end of the OEM warranty. In the appendix we argue that the average uncovered time can be interpreted as a first-order approximation to the mismatch cost. In Table 1 we obtain the lower bound by assuming that WSP knows all the future warranty claims and the time remaining on the customer warranty

Device	WSP's current policy	Random policy	YOF policy	OOF policy	Sampling policy	Lower bound
A	35.70	39.30	10.61 (73%)	8.89 (77%)	6.57 (83%)	2.85
B	35.20	40.79	9.39 (77%)	7.89 (81%)	5.43 (87%)	2.58
C	36.85	35.92	7.71 (79%)	7.15 (80%)	5.02 (86%)	3.14
D	25.86	34.54	7.53 (78%)	8.67 (75%)	5.55 (84%)	2.53
E	54.40	42.97	11.01 (74%)	12.59 (71%)	7.44 (83%)	2.85

Table 1 Average uncovered time per device (in days) and percentage reduction compared to the random assignment policy

of each future claim. Hence, this is a "clairvoyant" bound where all information is known and the problem becomes a large-scale deterministic assignment problem.

Table 1 indicates that WSP's current policy is consistent with random assignment, with the average uncovered time being on the order of a month or more. Furthermore, our policies significantly reduce the average uncovered time, with the sampling policy performing best. Note that there does not seem to be a strict ordering of the performance of the myopic YOF policy and the farsighted OOF policy. Finally, the performance of the sampling policy is close to the lower bound. In Section 5 we revisit Table 1 and discuss in detail how the performance of each policy depends on the problem settings. We also discuss the economic cost implications for WSP of switching the assignment policy and estimate potential savings on the order of $16 million per year, and a reduction of expected out-of-OEM-warranty returns between 73% and 91%.

1.1. A note on nomenclature

Throughout the paper we use the term *old customers* to refer to customers that have little time left in their customer warranty. Since the customer warranty has, in general, about a 12-month duration, an old customer is a customer with only 1 or 2 months of warranty left. Conversely, *young customers* are customers who still have most of their warranty left.

Similarly, we refer to a device as an *old device* if it has little time left in the OEM warranty, and as a *young device* if it has most of the OEM warranty left.

We define the *termination date*, of a customer or device, as the date that the customer or device warranty ends. Thus, if a customer's warranty termination date is $\bar{t}$ then, at time t, the customer has $(\bar{t}-t)^+$ time left in its warranty. We will also assume that both the customer and OEM warranty have equal length T_w. We define the *age* of a customer or device that is under warranty as the amount of the warranty that has passed. Thus, at time t, a customer or device with warranty termination date $\bar{t}$ has age $t-\bar{t}-T_w$. Finally, when a customer is beyond the customer warranty, the customer has no claims on the WSP, and is not included in our model and analysis.

The remainder of the paper is structured as follows. In Section 2, we present a literature review on matching. In Section 3, we formulate the warranty matching problem and in Section 4, we analyze different assignment strategies assuming a stationary failure time distribution. Finally, in Section 5, we evaluate the different policies through various numerical experiments.

2. Literature Review

Research in assigning items and individuals has been very active during the last 50 years. One of the seminal works in this area is on the Hungarian algorithm (Kuhn 1955). The stable matching problem, as discussed in Gale and Shapley (1962) and Gusfield and Irving (1989), has also received significant attention in the literature. A survey on the relationship between market design and matching problems can be found in Roth and Sotomayor (1992), and an application to matching medical students to residency programs is covered in Roth (2003). More recent applications include online dating (Hitsch et al. 2010), child adoption (Slaugh et al. 2016), and refugee resettlement (Aziz et al. 2018). The majority of papers in the assignment literature address a static problem, while we consider a dynamic assignment problem. Dynamic assignment models have been used to optimize kidney exchange (Ünver 2010), day care assignment (Kennes et al. 2014), refugee resettlement (Andersson et al. 2018), and house allocation (Kurino 2014).

In the operations management literature, a classic application of matching is blood bank inventory management (Jennings 1973, Pegels and Jelmert 1970, Cohen and Pierskalla 1975, Jagannathan and Sen 1991, Sarhangian et al. 2018). The role of matching strategies in kidney-exchange programs has also been an active area of research. A framework for studying this problem from a game-theoretic point of view can be found in Ashlagi et al. (2011) and in Ashlagi et al. (2015). Furthermore, queuing theory, more specifically so-called "double matching queues", is a proposed methodology by Boxma et al. (2011) to manage waiting lists for organ transplants. However, the proposed matching strategies in these papers are not readily applicable to closed-loop systems.

During the last few years, matching policies have received considerable attention to improve ridesharing platforms such as Uber and Lyft. A myopic policy commonly used in practice is to offer an arriving customer the closest driver. Korolko et al. (2018) and Özkan and Ward (2018) demonstrate that such a myopic policy to match drivers and customers may perform poorly. The study of Hu and Zhou (2016), which is applicable to ridesharing, considers a two-sided, discrete-time matching system where both supply and demand can abandon the system. Similar to our work, the optimal policy is state-dependent and complex.

We analyze a matching problem in the context of a closed-loop supply chain. Recent overviews of the literature on closed-loop supply chains can be found in Guide and Van Wassenhove (2009) and Souza (2013) and a recent typology of remanufacturing processes can be found in Abbey and Guide (2018). In contrast to Savaskan et al. (2004) and Galbreth and Blackburn (2006), we do not model in detail the collection of cores or the decision related to which cores should be remanufactured or scrapped (Calmon and Graves (2017) addresses such concerns). Furthermore, the type of matching problem that we consider in this paper has yet to be covered in the closed-loop supply chain literature. A multiperiod single-product inventory problem which faces demand both for new items and for warranty replacement is found in Huang et al. (2008). Similar to our setting, the number and ages of the failed items under warranty are available. However, we take the closed-loop nature into account by considering the repair and refurbishment of the failed items for reuse and matching. Pince et al. (2016) demonstrate how an OEM should dynamically allocate customer returns between fulfilling warranty claims and remarketing refurbished products over the product's life-cycle, but do not examine the problem of assigning devices to customers.

Although we build upon these streams of research, our work considers a distinct setting. Furthermore, our set-up is dynamic, and decisions made in one period impact the system's state in the future. Our goal is not to find an equilibrium allocation, but to find a cost minimizing assignment. Hence, our work also draws from and contributes to the literature on online allocation. An example of an online algorithm for minimum metric bipartite matching is Meyerson et al. (2006). Also, in Bansal et al. (2007), Manshadi et al. (2012) and Truong and Wang (2018), online algorithms for bipartite matching and their competitive ratios are examined. Recently, Huang et al. (2018) obtains tight competitive ratios for two algorithms of the "fully online matching problem", which is a generalization of the classic online bipartite matching problem. A review of tools for analyzing online algorithms, which encompass the type of policies that we use in this paper, can be found in Albers (2003).

3. The Warranty Matching Problem

We formulate the problem of dynamically assigning devices to customers as a discrete-time stochastic dynamic programm for which, in each period, customers of different ages place requests for replacement devices, and devices in the centralized inventory are then assigned to these customers. For our partner WSP, a time period is one day.

We assume that the inventory management policy is exogenous to the problem of assigning devices to customers, and is thus an input to our warranty matching model. This inventory management policy determines when to source new devices into inventory, as well as when to sell excess inventory through a side sales channel, as discussed in Calmon and Graves (2017). Furthermore, the policy assumes that all demand is met: if there is an insufficient supply of inventory, the customer will receive a new device or an upgrade as a replacement. In modeling the matching problem we do not explicitly account for the costs or revenues associated with the inventory management policy; that is, we do not include the costs for sending upgrades (or new phones) when there are shortfalls, nor the revenues from side sales.

We present the components of our discrete-time dynamic optimization model and the sequence of events below. We assume that the OEM warranty can have m possible termination dates, while the customer warranty can have n possible termination dates.

Customer demand process. We denote by d_{ij}^t the number of customers that file a warranty replacement request at the start of day t that have a customer warranty termination date $j \geq t$ and are returning a device with OEM warranty termination date i. We assume d_{ij}^t is stochastic with known distribution. We denote the total demand for replacements in day t by d^t and the matrix with the demand at time t by $\boldsymbol{d}_t = (d_{ij}^t)$.

Inventory state. The amount of inventory in the beginning of period t *before device arrivals* is given by $\boldsymbol{y}_t = (y_1^t, \ldots, y_m^t)$ where y_i^t is the number of devices in inventory at time t with OEM warranty termination date i. We denote the total inventory at time t by y^t.

Device arrival process. For simplicity, we assume there is a fixed deterministic lead time l that covers the time for the customer to return the failed device to the reverse logistics facility, plus the time to refurbish or repair the device. We also assume that each returned device that is covered by the OEM warranty has a probability α of being repairable[1] (with probability 1-α the device is not repaired and leaves the system). We denote the number of repaired devices that arrive to inventory at time t and that have an OEM warranty termination date i by the random variable ξ_i^t, which follows a Binomial distribution with parameters $(\sum_{j=1}^n d_{ij}^{t-l}, \alpha)$. Let ξ^t be the total number of repaired devices arriving into inventory at time t, i.e. $\xi^t = \sum_{i=1}^m \xi_i^t$.

Furthermore, we assume that there are two additional streams of exogenous device arrivals into inventory. The first is seed-stock from the OEM, often as specified by a contractual agreement between the OEM and WSP (usually 1% of sales). We denote the seed-stock with OEM warranty termination date i arriving in period t by s_i^t. The second stream are due to the inventory management policy, e.g., as described in Calmon and Graves (2017). According to this policy, on the one hand, WSP adds devices to inventory if the inventory level at the beginning of the period together with the arrivals of refurbished devices and seed-stock are deemed to be insufficient to efficiently satisfy demand. On the other hand, if WSP deems that it has excess inventory, a certain number of devices in inventory is sold through a side sales channel. We denote the net change to inventory at time t due to the inventory management decisions by b_i^t, for devices with OEM warranty termination date i, where we allow $b_i^t < 0$ to capture side sales. Thus, we denote the arrival of devices by $\boldsymbol{a_t} = (a_1^t, \ldots, a_m^t)$ where

$$a_i^t = \xi_i^t + s_i^t + b_i^t, \forall i = 1, \ldots, m.$$

As noted above, this specification includes side sales, and hence we allow the arrivals to be negative. WSP cannot sell more devices through the side sales channel than the inventory available. Thus,

[1] A device is *repairable* if the purported failure is covered under the OEM warranty.

$y_i^t + a_i^t \geq 0$, $\forall i$ (or, equivalently, $y_i^t + \xi_i^t + s_i^t + b_i^t \geq 0$). We observe that for any side sales quantity, the optimal side sales policy will sell the oldest devices in inventory. We denote the total arrivals into inventory at time t by a^t.

As discussed in Calmon and Graves (2017), we assume that all demand in a period has to be satisfied (no backlogging is allowed). Thus, $y^t + a^t \geq d^t$. We also note that demand for replacement devices and the arrival of devices into inventory are not independent.

Assignment decision. In each period, WSP must assign inventory to warranty replacement requests. We define the decision variable x_{ij}^t as the number of devices with OEM termination date i that are allocated to customers with customer warranty termination date j. Thus, in each period t, $\boldsymbol{x}_t = (x_{ij}^t) \in \mathbb{R}^{m \times n}$ must satisfy the demand/supply constraints

$$\sum_{i=1}^{m} x_{ij}^t = \sum_{i=1}^{m} d_{ij}^t, \forall j = 1, \ldots, n,$$
$$\sum_{j=1}^{n} x_{ij}^t \leq y_i^t + a_i^t, \forall i = 1, \ldots, m,$$
$$x_{ij}^t \geq 0, \forall i = 1, \ldots, m, \text{ and } \forall j = 1, \ldots, n.$$

We denote the set formed by these constraints by $\mathcal{X}_t(\boldsymbol{y}_t, \boldsymbol{a}_t, \boldsymbol{d}_t)$.

Inventory update (state equation). We assume the following sequence of events in each period: (*i*) Refurbished devices and seed-stock arrive into inventory, along with any side sales; (*ii*) Demand for replacements is observed and new devices are added into inventory if needed to avoid any shortfalls; (*iii*) Devices are assigned to customers. The inventory update equation at time t is

$$y_i^{t+1} = y_i^t + a_i^t - \sum_{j=1}^{n} x_{ij}^t, \forall i = 1, \ldots, m.$$

Matching costs. We denote the cost of assigning a device with OEM warranty termination date i to a customer with warranty termination date j at time t by $c_t(i,j)$. This cost represents the expected cost of an out-of-OEM-warranty return. If $i \geq j$, then this cost is zero, as there cannot be an out-of-OEM-warranty return; but if $i < j$, then the OEM warranty ends before the customer warranty, and there is some chance of an out-of-OEM-warranty return. Specifically, there is an out-of-OEM-warranty return if the next device failure occurs in the interval $(\max(i,t), j]$. We assume that WSP's cost is proportional to the *uncovered warranty time* $(j - \max(i,t))^+$. Namely, for some $\beta \geq 0$, we assume

$$c_t(i,j) = \beta \cdot (j - \max(i,t))^+. \quad (1)$$

We illustrate this uncovered mismatch cost in Figure 2 for $\beta = 1$.

Effectively in (1) we are assuming that the likelihood that the next device failure occurs in the time interval $(\max(i,t), j]$ is proportional to the length of the interval. In Appendix EC.1 we provide

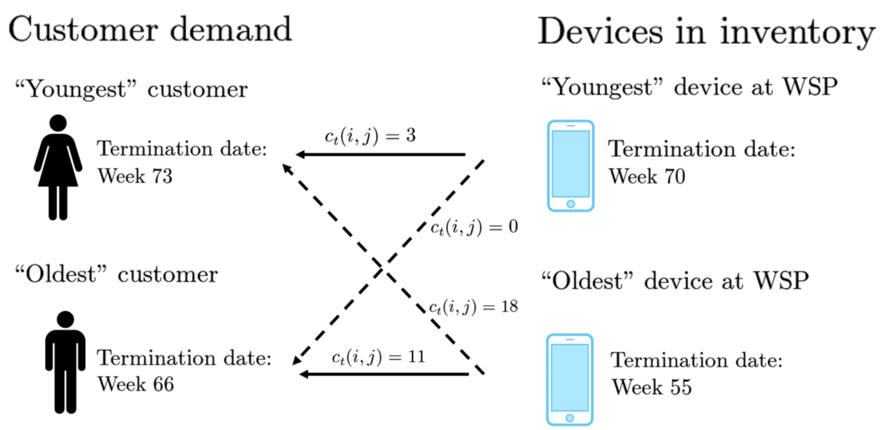

Figure 2 **Example of the uncovered mismatch cost for two replacement devices being assigned at time period t=50. The solid lines represent the optimal assignment and leads to a mismatch cost of 14 while the total cost of the assignment in dashed lines equals 18.**

some justification for this assumption when device failure times follow an Exponential or Weibull distribution. All our analysis would hold if the cost were of the form $c_t(i,j) = f\left((j - \max(i,t))^+\right)$ where f is a convex increasing positive function.

Dynamic program formulation. In theory we can formulate and solve this closed-loop assignment problem as a dynamic program. To do so we assume that we have a finite planning horizon of T periods, corresponding to the life-cycle for the device. We denote J_t as the cost-to-go at time t, with the boundary condition $J_{T+1} = 0$. We can obtain the optimal allocation policy by solving

$$J_t(\boldsymbol{y}_t, \boldsymbol{a}_t, \boldsymbol{d}_t) = \underset{\boldsymbol{x}_t \in \mathcal{X}_t(\boldsymbol{y}_t, \boldsymbol{a}_t, \boldsymbol{d}_t)}{\text{minimize}} \quad \sum_{i=1}^{m} \sum_{j=1}^{n} c_t(i,j) x_{ij}^t + E[J_{t+1}(\boldsymbol{y}_{t+1}, \boldsymbol{a}_{t+1}, \boldsymbol{d}_{t+1})] \tag{2}$$

where $\boldsymbol{y}_{t+1}$ is a function of $(\boldsymbol{y}_t, \boldsymbol{x}_t, \boldsymbol{a}_t)$ and the expectation is taken over $(\boldsymbol{a}_{t+1}, \boldsymbol{d}_{t+1})$.

We cannot realistically solve (2), due to the curse of dimensionality. The state space is very complex, as it includes, for each period, the on-hand inventory delineated by termination date, as well as the realization of the arrivals of devices and customer requests. Actually this is already a simplification, as the state space should also include the distribution of device ages in use by customers, as well as the number of devices currently being repaired and refurbished with the OEM. In this sense, we view (2) as an approximate dynamic program. Still, for any realistic specification of the given state space, it is much too large to attempt to solve.

Nevertheless this formulation provides some insight that can guide the development of implementable policies for our partner WSP. The computation of the cost-to-go function at time t entails the sum of two terms: "the current period cost" and "the future cost-to-go". A myopic policy considers just the current period cost, while a farsighted policy optimizes the future cost-to-go.

However, by design, neither policy can guarantee the minimization of the system-wide costs. Therefore, in addition, we propose a heuristic policy which balances the current period cost with the future cost-to-go.

4. Assignment Policies

In this section we introduce and examine properties of various assignment policies. We first examine the **random assignment policy** where devices are assigned to customers ignoring the time remaining in both the customer and OEM warranties. This is representative of WSP's current policy for assigning devices to customers. We bound the average behavior of this policy, and this bound is independent of the demand distribution and can be used by practitioners as a starting point to decide if a more structured assignment policy should be pursued or not.

We then examine a policy within the class of **myopic assignment policies**. Specifically, at time t a myopic policy solves

$$\underset{\boldsymbol{x}_t \in \mathcal{X}_t(\boldsymbol{y}_t, \boldsymbol{a}_t, \boldsymbol{d}_t)}{\text{minimize}} \quad \sum_{i=1}^{m} \sum_{j=1}^{n} c_t(i,j) x_{ij}^t. \tag{3}$$

Within the class of myopic policies, we analyze the *Youngest-Out-First (YOF) assignment policy* and prove a worst-case bound on its average performance.

The third policy we analyze, called the *Oldest-Out-First (OOF) assignment policy*, belongs to the class of **farsighted assignment policies**. In every period t, we minimize the future expected cost:

$$\underset{\boldsymbol{x}_t \in \mathcal{X}_t(\boldsymbol{y}_t, \boldsymbol{a}_t, \boldsymbol{d}_t)}{\text{minimize}} \quad E[J_{t+1}(\boldsymbol{y}_{t+1}(\boldsymbol{y}_t, \boldsymbol{x}_t, \boldsymbol{a}_t), \boldsymbol{a}_{t+1}, \boldsymbol{d}_{t+1})]. \tag{4}$$

In general, solving the problem above has the same complexity as solving the full dynamic program. However, in our setting, the cost function and the state space follow a specific structure, and the problem in (4) has a simple solution. The OOF policy is the policy induced by this solution.

Finally, we introduce and examine a **heuristic policy** which blends the YOF and OOF policies. This heuristic policy is the *Sampling Policy*, where in every period t we sample d^t devices from inventory, and then sort and match to incoming demand.

The structure of the cost function and state space that we exploit in our analysis is, simply put, that *the younger the devices in inventory, the lower the expected assignment cost*. Formally, let $\preceq$ be a partial order such that for any two inventory vectors $\boldsymbol{y}_t$ and $\tilde{\boldsymbol{y}}_t$ where $y^t = \tilde{y}^t$ we have that

$$\boldsymbol{y}_t \preceq \tilde{\boldsymbol{y}}_t \quad \text{if} \quad \sum_{i=k}^{m} y_i^t \geq \sum_{i=k}^{m} \tilde{y}_i^t, \forall k = 1, \ldots, m.$$

Then, for every period t, since the cost function $c_t(i,j)$ given in (1) is non-increasing in i we can argue that for inventory vectors $\boldsymbol{y}_t$ and $\tilde{\boldsymbol{y}}_t$, and arrival vectors $\boldsymbol{a}_t$ and $\tilde{\boldsymbol{a}}_t$, where $y^t + a^t = \tilde{y}^t + \tilde{a}^t$ we have

$$\boldsymbol{y}_t + \boldsymbol{a}_t \preceq \tilde{\boldsymbol{y}}_t + \tilde{\boldsymbol{a}}_t \implies J_t(\boldsymbol{y}_t, \boldsymbol{a}_t, \boldsymbol{d}_t) \leq J_t(\tilde{\boldsymbol{y}}_t, \tilde{\boldsymbol{a}}_t, \boldsymbol{d}_t), \forall \boldsymbol{d}_t. \tag{5}$$

We will use this observation to describe the behavior of the YOF and OOF policies, and also to discuss the intuition behind each heuristic.

For the analysis and results presented next, we model the customer warranty termination date for each customer that files a warranty request at time t as a random variable denoted by D_t. We assume that the distribution of D_t does not change dramatically from one period to the next. Formally, we assume that

$$\max_{k=1,\ldots,m} |\Pr(D_{t+1} \leq k) - \Pr(D_t \leq k)| \leq \epsilon, \forall t. \tag{6}$$

Namely, we assume that the Kolmogorov-Smirnov statistic of the termination date distribution of demand changes by at most ϵ from one period to another. For WSP, a time period corresponds to a day and we estimate that $\epsilon \leq 0.01$ for all the devices for which we had data.

In addition we assume that $D_1, \ldots, D_T$ are independent random variables, but not necessarily identically distributed. Finally, for our theoretical results, we assume that, the OEM warranty termination date of each device returned by a customer for refurbishment is the same as the customer warranty termination date. In effect we assume that the two warranties are matched prior to any device failures.

4.1. Random Assignment

Due to the large volume of incoming refurbished devices and requests (on the order of tens of thousands per day across a couple of dozen device models), WSP does not sort devices nor customer requests. In fact, as refurbished devices are received from the OEM, the devices are stored randomly in large boxes or containers in inventory at WSP's reverse logistics facility. As these refurbished devices are held separately from the seed-stock received from the OEM, WSP can prioritize one class of inventory relative to the other. However, in order to describe this scenario from a conceptual level, we will not differentiate between seed-stock and refurbished devices but, as a simplified model, will examine the case where devices in inventory are randomly assigned to replacement requests. Our intent is to develop a benchmark, as well as to analyze the potential downsides of random assignment.

In order to obtain the worst-case performance of random assignment, we will make a few simplifying assumptions. We assume no side sales ($b^t \geq 0$), that when a new device is sold its customer warranty and OEM warranty start at the same time, and that all returned devices are repaired

($\alpha = 1$). Relaxing these assumptions does not change our insights but significantly increases the amount of notation we have to use. Finally, we assume that in each period the total demand and inventory are constants, namely $d^t = d, y^t = y, \forall t$ and that $y \geq 0$.

For each period t, *after device arrivals*, we model the OEM warranty termination date of a randomly-chosen device in inventory by a random variable Y_t. Hence, for a random assignment policy, the expected cost for each assignment in period t is

$$E[c_t(Y_t, D_t)] = \beta \cdot E[(D_t - \max(Y_t, t))^+].$$

Due to the closed loop system, and the assumption that the warranties are initially matched, we can model the random variable Y_t in terms of D_t. Each device in inventory in period t was returned by a customer at time $t - l - Z$ for refurbishment, where Z is a random variable that represents the amount of time a device remains in inventory. Hence, we can model the random variable Y_t in terms of random variables D_t and Z, namely $Y_t = D_{t-l-Z}$. If devices in inventory are chosen at random to be used as replacements then, in every period t, a device in inventory has a probability $d/(y+d)$ of being used as a replacement and probability $1 - d/(y+d)$ of being carried over to the next period. In the limit as t increases, Z becomes a geometric random variable with parameter $d/(y+d)$. Since we assume a fixed lead time l, we can express the expected cost as

$$E[c_t(Y_t, D_t)] = E[c_t(D_{t-l-Z}, D_t)] = \beta \cdot E[(D_t - \max(D_{t-l-Z}, t))].$$

We assume the limiting distribution for Z to develop the worst-case upper bound on the expected cost in the proposition below. All our proofs are in the Appendix.

Proposition 1. *Let D_t be a random variable that represents the customer warranty termination date of a randomly chosen customer at time t. Furthermore, assume that the distributions of $\{D_t\}$ satisfy the distance constraint given in (6). Let l be the lead time and Z be a geometric random variable with parameter $d/(y+d)$. Then, assuming a warranty length of T_w, and assuming $d_t = d$, $y_t = y, \forall t \in [1, T]$, the following distribution-free bound holds*

$$E[c_t(Y_t, D_t)] \leq \frac{(1 + a_r \cdot \epsilon)^2 + b_r \epsilon^2}{4} \cdot \beta \cdot T_w,$$

where $a_r = l + \frac{y+d}{d}$ and $b_r = \frac{y(y+d)}{d^2}$. This bound is tight, i.e., there is a distribution of D_t that achieves this bound. Also, let U_t^R be a random variable that represents the uncovered time of a customer that requires a replacement device at time t under the random assignment policy. Then, the average uncovered time satisfies

$$\frac{1}{T} \sum_{t=1}^{T} E[U_t^R] \leq \frac{(1 + a_r \cdot \epsilon)^2 + b_r \epsilon^2}{4} \cdot T_w.$$

For WSP, we estimate that $\epsilon = 0.2\%$ per day, and we assume that the lead time is 3 weeks and $T_w = 12$ months. In practice, WSP's total demand per period and the total inventory per period are not stationary. However, for the inventory management policy in our simulations, we note that, in general, the carry-over inventory at the start of any given day is usually less than the total demand in that day, i.e. $y^t \leq d^t, \forall t$ (although it must be that $y^t + a^t \geq d^t$ since no backlogging is allowed). Since the bound in the proposition is increasing in y, if we assume $y^t = d^t$ (which is the worst case for WSP's data) we have $\frac{1}{T}\sum_{i=1}^{T} E[U_t^R] \leq 100$ days.

The bound in Proposition 1 only has positive probability weights on the first and last termination dates of the distribution's support (i.e. it is a variance-maximizing discrete distribution) which does not match at all the actual distribution of termination dates at WSP. Thus, the bound is loose. However, if a practitioner deems that 100 days of uncovered time were acceptable, a more sophisticated policy or system simulation would not be required.

With these results in hand, we now proceed to analyze the *Youngest-Out-First* assignment policy, where devices in inventory and requests are sorted by age and matched youngest to oldest.

4.2. Myopic Assignment: The Youngest-Out-First Policy

As mentioned in the beginning of this section, myopic assignment policies solve, for each period t, the allocation problem in (3). If the costs $c_t(i,j)$ are as described in (1), then the *Youngest-Out-First* (YOF) policy solves the problem in (3) for any $(\boldsymbol{y}_t, \boldsymbol{a}_t, \boldsymbol{d}_t)$ where $y^t + a^t \geq d^t$. This policy is depicted in Figure 3 and described in Algorithm 1.

Algorithm 1 Youngest Out First Policy

for all t **do**

1) Sort customer requests by customer warranty termination date (ignore the age of the device that the customer is returning)

2) Sort devices in inventory by OEM warranty termination date

3) Match the youngest device in inventory (with highest OEM warranty termination date) to the youngest customer (with highest customer warranty termination date).

4) Repeat Step 3 until all customer requests are fulfilled.

end for

For a single period, the YOF policy is optimal, as stated in the following proposition.

Proposition 2. *The Youngest-Out-First policy is optimal for the problem in Equation 3.*

Suppose in period t that the total demand equals the total arrivals and any inventory, i.e., $a^t + y^t = d^t$. In this case all available inventory in period t is used to fulfill demand and hence, $y^{t+1} = 0$ independent of the policy used at time t. Thus, the assignment in period t has no effect on

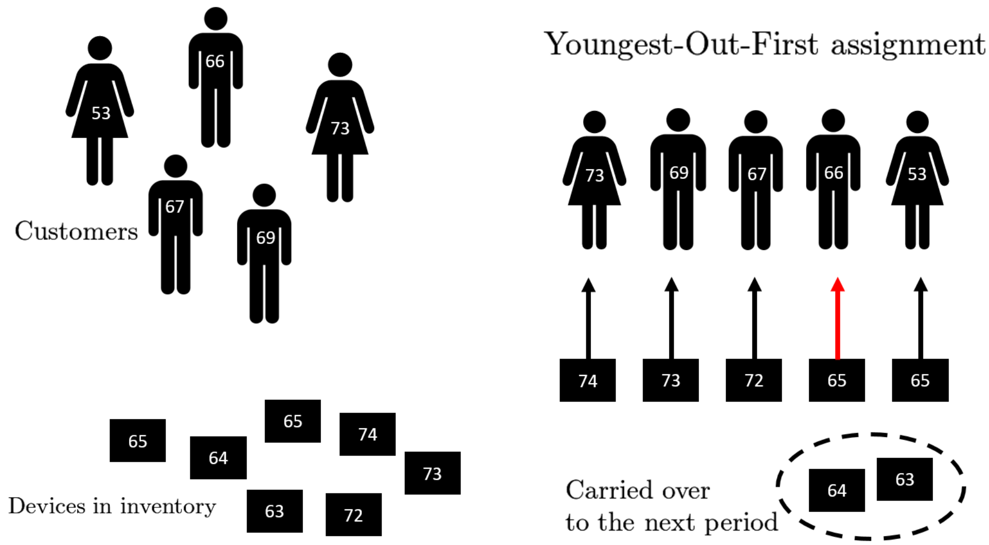

Figure 3 **Depiction of the Youngest-Out-First assignment strategy. In this example, $y^t + a^t = 7$ and $d^t = 5$. Customers and devices are sorted by termination date and matched youngest to oldest. Unused devices in inventory are carried over to the next period and age. Note that the only assignment that incurs a cost is the assignment of the device with OEM warranty termination date 65 to the customer with customer warranty termination date 66.**

the assignments in periods $[t+1, T]$. The YOF policy is then optimal in period t since it minimizes the costs in the current period.

The YOF policy will also have an adequate practical performance in cases where there is inventory buildup and some of that inventory never needs to be used for replacement, e.g., when the number of device arrivals exceeds the demand for replacement over the device's life-cycle and excess inventory is sold through side sales. In fact, under certain assumptions, we can bound the expected uncovered warranty time for a given period as stated in the following proposition.

Proposition 3. *Let U_t^{MY} be a random variable that represents the uncovered time of a customer chosen at random at time t that requires a replacement device under some myopic policy. If $a^t \geq \xi^t \geq d^t$, then for any distribution of termination dates of the demand, we have, for a warranty length of T_w, that*

$$E[U_t^{MY}] \leq T_w \left(\frac{1}{\sqrt{d^t}} + l\epsilon \right).$$

For WSP, $\epsilon = 0.2\%$ per day, $T_w = 1$ year and $l = 21$ days. Using the proposition above, in periods where $\xi^t \geq d^t$, the expected uncovered time under the myopic policy is at most $365 \left(\frac{1}{\sqrt{d^t}} + 0.042 \right)$ days. If $d^t = 250$, which is about the average daily replacement demand for many devices sold by WSP, then the average uncovered time in period t is less than 38 days. Although this bound requires $\xi^t \geq d^t$, it is very simple to calculate and can be used as a starting point to examine the benefits of implementing a myopic assignment policy.

Conversely, if inventory and arrivals exceed total demand, i.e., $a^t + y^t > d^t$, inventory is carried over and using the YOF policy could lead to a poor inventory state. Then, a high assignment cost might be incurred if the carry-over inventory has to be used in some future period. For example, assume for some interval $[t_1, t_2]$ that $b^t \geq 0, \forall t \in [t_1, t_2]$ (there are no side sales) and $a^t = d^t$, i.e., the total number of arrivals into inventory in each period is the same as total demand, and the number of devices in inventory remains the same over this interval. Then, under the YOF policy, the inventory gets older over this interval, namely $\boldsymbol{y}_{t_1} \preceq \boldsymbol{y}_{t_1+1} \preceq \ldots \preceq \boldsymbol{y}_{t_2}$ for any realization of demand and arrivals. Note that if in any period t there were an arriving device that was older than a device in inventory, the YOF policy would use the younger device in inventory and carry over the older arriving device. Hence, although the myopic policy minimizes the current period costs, it will accumulate "old" devices in inventory, which is a downside of this policy. If in some period $t > t_2$ there is some shock on total demand and $d^t > a^t$, some of the older inventory will be used, leading to high assignment costs.

As an alternative to the class of myopic policies, we consider next the class of farsighted policies, which aims to put the system in the best possible inventory state.

4.3. Farsighted Assignment: The Oldest-Out-First Policy

Myopic polices do not take into account information about future demand. This may lead to policies that are simple and, if well designed, have reasonable practical performance. In this section, we examine a policy that does the opposite of the myopic policy: instead of taking into account costs in the current period, it *only takes into account future costs*, i.e., in every period this policy solves the problem in (4). Thus, this is a *farsighted policy.*

In general, farsighted policies are not optimal and, at a first glance, computing a farsighted policy might seem to have the same complexity as computing the optimal policy. However, due to the one-sided cost structure at WSP, and the resulting partial order of the state space described in (5), we can easily derive a farsighted policy for the problem, which we call the *Oldest-Out-First (OOF) policy.* Like the YOF policy, the OOF policy also sorts devices in inventory by warranty termination date but now it matches devices in inventory to demand from *oldest to youngest.* The OOF policy is described in Algorithm 2.

We will first show that the OOF policy is indeed a farsighted policy in the proposition below.

Proposition 4. *For any state $(\boldsymbol{y}_t, \boldsymbol{a}_t, \boldsymbol{d}_t)$, the OOF policy solves the optimization problem in (4).*

When all devices in inventory and arrivals are used to satisfy demand, i.e., $y^t + a^t = d^t$ the OOF policy is the same as the myopic policy and, hence, is optimal. The OOF policy could also be effective when the rate of demand is about the same as the arrival rate, i.e., if $d^t \approx a^t$. This will guarantee that the device inventory will not grow older, which occurs when the YOF policy is used.

Algorithm 2 Oldest-Out-First Policy

for all t **do**

1) Sort customer requests by customer warranty termination date (ignore the age of the device that the customer is returning)

2) Sort devices in inventory by OEM warranty termination date

3) Match the oldest device in inventory (with lowest OEM warranty termination date) to the oldest customer (with lowest customer warranty termination date).

4) Repeat Step 3 until all customer requests are fulfilled.

end for

In fact, consider the same example discussed for the YOF policy. Assume for some interval $[t_1, t_2]$ that $b^t \geq 0, \forall t \in [t_1, t_2]$ (there are no side sales) and $a^t = d^t$, i.e., the total number of arrivals into inventory in each period is the same as total demand. Then, under the OOF policy we have the inventory getting younger, in that $\boldsymbol{y}_{t_2} \preceq \boldsymbol{y}_{t_2-1} \preceq \ldots \preceq \boldsymbol{y}_{t_1}$ for any realization of demand and arrivals. Hence, the OOF policy will accumulate "young" devices in inventory and, will be better positioned to respond to a spike in total demand that will dip into the inventory. This is in complete contrast to how the YOF policy behaves.

The OOF policy might perform poorly when there is excessive carry-over inventory, i.e., $d^t << y^t + a^t, \forall t$. In this case the OOF policy will be unable to clear old devices in inventory. For example, consider a set-up where there is a large device inventory, ranging in age from young to old, and only one demand request arrives per time-period. There is also one device arrival each period, and the arriving device is young. The optimal policy would have a cost of (near) zero as we can always assign a young device to the demand. However, the cost of the OOF policy could be very high as we always assign the oldest device; and effectively any young device needs to wait its turn until it becomes the oldest device and can then be assigned to the demand.

Ultimately, the performance of the OOF policy depends on the inventory management policy. The OOF policy performs well as long as the amount of carry-over inventory is moderate or small, in that it gets used and does not accumulate.

In our tests, the OOF policy has a reasonable performance and can perform better than the myopic policy, as displayed in Table 1. This is because, at WSP, the demand is non-stationary, so there is inventory build-up at the start of a device's life-cycle, and then this inventory is used when the demand for replacement devices peaks. In the following example, we present a sample path of demand where the OOF policy performs better than any myopic policy.

Example where OOF is better than any myopic policy. Consider a setting where incoming demand can have three termination dates $\{t_1, t_1 + \epsilon, t_2\}$, with $t_2 > t_1 + \epsilon >> t$, ϵ is small, lead time

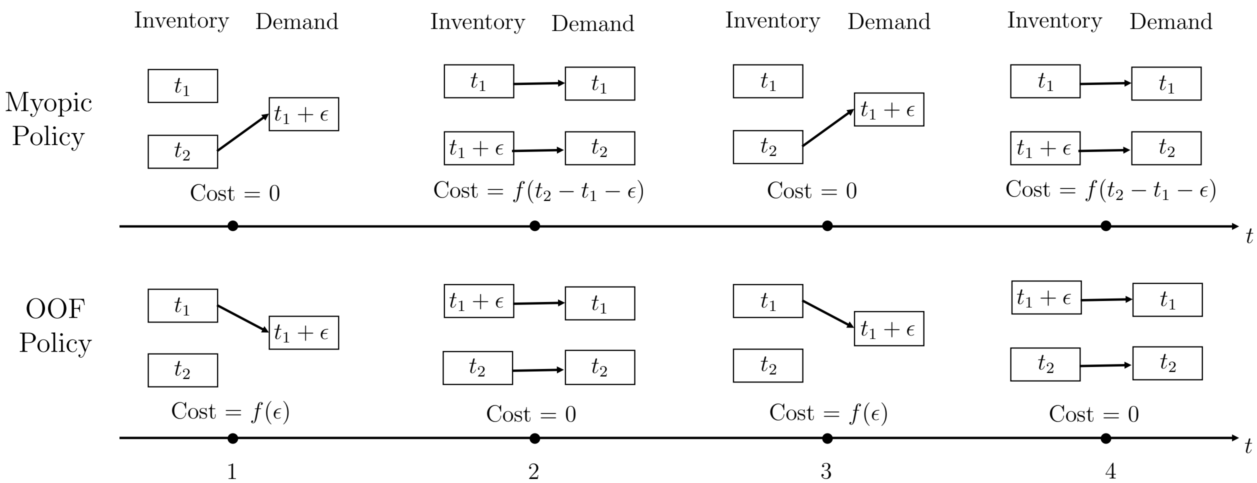

Figure 4 Example of a setting where the OOF policy is better than any myopic policy.

is 1, all devices are refurbished, and the OEM warranty of devices returned by customers has the same termination date as the customer warranty. Assume the following demand pattern: when t is odd, the demand is a single customer with termination date $t_1+\epsilon$, and when t is even, the demand is two customers with termination dates t_1 and t_2. Hence, starting at $t=2$, when t is even, there will be one device with OEM warranty termination date $t_1+\epsilon$ arriving into inventory, and when t is odd, there will be two devices with warranty termination dates t_1 and t_2 arriving into inventory. Also, assume that the initial inventory in period $t=1$ has two devices: one with termination date t_1 and the other with termination date t_2. We illustrate this setting in Figure 4.

We now compare the costs of a myopic policy and the OOF policy for this demand path. In period $t=1$, a myopic policy will match the device with termination date t_2 to the demand, resulting in an assignment cost of 0 and an inventory in the start of period $t=2$ consisting of one device with termination date t_1 and another with termination date $t_1+\epsilon$. Then, in $t=2$, both devices in inventory are used and the assignment cost will be $\beta\cdot(t_2-t_1-\epsilon)$. In period $t=3$ the demand and inventory after arrivals are the same as period $t=1$. Hence, for a horizon T where $T<t_1$ and T is even, the cost of the system will be $T\beta\cdot(t_2-t_1-\epsilon)/2$.

Conversely, in period $t=1$, the OOF policy will match the device with termination date t_1 to the demand, resulting in an assignment cost of $\beta\epsilon$, and an inventory in the start of period $t=2$ consisting of one device with termination date t_2 and another with termination date $t_1+\epsilon$. Then, in period $t=2$, both devices are used and, since the termination dates of the incoming demand is t_2 and t_1, the assignment cost is 0. In period $t=3$ the demand and inventory after arrivals is the same as in $t=1$. In this case, for a horizon of T where $T<t_1$ and T is even, the cost of the system will be $T\beta\cdot\epsilon/2$. The ratio between the total costs of any myopic policy and the OOF policy is $(t_2-t_1-\epsilon)/\epsilon$. Since the difference t_2-t_1 can be made arbitrarily large and ϵ arbitrarily small, the OOF policy can perform arbitrarily better than any myopic policy.

This example highlights a key tension of the warranty matching problem: a low-cost assignment in the current period can lead to high-cost states of the system in the future. In fact, completely ignoring future inventory position might lead to an arbitrarily poor system performance. We now introduce a heuristic policy that attempts to balance between short-term and long-term costs.

4.4. Heuristic Policies: The Sampling Policy

We introduce a heuristic policy that is simple and incorporates elements from both the YOF and OOF policy. This heuristic policy, which we denote the *Sampling Policy*, involves sampling, sorting, and then matching devices in inventory to demand.

For both the YOF and the OOF policy, in any period the inventory usually has a very different termination date distribution than the demand: under the YOF policy the inventory will tend to be older than the incoming demand (they are "old" devices), while under the OOF policy it will tend to be younger. Conceivably a better policy would manage the inventory so that its termination date distribution was matched to that of the demand. More specifically, in period t, one would like to achieve a low assignment cost in the current period while ensuring that the termination date distribution of devices in $\boldsymbol{y}_{t+1}$ is similar to the distribution of termination dates in $\boldsymbol{d}_{t+1}$.

Maintaining similarity between the distribution of $\boldsymbol{y}_t$ and $\boldsymbol{d}_t$ in every period t is the goal of the *Sampling Policy*. This policy samples without replacement d^t devices from the $y^t + a^t$ devices in inventory in period t, and then assigns these devices to demand to minimize the costs in the current period; as previously shown this can be done by the YOF policy, which is equivalent to the OOF policy in this case. Then, since the termination date distribution of demand changes fairly slowly, the devices carried over (if any) should have a similar distribution to future demand. The Sampling Policy is described formally in Algorithm 3. Furthermore, as displayed in Table 1, the sampling policy leads to lower uncovered costs than either the YOF or OOF policy.

Algorithm 3 The Sampling Policy

for all t **do**

1) Sort customer requests by customer warranty termination date (ignore the age of the device that the customer is returning)

2) Sample without replacement d^t devices from inventory and sort them by OEM warranty termination date

3) Assign the sampled devices to demand following the YOF policy (which, in this case, is the same as the OOF policy).

end for

By randomly selecting devices in inventory to satisfy demand, the distribution of warranty termination dates of *not selected* devices that are carried over to the next period should not be too skewed and should be similar to the termination dates of incoming demand (as long as this distribution does not change too rapidly). This feature of the sampling policy contrasts with the YOF policy, where carried over devices are "old" relative to the incoming demand.

Furthermore, if total demand for replacements is high enough, the devices that are *selected* by the sampling policy have roughly the same distribution of termination dates as the distribution of termination dates of incoming demand. Once sampled devices are sorted and matched, this leads to a low mismatch cost. This policy feature contrasts with the OOF policy, where selected devices are "old" (so that carried over devices are "young").

Thus, the sampling policy mitigates the downsides of the YOF and OOF policy and tries to maintain a balanced distribution of termination dates in the carried over inventory. The sampling policy will not be near optimal if there is very low incoming demand in some period (for example, if there is demand for only one replacement device). In such case, the sampling policy will be similar to the random assignment policy.

Finally, we note that since WSP holds inventory in a centralized reverse logistics facility that processes all warranty claims and device returns, the YOF, OOF, and sampling policies are practical. The YOF and OOF policies could be implemented as long as devices can be easily sorted by termination date. We believe that the sampling policy is potentially the easiest policy to implement, as it only requires the sampled devices to be sorted. This set of sampled devices can be obtained by any picking strategy that picks phones independent of their age (e.g., FIFO).

5. Numerical Experiments

We now compare and contrast the assignment policies introduced in the previous section through numerical experiments and we investigate the sensitivity of the assignment policies to changes in seed-stock. In addition, this section concludes by simulating the performance of these policies with real data from our partner WSP. For the simulations in this section, we built a discrete-time simulator using the Julia programming language (Lubin and Dunning 2015). This simulator allows the warranty claims, devices and their termination date to be individually realized.

We will use the *average uncovered time per replacement device shipped* as a metric for measuring the performance of the different assignment policies. That is, we set the assignment cost $c_t(i,j) = (j - \max(i,t))^+$ for when a device with OEM-warranty termination date i is assigned to a customer with termination date j. Since a customer can experience multiple device failures, this is a measure of WSP's exposure to out-of-OEM-warranty returns.

We also compute a lower bound for the assignment problem where we assume that WSP knows all future warranty claims and their corresponding customer warranty termination dates. The lower bound is then the solution of the resulting large-scale assignment problem, and is a very conservative bound since it assumes that all future information is available to WSP.

5.1. Comparing the Different Assignment Policies

We assume a simulation horizon of two years and we let one discrete period correspond to one week. We also assume that when a customer purchases a new device, both the OEM and customer warranties start simultaneously, and that both the customer and the OEM warranty have a length of 52 weeks. In all our experiments, we replicate the simulation of the entire life-cycle 100 times.

We assume that 100,000 devices are sold over a product's life-cycle, and that all sales occur during the first 32 weeks after a product's launch date and decrease linearly over the sales horizon. Furthermore, we assume that prices of new and refurbished devices depreciate exponentially at a rate of 50% per year[2]. We assume that the initial wholesale price of devices that WSP sources is $100 per device. We also assume the WSP sells a refurbished device for an initial price of $75 per device. For each time period, the number of refurbished devices that are sold through side sales channels instead of being used to meet warranty demand, is based on the certainty-equivalent approximation in Calmon and Graves (2017). This approximation calculates the selldown level in each time period by assuming that the demand and arrival processes of devices are equal to their average values.

When a failed device is sent to the OEM for repair/refurbishment, the device keeps aging and consuming OEM warranty. Upon completion of the refurbishment process, the device is returned to the inventory at WSP's reverse logistics facility. In each time period, refurbished devices can be used to meet warranty demand or carried over to the next period or sold. Any device failures that occur beyond a customer's warranty period are not replaced. In addition, new devices can be purchased from the OEM in case of shortages.

We assume a three-week deterministic lead time from when a device fails until it completes repair and is returned to inventory, and a 20% probability that a device sent to the OEM is not repaired and leaves the system. We assume that in each period the reverse logistics facility receives into its inventory a seed-stock of new devices, equal to 5% of the devices sold during that period. We also assume that the failure distribution of new devices follows an exponential distribution with a mean failure time of 48 months. This leads to a constant weekly hazard rate of 0.0051 for the devices and approximately 22% of new devices fail under warranty. Finally, to simplify our experiment, we only simulate the first device failure and measure its assignment cost[3]. In Appendix EC.6 we discuss the case where devices have an increasing failure rate.

[2] Two data sources for phone price depreciation are (i) the depreciation of used phones (`https://www.bankmycell.com/blog/phone-depreciation/`) and (ii) CamelCamelCamel, a website that tracks Amazon prices (`https://camelcamelcamel.com`)

[3] We have found that incorporating the second and subsequent failures of a device will have insignificant impact on the general findings from these experiments.

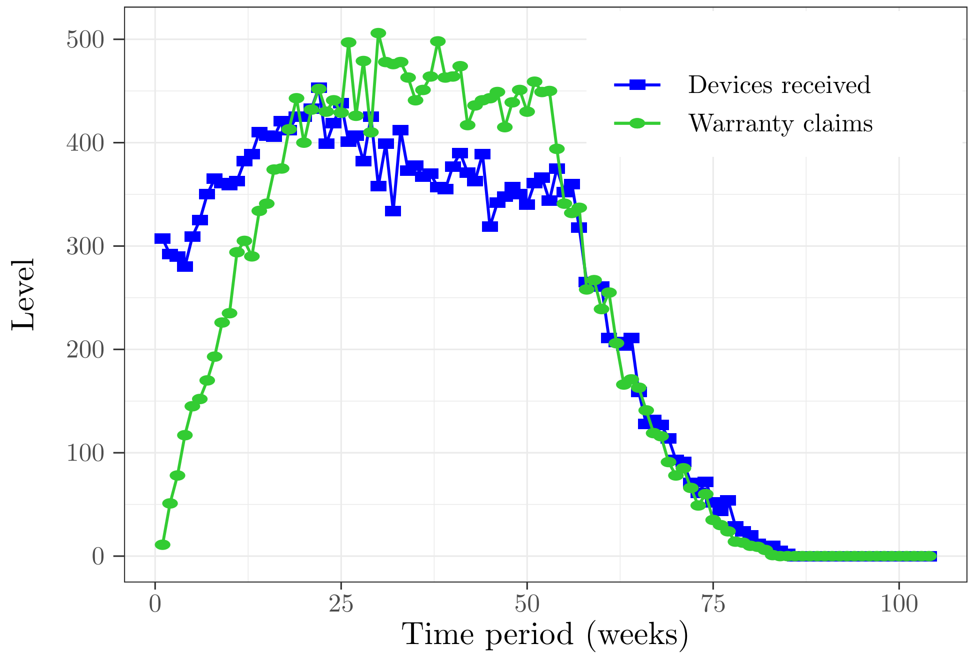

Figure 5 **Average number of devices received into inventory and warranty claims per period for 100 replications of the life-cycle.**

The number of devices received into inventory, the inventory level, and the number of warranty claims are common for all assignment policies as they only depend on new device sales and their failure times[4]. For each time period, Figure 5 depicts a sample path of the number of devices received into inventory and the number of warranty requests. Observe that the largest inventory build-up and depletion occurs during the first part of the simulation horizon, which is representative for many devices of our partner WSP. We plot the correspoding inventory level for this sample path in Figure 6a and display the number of new devices needed to cover shortages in Figure 6b.

For the time window $t \in [10, 32]$, Figure 7 depicts the average age of the inventory for each policy for each week after the assignments, and the average age of the incoming warranty claims, which is the same for all policies. We observe that a set of old devices gets "stuck" in inventory under the YOF policy, while the OOF policy results in the youngest inventory over this interval. For the sampling and random policies, we observe that the average age of the inventory is fairly close to the average age of the incoming warranty claims.

The age distribution of the inventory dictates the performance of the assignment policies. Initially, relative to the OOF policy, the YOF policy's average uncovered time is lowest. However, when inventory becomes scarce, the OOF policy outperforms the YOF policy due to having younger

[4] This condition would not hold if we included second and subsequent failures of a device.

Electronic copy available at: https://ssrn.com/abstract=3357683

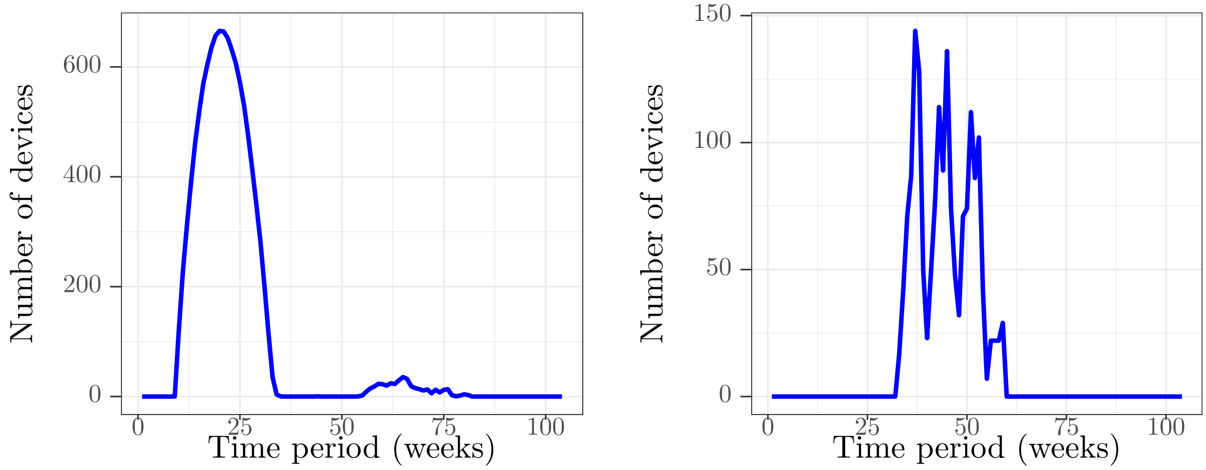

(a) Inventory level

(b) New devices used

Figure 6 **Inventory level and new devices used.**

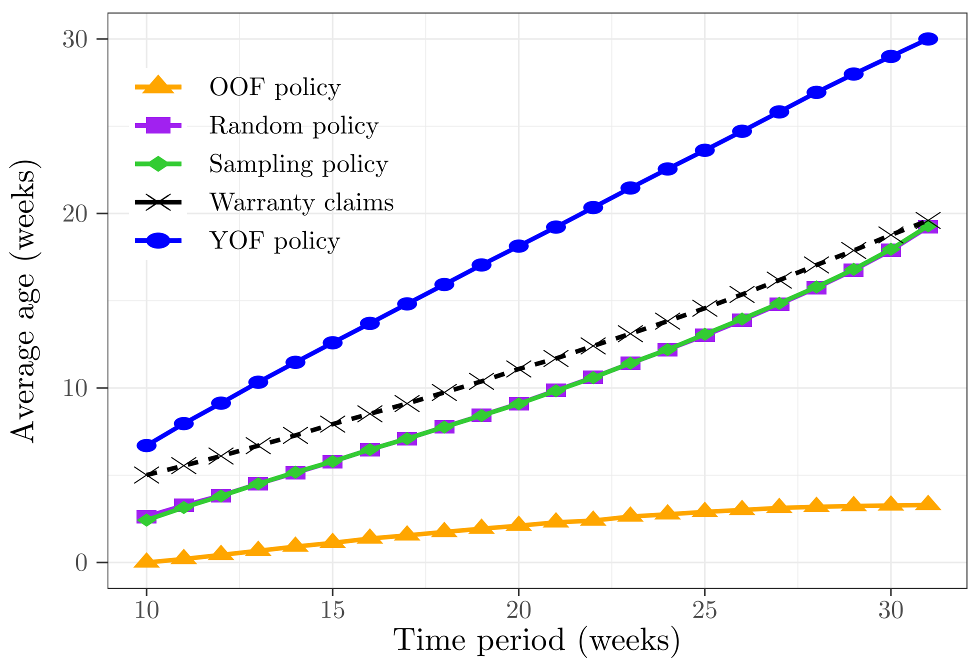

Figure 7 **Average age of the inventory and warranty claims for time window $t \in [10, 32]$.**

devices in inventory. When there is minimal inventory (e.g., after week 32), the YOF, OOF and sampling policy all perform very similarly[5].

[5] We note that there is minimal inventory towards the end of the device's life-cycle because, at this stage, the demand for replacement devices is decreasing over time (as observed in Figure 5), such that there are more incoming refurbished

This is the general pattern given an inventory build-up as shown in Figure 6a. The inventory build-up depends on the imbalance between received devices into inventory and warranty claims, both of which depend on the sales pattern for new devices. On the one hand, a larger inventory peak may affect the YOF policy's performance in a negative way as it may generate a larger number of old devices that eventually need to be used to meet warranty demand; on the other hand, the OOF policy's performance might be affected by a set of old devices that cannot be immediately cleared.

Over 100 replications, the average uncovered time per device is 2.89 weeks for the random assignment policy, 0.41 weeks for the YOF policy, 0.31 weeks for the OOF policy, and 0.22 weeks for the sampling policy. The OOF and sampling policies seem to outperform the YOF policy by avoiding the accumulation of old devices in inventory during the inventory build-up phase at the start of the life-cycle, as shown in Figure 7.

5.2. Sensitivity analysis

For some OEMs, the contract terms stipulate a seed-stock of new devices that the OEM provides to WSP to be used as replacement devices for the initial warranty claims. We show that variations of the seed-stock percentage affect the size and length of the inventory build-up and thus the assignment policies' performance.

Figure 8 shows the impact of changing the seed-stock percentage from 5% of sales to 3% or to 7% of sales. For the decrease to 3%, we find that more devices are being kept to meet warranty requests during the first periods of the life-cycle and, similarly, we observe more devices being sold in case of an increase to 7%.

Table 2 shows the sensitivity to a change in the seed-stock percentage for 100 replications. On average, the YOF policy outperforms the OOF policy for a seed-stock percentage of 3% of sales, whereas the OOF policy outperforms the YOF policy for a seed-stock percentage of 5% and 7% of sales. For all cases, the sampling policy does best[6].

The size and length of the inventory build-up plays an important role on the performance of the assignment policies. When comparing the 3% and 7% cases under the YOF policy, we find that, on average, older devices accumulate in inventory and eventually are used to meet warranty demand at a high assignment cost. This especially harms the performance of the YOF policy in the latter case. Conversely, higher seed-stock improves the performance of the OOF policy, since with higher seed-stock, younger devices accumulate in inventory.

devices from the OEM than demand for replacement devices in every period, and the inventory management policy sells the excess incoming devices.

[6] We conducted a Kolgomorov-Smirnov hypothesis test to verify whether the samples are drawn from the same distribution. The outcome of each pairwise test rejects the hypothesis with a p-value < 0.001.

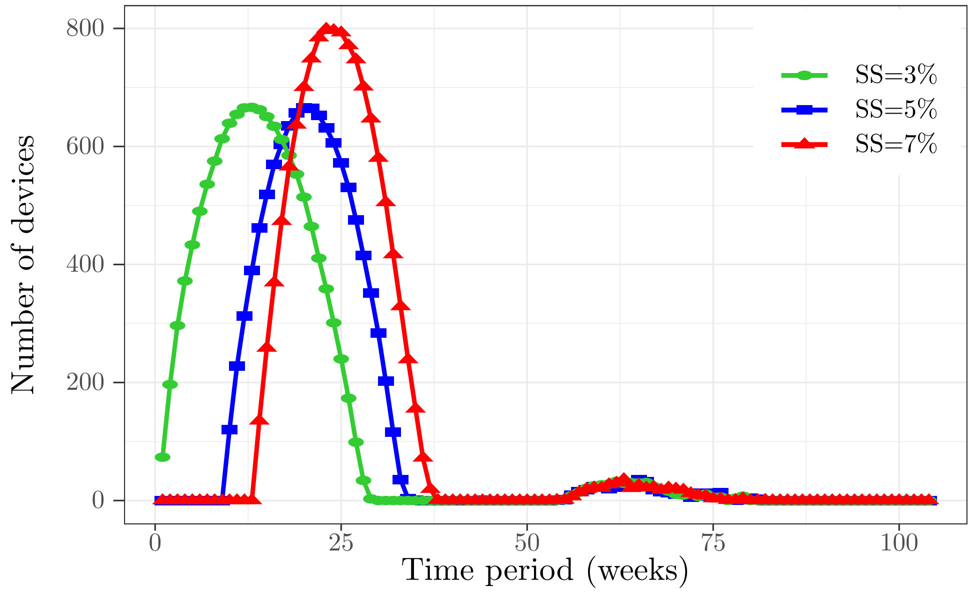

Figure 8 **Inventory level for different seed-stock rates as a percentage of sales.**

Assignment policy	3% seed-stock	5% seed-stock	7% seed-stock
Random policy	2.70 (0.04)	2.89 (0.04)	2.88 (0.04)
YOF policy	0.36 (0.02)	0.41 (0.03)	0.47 (0.03)
OOF policy	0.40 (0.02)	0.31 (0.02)	0.29 (0.02)
Sampling	0.22 (0.02)	0.22 (0.02)	0.23 (0.02)
Lower bound	0.16 (0.02)	0.16 (0.02)	0.16 (0.02)

Table 2 **Average and standard deviation (between brackets) of the uncovered time (in weeks) of each assignment policy for a 2% change in seed-stock**

5.3. Simulation based on data from WSP

We now compare the performance of the assignment policies when applied to sales and failure data for five of WSP's best-selling models, which we refer to as model A to model E. For every device of each model that is sold we have the device's sales date and failure date (if the device failed at all). We depict the number of sold and failed devices for each model in Figure EC.3 in Appendix EC.9, and the hazard rates of three models in Figure EC.2 in Appendix EC.7.

In practice, over 10% of the devices sold for each model failed under warranty. For these tests, we assume an 80% probability that a failed device is refurbished and that seed-stock corresponds to 1% of sales. We set the OEM lead time to three weeks. We assume that the OEM warranty for each device starts on its sales date, as does the customer warranty; and we assume that both warranties are for twelve months.

Similar to our previous simulations, we assume that prices of new and refurbished devices depreciate at a rate of 50% per year over the entire simulation horizon and we use the certainty-equivalent approximation to calculate the selldown levels (Calmon and Graves 2017). We assume

Electronic copy available at: https://ssrn.com/abstract=3357683

Model	Random policy	YOF policy	OOF policy	Sampling policy
A	100%	25%	22%	15%
B	100%	23%	19%	13%
C	100%	17%	15%	10%
D	100%	21%	24%	17%
E	100%	23%	28%	17%

Table 3 Expected number of out-of-OEM-warranty repairs as a percent of the out-of-OEM-warranty repairs under the random assignment policy

a daily time period so that there are failures and matching on a daily basis. The simulation time horizon for each model is 972 days.

In Table 1 we display the average uncovered time for each assignment policy for each model. Since we have data on which device was sent to which customer, we can estimate the average uncovered time for WSP's current policy [7]. In Table 3 we report the expected number of out-of-OEM-warranty repairs for each assignment policy for each model as a percent of the out-of-OEM-warranty repairs of the random assignment policy (we omit the absolute numbers to preserve the confidentiality of our industrial partner's data). To estimate the expected number of out-of-OEM-warranty repairs, we first calculate each device's empirical failure age distribution using WSP's data. We assume that each replacement device has age zero when sent to a customer. Then, we use the empirical failure age distribution to compute the probability that the device will fail after the device's OEM warranty ends but before the end of the customer's warranty. The expected number of out-of-OEM-warranty failures is the sum of these probabilities across all device-customer assignments. When comparing our assignment policies to the random policy, we find an average percentage reduction of the average number of out-of-OEM-warranty repairs between 78% to 86%.

For all device models, the best performing policy is the sampling policy. As discussed in Section 4, the sampling policy allows for the distribution of warranty termination dates of the devices carried over to be similar to the distribution of customer warranty termination dates of incoming demand. Furthermore, since in WSP's setting there is a large volume of warranty claims per day, the sampling policy will have a low assignment cost in every period. If there were very few warranty claims per day, the performance of the sampling policy would degrade. We examine the sensitivity of the sampling policy to the number of devices sold by WSP in Appendix EC.8.

In Figure 9 we show the inventory levels for models A and E. Similar to our numerical analysis in the previous subsection, we observe a single inventory peak for model A. The inventory build-up leads to a set of old devices getting stuck in inventory under the YOF policy; the YOF policy then

[7] The actual uncovered time is likely longer due to two reasons: the OEM warranty may start before the device was sold to the customer, as we assume in the simulation; and some customers will buy an extended warranty that goes beyond the assumed one year.

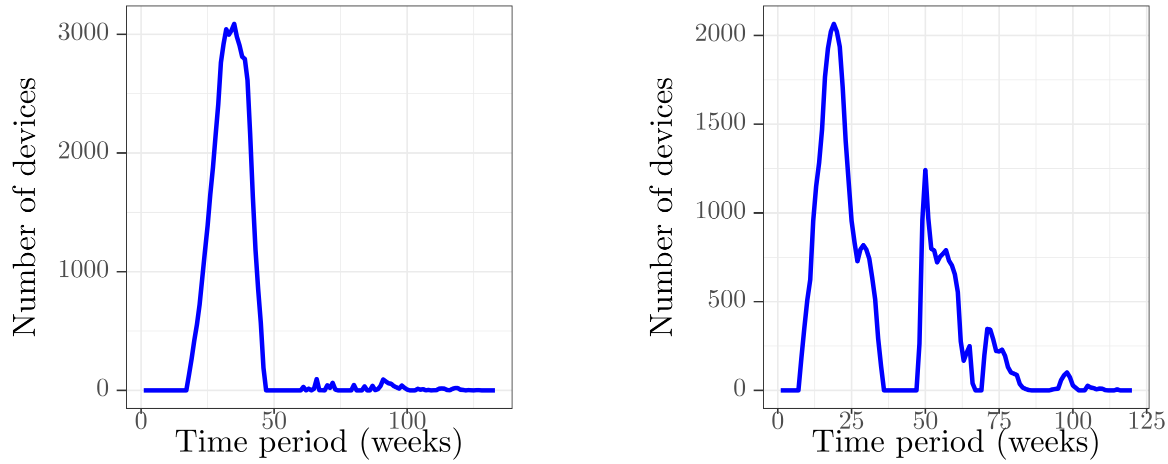

(a) Model A (b) Model E

Figure 9 **Inventory level for device models A and E.**

must assign these old devices during the following inventory depletion period. This results in a poor performance of the YOF policy compared to the OOF policy. Compared to model A, model E experiences three shorter inventory peaks which results in a better performance for the YOF policy. However, for model E, the inventory accumulations that occur later in the time horizon hurt the performance of the OOF policy, as this policy must increasingly assign older devices to warranty claims that are declining over time.

5.4. Economic cost of different policies

In this section, we estimate and compare the *economic cost* of the different policies. Although our main insights do not change, this cost analysis allows us to: (i) take into account the potential decrease of refurbishment costs during the life-cycle of the device, and (ii) estimate potential savings for WSP. We note that, at any given moment, WSP is managing a customer base of over 150 million customers[8].

We assume that the cost of repairing an out-of-OEM-warranty return is one third of the wholesale price of a new device at the time the warranty claim is filed; that is, at time t the refurbishment cost $r_t = p_t/3$, where p_t is the wholesale price of a new device. We expect that this assumption is conservative; for instance, Apple charges consumers over 50% the MSRP of devices for out-of-warranty repairs[9]. Then, when a device with OEM warranty termination date i is assigned to a

[8] Data on the largest American WSPs is publicly available: `https://en.wikipedia.org/wiki/List_of_United_States_wireless_communications_service_providers`

[9] `https://support.apple.com/iphone/repair/service/pricing`

Device	Random Policy	YOF Policy	OOF Policy	Sampling Policy
A	0.387	0.094	0.077	0.050
B	0.516	0.117	0.095	0.061
C	0.324	0.066	0.061	0.041
D	0.423	0.084	0.101	0.066
E	0.214	0.047	0.058	0.034
Average	0.374	0.082	0.079	0.050

Table 4 Expected out-of-OEM repair costs weighted by the total number of devices sold of each model for different assignment policies. The unit of the values in the table is $/device sold.

customer with customer warranty termination date j, where $j > i$, the expected cost is

$$\text{Expected cost} = \sum_{k=\max(i,t)}^{j} r_k \cdot \Pr(\text{device fails at time } k)$$

where Pr(device fails at time k) is estimated from the device's empirical failure age distribution. Based on these assumptions, we use the warranty data to estimate the total expected out-of-OEM-warranty repair cost for each device. We assume that a new phone costs $750 and that WSP sources a new phone for a wholesale price of about $650 for a 15% gross margin. Therefore, the initial out-of-OEM-warranty repair cost is $217 per device and we assume that this cost depreciates at a rate of 50% per year[10].

We report the expected out-of-OEM-warranty repair cost per device sold for each model in Table 4. By weighting by the number of devices sold for each model we obtain the average expected cost per device sold for each policy; for the random assignment policy, the average is $0.374 per device sold and for the sampling policy it is $0.050 per device sold, a cost reduction of 87%.

If we assume that consumers change devices once every three years on average, then WSP sells about 50 million devices per year. If we then use the costs from Table 4, we can estimate that switching from random assignment to the sampling policy would lead to savings of about $16 million per year (namely 50 million $\cdot$ ($0.374 − $0.050)).

6. Conclusion

In this paper, we model and analyze the problem of assigning devices to customers in a closed-loop system when there are two warranties in place: (a) the customer warranty offered by the retailer to the customer, and the (b) OEM warranty, provided by the OEM to the retailer. Ideally the two warranties would be matched, i.e., the customer would have the same time left in his customer

[10] This estimate is conservative. Data from Apple (`https://support.apple.com/iphone/repair/service/pricing`) indicates that the repair costs of phones depreciates at a rate lower than 50% per year. Furthermore, the $250 initial cost for repairs would imply that a company such as Apple has gross margin from servicing of over 100% for most devices, which is likely too high.

warranty as the device would have left in the OEM warranty. A mismatch between these warranties incurs costs to the retailer: when a customer still covered by the customer warranty has a device that fails, and this device is not covered by the OEM warranty.

We formulated the problem as a discrete-time stochastic dynamic program that takes into account the closed-loop nature of the system, as well as the termination date of customers and devices. Given this setting, we analyzed different assignment strategies and how they impact the average uncovered time. In particular, we focused on four different policy classes that consider different trade-offs in this dynamic optimization problem.

First, we analyzed random assignment both as as a proxy for WSP's current practice and as a benchmark. We proved a distribution-free upper bound on the expected mismatch cost for this policy. This bound has a practical interpretation and can help a manager decide whether it is worth investing in an assignment policy other than random assignment.

Second, we considered the Youngest-Out-First policy, where customers and devices are sorted by age and matched from youngest to oldest. This policy is optimal if we run out of inventory in any period, and it performs well if the inventory build-up never needs to be used for replacements. However, this policy has a major drawback, namely the YOF policy may lead to an accumulation of old devices in the system. These devices might become out-of-OEM-warranty returns when sent to customers.

Third, we addressed this drawback through the Oldest-Out-First policy. This policy also sorts devices and customers by age but now matches them from oldest to youngest. This policy aims to avoid allowing devices to "age" in inventory. The OOF policy is also an optimal policy if we run out of inventory in any period and it will outperform the YOF policy, if we build inventory that we then use in the near term for replacements. However, the OOF policy will perform very poorly when we accumulate a large inventory over the device's life-cycle.

Finally, we introduced the sampling policy where devices in inventory are sampled, sorted, and then matched to demand. This policy blends elements of the YOF and the OOF policy.

These policies are practical, since WSP has a centralized reverse logistics facility for processing warranty claims and device returns. Also, if perfectly sorting all devices is challenging, WSP can perform a "coarse" sorting where sampled devices are not perfectly sorted and are, instead, sorted by termination week (or even month). This "coarse" sort would still present significant benefits, albeit with some degradation on the order of a few days.

Through numerical simulation experiments, we observe that the YOF, OOF and sampling policies significantly decrease the average number of uncovered weeks compared to random assignment. The relative performance of the OOF and YOF policies depends on the length and size of the inventory build-up.

These findings were confirmed using sales and failure dates from five of WSP's best-selling models. We observe that the YOF, OOF and sampling policies present vast improvements over random assignment due to the *power of sorting* requests and devices as the average uncovered time, as well as the expected number of out-of-OEM-warranty failures, is reduced by between 75% to 90%.

Our analysis assumes a fairly simple structure of the assignment cost. Examining matching policies that account for time or age-dependent failure rates could lead to more effective assignment policies.

We note that the side-sales policy is key to the assignment policies' performance as it puts a bound on how much the inventory can grow. Thus, a promising future research direction is to examine both inventory management and device assignment with an integrated model. Another research opportunity might be the examination of other dynamic optimization problems with a cost and state-space structure that leads to simple farsighted policies performing better than myopic policies. Finally, in other practical settings the refurbished inventory might be held in multiple regional facilities, which would add another dimension to the assignment problem.

References

Abbey JD, Guide VDR (2018) A typology of remanufacturing in closed-loop supply chains. *International Journal of Production Research* 56(12):374–384.

Albers S (2003) Online algorithms: a survey. *Mathematical Programming* 97(1-2):3–26.

Andersson T, Ehlers L, Martinello A (2018) Dynamic refugee matching. *Working Paper 2018:7, Lund University, Department of Economics School of Economics and Management* 29 pages.

Ashlagi I, Fischer F, Kash IA, Procaccia AD (2015) Mix and match: A strategyproof mechanism for multi-hospital kidney exchange. *Games and Economic Behavior* 91:284–296.

Ashlagi I, Gilchrist DS, Roth AE, Rees MA (2011) Nonsimultaneous chains and dominos in kidney-paired donation—revisited. *American Journal of Transplantation* 11(5):984–994.

Aziz H, Chen J, Gaspers S, Sun Z (2018) Stability and pareto optimality in refugee allocation matchings. *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, 964–972.

Bansal N, Buchbinder N, Gupta A, Naor JS (2007) An o (log2 k)-competitive algorithm for metric bipartite matching. *Algorithms–ESA 2007*, 522–533 (Springer).

Boxma OJ, David I, Perry D, Stadje W (2011) A new look at organ transplantation models and double matching queues. *Probability in the Engineering and Informational Sciences* 25(2):135–155.

Calmon AP (2015) *Reverse logistics for consumer electronics: Forecasting failures, managing inventory, and matching warranties*. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.

Calmon AP, Graves SC (2017) Inventory management in a consumer electronics closed-loop supply chain. *Manufacturing & Service Operations Management* 19(4):568—-585.

Cohen MA, Pierskalla WP (1975) Management policies for a regional blood bank. *Transfusion* 15(1):58–67.

Galbreth MR, Blackburn JD (2006) Optimal acquisition and sorting policies for remanufacturing. *Production and Operations Management* 15(3):384–392.

Gale D, Shapley LS (1962) College admissions and the stability of marriage. *The American Mathematical Monthly* 69(1):9–15.

Guide VDR, Van Wassenhove LN (2009) OR FORUM: The evolution of closed-loop supply chain research. *Operations Research* 57(1):10–18.

Gusfield D, Irving RW (1989) *The stable marriage problem: structure and algorithms*, volume 54 (MIT press Cambridge).

Hitsch GJ, Hortaçsu A, Ariely D (2010) Matching and sorting in online dating. *American Economic Review* 100(1):130–163.

Hu M, Zhou Y (2016) Dynamic type matching. *Rotman School of Management Working paper No. 2592622* 71 pages, URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2592622.

Huang W, Kulkarni V, Swaminathan JM (2008) Managing the inventory of an item with a replacement warranty. *Management Science* 54(8):1441—-1452.

Huang Z, Peng B, Tang ZG, Tao R, Wu X, Zhang Y (2018) Tight competitive ratios of classic matching algorithms in the fully online model. *Working paper* 17 pages, URL https://arxiv.org/pdf/1810.07903.pdf.

Jagannathan R, Sen TK (1991) Storing crossmatched blood: A perishable inventory model with prior allocation. *Management Science* 37(3):251–266.

Jennings JB (1973) Blood bank inventory control. *Management Science* 19(6):637–645.

Kennes J, Monte D, Tumennasan N (2014) The day care assignment: A dynamic matching problem. *American Economic Journal: Microeconomics* 6(4):362–406.

Korolko N, Woodard D, Yan C, Zhu H (2018) Dynamic pricing and matching in ride-hailing platforms. *Working paper* 38 pages, URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3258234.

Kuhn HW (1955) The hungarian method for the assignment problem. *Naval research logistics quarterly* 2(1-2):83–97.

Kurino M (2014) House allocation with overlapping agents: A dynamic mechanism design approach. *American Economic Journal: Microeconomics* 6(1):258–289.

Lubin M, Dunning I (2015) Computing in operations research using Julia. *INFORMS Journal on Computing* 27(2):238–248.

Manshadi VM, Gharan SO, Saberi A (2012) Online stochastic matching: Online actions based on offline statistics. *Mathematics of Operations Research* 37(4):559–573.

Meyerson A, Nanavati A, Poplawski L (2006) Randomized online algorithms for minimum metric bipartite matching. *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithms*, 954–959.

Özkan E, Ward AR (2018) Dynamic matching for real-time ridesharing. *Working paper* 73 pages, URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2844451.

Pegels CC, Jelmert AE (1970) An evaluation of blood-inventory policies: a markov chain application. *Operations Research* 18(6):1087–1098.

Pince C, Ferguson M, Toktay B (2016) Extracting maximum value from consumer returns: Allocating between remarketing and refurbishing for warranty claims. *Manufacturing & Service Operations Management* 18(4):475–492.

Roth AE (2003) The origins, history, and design of the resident match. *JAMA* 289(7):909–912.

Roth AE, Sotomayor MAO (1992) *Two-sided matching: A study in game-theoretic modeling and analysis*. Number 18 (Cambridge University Press).

Sarhangian V, Abouee-Mehrizi H, Baron O, Berman O (2018) Threshold-based allocation policies for inventory management of red blood cells. *Manufacturing & Service Operations Management* 20(2):347–362.

Savaskan RC, Bhattacharya S, Van Wassenhove LN (2004) Closed-loop supply chain models with product remanufacturing. *Management Science* 50(2):239–252.

Slaugh VW, Akan M, Kesten O, Ünver MU (2016) The Pennsylvania adoption exchange improves its matching process. *Interfaces* 46(2):133–153.

Souza GC (2013) Closed-loop supply chains: A critical review, and future research. *Decision Sciences* 44(1):7–38.

Truong VA, Wang X (2018) Online two-sided bipartite matching. *Working paper* 24 pages, URL http://www.columbia.edu/~xw2230/papers/TwoSided_public.pdf.

Ünver MU (2010) Dynamic kidney exchange. *The Review of Economic Studies* 77(1):372–414.

Electronic Companion: Proofs and Additional Figures

EC.1. Motivation for Cost Function Proportional to Uncovered Warranty Time

In Section 3, we specify the cost of assigning a device with OEM warranty termination date i to a customer with warranty termination date j at time t by $c_t(i,j) = \beta \cdot (j - \max(i,t))^+$. That is, we assume that the expected out-of-OEM-warranty failure cost is proportional to the uncovered warranty time $(j - \max(i,t))^+$. We now provide motivation for this approximation for exponential and Weibull failure age distributions. We will use the terminology *failure age* for the time of failure, as we are modeling the failure time as it depends on the age of the device. In this analysis, we assume that when a replacement device is sent to a customer at time t, the age of the device at time t is zero since the device is not in use while in WSP's inventory.

EC.1.1. Exponential Failure Age Distribution

If the failure age follows an Exponential distribution with rate λ, the probability of an out-of-OEM-warranty failure is the probability that the device fails at time τ with $\max(i,t) \le \tau \le j$. Thus, if $\max(i,t) \le j$,

$$\begin{aligned}\Pr(\max(i,t) \le \tau \le j) &= \int_{max(i,t)}^{j} \lambda e^{-\lambda(x-t)} dx \\ &= e^{\lambda t} \cdot \left(e^{-\lambda \max(i,t)} - e^{-\lambda j}\right).\end{aligned}$$

If the out-of-warranty cost is proportional to the failure probability, we have

$$c_t(i,j) = \beta e^{\lambda t} \left(e^{-\lambda \max(i,t)} - e^{-\lambda j}\right)^+.$$

The function $e^{-\lambda x} - e^{-\lambda y}$ is neither convex nor concave for any positive values of x and y and for any rate λ.

The first-order approximation of $e^{-\lambda x}$ around some point a is $e^{-\lambda a} - \lambda e^{-\lambda a} \cdot (x - a)$. Thus, the first-order approximation of the cost $c_t(i,j)$ around t is

$$c_t(i,j) \cong \beta \cdot \lambda \cdot e^{\lambda t} \cdot e^{-\lambda t} \cdot (j - \max(i,t))^+ = \beta \cdot \lambda \cdot (j - \max(i,t))^+.$$

Hence, the first-order approximation of the out-of-warranty cost is proportional to the uncovered warranty time.

EC.1.2. Weibull Failure Age Distribution

Similar to the previous case, if the failure age follows a Weibull Distribution with scale $1/\lambda$ and shape parameter k, the probability of an out-of-OEM-warranty failure is the probability that the failure time τ occurs between $\max(i,t)$ and j. Thus, if $\max(i,t) \le j$,

$$\begin{aligned}\Pr(\max(i,t) \le \tau \le j) &= \Pr(\tau \le j) - \Pr(\tau \le \max(i,t)) \\ &= e^{-(\lambda(\max(i,t)-t))^k} - e^{-(\lambda(j-t))^k}.\end{aligned}$$

If the out-of-warranty cost is proportional to the failure probability, we have

$$c_t(i,j) = \beta \cdot \left(e^{-(\lambda(\max(i,t)-t))^k} - e^{-(\lambda(j-t))^k}\right)^+ = \beta \cdot (g(\max(i,t)-t, j-t))^+,$$

where $g(x,y) = e^{-(\lambda x)^k} - e^{-(\lambda y)^k}$. The function $g(x,y)$ is not necessarily convex or concave, since the Hessian of $g(x,y)$ is a diagonal matrix with eigenvalues

$$\left(e^{-(\lambda x)^k} k\lambda^k x^{k-2}(1-k(1-(\lambda x)^k)), -e^{-(\lambda y)^k} k\lambda^k y^{k-2}(1-k(1-(\lambda y)^k))\right).$$

The first-order approximation of the exponential in $g(x,y)$ around some point $x=y=a>0$ yields

$$g(x,y) \approx k e^{-(\lambda a)^k} \lambda^k a^{k-1}(y-x).$$

If we use this approximation in (1) with $a=t$ we obtain a first-order approximation that is proportional to the uncovered warranty time.

EC.2. Proofs for Random Matching

Before proving Proposition 1, we first introduce an auxiliary result.

Lemma EC.1. *Let X and Y be two discrete random variables with bounded support $[1,m]$. Let $\boldsymbol{p}$ be the probability distribution of X and $\tilde{\boldsymbol{p}}$ be the distribution of Y. Also, let X and Y have a bounded Kolmogorov-Smirnov distance, i.e.,*

$$\left|\sum_{i=1}^k p_i - \sum_{i=1}^k \tilde{p}_i\right| \le \epsilon, \ \forall k = 1,\dots,m.$$

Then, for any positive convex increasing function $f:\mathbb{R}\to\mathbb{R}$, where $f(x)=0$ for $x\le 0$, we have

$$E[f(X-Y)] \le \frac{(1+\epsilon)^2}{4} f(m-1).$$

Furthermore, this bound is tight.

Proof. Finding the upper bound is equivalent to solving.

$$\begin{aligned}
\max_{\boldsymbol{p},\tilde{\boldsymbol{p}}} \quad & \sum_{j=1}^{m}\sum_{i=1}^{m} \tilde{p}_i \cdot p_j \cdot f(j-i) \\
\text{s.t.} \quad & \sum_{i=1}^{m} p_i = 1,\ \sum_{i=1}^{m} \tilde{p}_i = 1, \\
& \left| \sum_{i=1}^{k} p_i - \sum_{i=1}^{k} \tilde{p}_i \right| \le \epsilon,\ \forall k = 1,\ldots,m, \\
& \boldsymbol{p}, \tilde{\boldsymbol{p}} \ge 0.
\end{aligned} \tag{EC.1}$$

The optimization problem is not convex since the objective function is not concave. In order to reformulate this problem as a convex optimization problem, we first prove that there is an optimal solution, $(\boldsymbol{p}^*, \tilde{\boldsymbol{p}}^*)$ where $p_2^* = \ldots = p_{m-1}^* = 0$ and $\tilde{p}_2^* = \ldots = \tilde{p}_{m-1}^* = 0$. Consider two random variables (X, Y) with distributions $(\boldsymbol{p}, \tilde{\boldsymbol{p}})$ that are feasible for the optimization problem above and have $p_k > 0$ or $\tilde{p}_k > 0$ for some $k \in [2, m-1]$. Then, consider alternative random variables (X', Y') with distributions $(\boldsymbol{p}', \tilde{\boldsymbol{p}}')$ such that

$$X' = X \cdot \mathbb{1}_{\{X \neq k\}} + Z \cdot \mathbb{1}_{\{X = k\}}, \qquad Y' = Y \cdot \mathbb{1}_{\{Y \neq k\}} + Z' \cdot \mathbb{1}_{\{Y = k\}}$$

where Z and Z' have the same distribution and $\mathbb{1}_{\{\cdot\}}$ is the indicator function. Let the distribution of Z be

$$Z = \begin{cases} 1, \text{ with prob. } \frac{m-k}{m-1} \\ m, \text{ with prob. } \frac{k-1}{m-1}, \end{cases}$$

and $E[Z] = k$. The distributions of X' and Y' are

$$p'_j = \begin{cases} p_j, \text{ if } X \notin \{1,k,m\}, \\ p_1 + p_k \frac{m-k}{m-1}, \text{ if } j = 1, \\ p_m + p_k \frac{k-1}{m-1}, \text{ if } j = m, \\ 0, \text{if } j = k, \end{cases} \qquad \text{and} \qquad \tilde{p}'_j = \begin{cases} \tilde{p}_j, \text{ if } Y \notin \{1,k,m\}, \\ \tilde{p}_1 + \tilde{p}_k \frac{m-k}{m-1}, \text{ if } j = 1, \\ \tilde{p}_m + \tilde{p}_k \frac{k-1}{m-1}, \text{ if } j = m, \\ 0, \text{if } j = k. \end{cases}$$

Note that $E[X'] = E[X]$ and $E[Y'] = E[Y]$, but $p_k = \tilde{p}_k = 0$. Also, since we are "splitting" the weights p_k and $\tilde{p}_k$ in the same way, all the constraints in (EC.1) are still satisfied. Since the function $(x-y)^+$ is jointly convex in (x, y) we can use Jensen's Inequality to obtain

$$\begin{aligned}
E[f(X'-Y')] &= E_{X,Y}[E_{Z,Z'}[f(X \cdot \mathbb{1}_{\{X \neq k\}} + Z \cdot \mathbb{1}_{\{X=k\}} - Y \cdot \mathbb{1}_{\{Y \neq k\}} - Z' \cdot \mathbb{1}_{\{Y=k\}})|X, Y]] \\
&\ge E_{X,Y}[f(X \cdot \mathbb{1}_{\{X \neq k\}} + E[Z] \cdot \mathbb{1}_{\{X=k\}} - Y \cdot \mathbb{1}_{\{Y \neq k\}} - E[Z'] \cdot \mathbb{1}_{\{Y=k\}})] \\
&= E_{X,Y}[f(X \cdot \mathbb{1}_{\{X \neq k\}} + k \cdot \mathbb{1}_{\{X=k\}} - Y \cdot \mathbb{1}_{\{Y \neq k\}} - k \cdot \mathbb{1}_{\{Y=k\}})] \\
&= E_{X,Y}[f(X-Y)].
\end{aligned}$$

Hence, (X', Y') attains a higher objective than (X, Y). We can repeat this procedure for any other positive component in the interval $[2, m-1]$ of either $\boldsymbol{p}$ or $\tilde{\boldsymbol{p}}$. Since this procedure applies for any discrete distribution with support $[1,m]$, we can rewrite the optimization problem only considering distributions that have positive weights on 1 and m. Thus, the optimization problem becomes

$$\begin{aligned} \max_{p_1,\tilde{p}_1} \quad & f(m-1)(1-p_1)\cdot\tilde{p}_1 \\ \text{s.t.} \quad & 0\le p_1\le 1,\ 0\le\tilde{p}_1\le 1, \\ & |p_1-\tilde{p}_1|\le\epsilon. \end{aligned} \tag{EC.2}$$

and the optimal solution is $p_m^*=\tilde{p}_1^*=1/2+\epsilon/2$, $p_1^*=\tilde{p}_m^*=1/2-\epsilon/2$ and $p_k^*=\tilde{p}_k^*=0, \forall k\in[2,\ldots,m-1]$. The optimal objective is $\frac{(1+\epsilon)^2}{4}f(m-1)$ and, by construction, the bound is tight. □

With the result above in hand, we now prove Proposition 1.

Proof of Proposition 1.

First, using the triangle inequality and (6), note that

$$\max_{k=1,\ldots m}|\Pr(D_t\le k)-\Pr(D_{t-l-z}\le k)|\le(l+z)\epsilon, \forall t. \tag{EC.3}$$

To keep our notation aligned with the previous proposition, we assume that the cost of assigning a device with OEM warranty termination date i to a customer with warranty termination date j at time t is $c_t(i,j)=f(j-\max(i,t))$, where f is a convex non-decreasing function and $f(x)=0$ for $x\le 0$ (the uncovered warranty time satisfies this assumption). Note that since $f(x)=0$ if $x\le 0$, we have

$$E[f(D_t-\max(D_{t-l-Z},t))]=E[f(\max(D_t,t)-\max(D_{t-l-Z},t))].$$

Thus, we can consider, WLOG, discrete distributions with support starting at t. Furthermore, since devices and customers have a warranty length of T_w, at time t, customers and devices will have latest warranty termination date of $t+T_w$ (corresponding to devices and customers whose warranty starts at t and ends at $t+T_w$). Thus, we can assume, WLOG, discrete distributions with support at most $[t,t+T_w]$.

Then, once again using the fact that f is a convex function and $f(x)=0$ for $x\le 0$, we use Lemma EC.1 to obtain

$$\begin{aligned} E[f(D_t-\max(D_{t-l-Z},t))] &= E_Z[E[f(D_t-\max(D_{t-l-Z},t))|Z]] \\ &\le E_Z\left[\frac{(1+(l+Z)\epsilon)^2}{4}f(T_w+t-t)\right] \\ &= \frac{E[(1+(l+Z)\epsilon)^2]}{4}f(T_w). \end{aligned}$$

Since $E[(l+Z)^2]=var(Z)+(E[l+Z])^2$ we have

$$\begin{aligned} E[(1+(l+Z)\epsilon)^2] &= 1+2(l+E[Z])\epsilon+E[(l+Z)^2]\epsilon^2 \\ &= 1+2(l+E[Z])\epsilon+var(Z)\epsilon^2+(l+E[Z])^2\epsilon^2 \\ &= (1+(l+E[Z])\epsilon)^2+var(Z)\epsilon^2. \end{aligned}$$

Recall that Z follows a geometric distribution with parameter $d/(y+d)$. Thus, $E[Z]=(y+d)/d$ and $var(Z)=y(y+d)/d^2$, which gives the constants in the statement of the proposition.

To show that the bound is tight, we use an argument similar to the end of the proof of Lemma EC.1. First, assume that D_t follows a two-point distribution (as in Lemma EC.1) with positive weights on t and $t+T_w$, and that, for any value of Z, D_{t-l-Z} follows the same distribution. Then for this two-point distribution, we have $E[f(D_t-\max(D_{t-l-Z},t))|Z]=\frac{(1+(l+Z)\epsilon)^2}{4}f(T_w)$. Since $f(T_w)$ does not depend on the realization of Z (it depends only on the warranty length), we can take the expectation over Z of both sides of the equality which yields the tight bound.

In order to obtain the expectation $E[U_t^R]$ we note that the bound above is independent of t. □

EC.3. Proofs for the YOF Policy

Proof of Proposition 2.

For the assignment problem given by (3), it is always optimal to use the d^t youngest devices in inventory (i.e., the d^t devices with highest termination date) to satisfy customer demand.

We now argue that it is optimal to sort the d^t youngest devices in inventory and the customer demand and match them from highest termination date to lowest termination date. We argue by contradiction. Assume that there is an optimal matching of devices to warranty claims that differs from the YOF matching and has a total cost that is strictly less than the one generated by the YOF policy. By assumption, there must be at least one "crossing", i.e., two devices in stock, with termination dates i_1 and i_2, that are each matched to warranty claims from customers with termination dates j_1 and j_2 , respectively, and $i_1 \geq i_2$ and $j_1 \leq j_2$. The term "crossing" is used here because if we arranged all the requests and devices in stock by termination date, and then looked at the network formed by matching devices in stock to requests, we would have crossing edges in any solution other than the one generated by the YOF policy.

To simplify notation, we define $c_t(i,j)=f(j-i)$, where f is a non-decreasing and convex function. Thus, since f is non-decreasing and convex in $j-i$, we can match i_1 with j_2 and i_2 with j_1, and potentially improve the cost, obtaining a contradiction. To show this, let $\Delta_j=j_2-j_1$. Then,

$$\begin{aligned}
&c_t(i_1,j_1)+c_t(i_2,j_2)-c_t(i_1,j_2)-c_t(i_2,j_1)=\\
&\quad=f(j_1-\max(i_1,t))+f(j_2-\max(i_2,t))-f(j_2-\max(i_1,t))-f(j_1-\max(i_2,t))\\
&\quad=f(j_1-\max(i_1,t))-f(j_1-\max(i_1,t)+\Delta_j)+f(j_1-\max(i_2,t)+\Delta_j)-f(j_1-\max(i_2,t))\\
&\quad\geq f'(j_1-\max(i_1,t))(-\Delta_j)+f'(j_1-\max(i_2,t))\Delta_j\\
&\quad=\big(f'(j_1-\max(i_2,t))-f'(j_1-\max(i_1,t))\big)\Delta_j\\
&\quad\geq 0.
\end{aligned}$$

The first inequality comes from the fact that f is convex (if f were not differentiable, a subgradient can be used), and the last inequality comes from the assumptions that $i_1 \geq i_2$ and that f is non-decreasing and convex, such that $f'(x) \geq f'(y)$ for any $x \geq y$.

We can repeat this procedure iteratively for all the "crossings" (and crossings that can be potentially added during the procedure), and eventually end in a solution where no crossings exist, which is exactly the one generated by the YOF policy, which is a contradiction. Since inventory is renewed in each period, the proof is complete. □

Proof of Proposition 3.

Since $\xi^t \geq d^t$, we upper-bound the mismatch cost by assuming that only incoming devices into inventory are used. If $\xi^t > d^t$, then we assume that d^t of the incoming devices are chosen at random to be sorted and assigned to the incoming demand. When matching the d^t arriving devices to the demand, the YOF policy is optimal and any myopic policy will do the same procedure. If carried over devices or seed-stock could be used as well (which is the case in any myopic policy including the YOF policy) the cost would be potentially lower.

For each period t, let $\{D_t\}_{i=1}^{d^t}$ be a set with the warranty termination dates of the d^t customers in the demand, and let $D_t^{(i)}$ denote the termination date of the i-th youngest customer in the demand. Also, let $\hat{F}_t$ be the empirical CDF of the warranty termination dates $\{D_t\}_{i=1}^{d^t}$. We note that, for any x, $\hat{F}_t(x)$ is a random variable and

$$\hat{F}_t(x) = \frac{1}{d^t} \sum_{i=1}^{d^t} \mathbb{1}_{\{D_t \leq x\}},$$

where $\mathbb{1}_{\{\cdot\}}$ is the indicator function. Let F_t be the CDF of the termination dates D_t in period t. Then, $\mathbb{1}_{\{D_t \leq x\}}$ is a Binomial random variable with parameter $F_t(x) = \Pr(D_t \leq x)$. Thus, $E[\hat{F}_t(x)] = F_t(x)$ and $var(\hat{F}_t(x)) = \frac{F_t(x)(1-F_t(x))}{d^t}$.

Furthermore, since we choose d^t of the arriving devices at random to be used as replacements, and since we assume that the warranty termination dates of customers in the incoming demand is the same as the warranty termination date of the devices that they return for repair, the termination dates of devices used as replacement at time t follows the same distribution as the termination dates of the demand in period $t-l$. Hence, denote by $D_{t-l}^{(i)}$ the warranty termination date of the i-th youngest device among the d^t devices sampled to be used as replacement at time t. Similarly, we denote by $\hat{F}_{t-l}$ the empirical CDF of the warranty termination dates of the d^t devices used as replacement, and note that $E[\hat{F}_{t-l}(x)] = F_{t-l}(x)$ and $var(\hat{F}_{t-l}(x)) = \frac{F_{t-l}(x)(1-F_{t-l}(x))}{d^t}$.

We have that

$$E[U_t^{MY}] \leq \frac{1}{d^t} \sum_{i=1}^{d^t} E\left[\left(D_t^{(i)} - \max(D_{t-l}^{(i)}, t)\right)^+\right], \forall t.$$

Electronic copy available at: https://ssrn.com/abstract=3357683

Since

$$\left(D_t^{(i)} - \max(D_{t-l}^{(i)}, t)\right)^+ = \left(\max(D_t^{(i)}, t) - \max(D_{t-l}^{(i)}, t)\right)^+,$$

and since the warranty has length T_w, we will assume that the support of D_t and D_{t-l} is $[t, t+T_w]$. Then, for any realization of $\{D_t^{(i)}\}$ and $\{D_{t-l}^{(i)}\}$ we have

$$\frac{1}{d^t}\sum_{i=1}^{d^t}\left(D_t^{(i)} - D_{t-l}^{(i)}\right)^+ = \int_t^{t+T_w}\left(\hat{F}_{t-l}(x) - \hat{F}_t(x)\right)^+ dx.$$

The equality can be deduced by noting that $\hat{F}_t$ only increases at the points $\{D_t^{(i)}\}$ and then comparing the integration areas.

Note that, for any x in the interval $[t, t+T_w]$,

$$\begin{aligned}\left(\hat{F}_{t-l}(x) - \hat{F}_t(x)\right)^+ &\le \left(\hat{F}_{t-l}(x) - F_{t-l}(x)\right)^+ + \left(F_t(x) - \hat{F}_t(x)\right)^+ + (F_{t-l}(x) - F_t(x))^+ \\ &\le \left|\hat{F}_{t-l}(x) - F_{t-l}(x)\right| + \left|F_t(x) - \hat{F}_t(x)\right| + |F_{t-l}(x) - F_t(x)|. \end{aligned} \tag{EC.4}$$

Using the assumption in (6) and since the random variables have bounded support, we take the expectation over $\hat{F}_t(x)$ and $\hat{F}_{t-l}(x)$ in (EC.4) and then integrate over the support $[t, T_w]$:

$$\begin{aligned}&\int_t^{t+T_w} E_{\hat{F}_{t-l}(x)}\left[\left|\hat{F}_{t-l}(x) - F_{t-l}(x)\right|\right] + E_{\hat{F}_t(x)}\left[\left|F_t(x) - \hat{F}_t(x)\right|\right] + |F_{t-l}(x) - F_t(x)|\, dx \\ &\le \int_t^{t+T_w} E_{\hat{F}_{t-l}(x)}\left[\left|\hat{F}_{t-l}(x) - F_{t-l}(x)\right|\right] + E_{\hat{F}_t(x)}\left[\left|F_t(x) - \hat{F}_t(x)\right|\right] dx + T_w l\epsilon \\ &\le \int_t^{t+T_w} \sqrt{E_{\hat{F}_{t-l}(x)}\left[\left(F_{t-l}(x) - \hat{F}_{t-l}(x)\right)^2\right]} + \sqrt{E_{\hat{F}_t(x)}\left[\left(F_t(x) - \hat{F}_t(x)\right)^2\right]} dx + T_w l\epsilon.\end{aligned}$$

The first inequality comes from (6). The second inequality is a direct application of Jensen's inequality since $\sqrt{x}$ is concave. The expectations in the last expressions are the variances of $\hat{F}_t(x)$ and $\hat{F}_{t-l}(x)$. Thus, from the discussion at the start of the proof,

$$E\left[\left(F_t(x) - \hat{F}_t(x)\right)^2\right] = \frac{F_t(x)\cdot(1 - F_t(x))}{d^t} \le \frac{1}{4d^t}.$$

Putting all the results together we obtain

$$E\left[\int_t^{t+T_w}\left|\hat{F}_{t-l}(x) - F_{t-l}(x)\right| + \left|F_t(x) - \hat{F}_t(x)\right| + l\epsilon dx\right] \le T_w\left(\frac{1}{\sqrt{d^t}} + l\epsilon\right).$$

Hence,

$$E[U_t^{MY}] \le T_w\left(\frac{1}{\sqrt{d^t}} + l\epsilon\right),$$

completing the proof. □

EC.4. Proofs for the OOF Policy

Proof of Proposition 4. For some time t, let $\boldsymbol{x}_t$ be the allocation obtained using the OOF policy and let $\boldsymbol{y_{t+1}}$ be the starting inventory in period $t+1$. Assume, for contradiction, that there is some other feasible allocation $\tilde{\boldsymbol{x}}_t$ with next period inventory $\tilde{\boldsymbol{y}}_{\boldsymbol{t+1}}$, that has a strictly lower objective cost for problem in Equation 4. However, since, by construction, in the OOF policy the d^t "oldest" devices are used, it must be that $\boldsymbol{y_{t+1}} \preceq \tilde{\boldsymbol{y}}_{\boldsymbol{t+1}}$ which implies $J_{t+1}(\boldsymbol{y}_{t+1}, \boldsymbol{a}_{t+1}, \boldsymbol{d}_{t+1}) \leq J_{t+1}(\tilde{\boldsymbol{y}}_{t+1}, \boldsymbol{a}_{t+1}, \boldsymbol{d}_{t+1}), \forall (\boldsymbol{a}_{t+1}, \boldsymbol{d}_{t+1})$, as stated in (5). Hence, we have a contradiction on the optimality of $\tilde{\boldsymbol{x}}_t$. □

EC.5. Hazard rate distributions

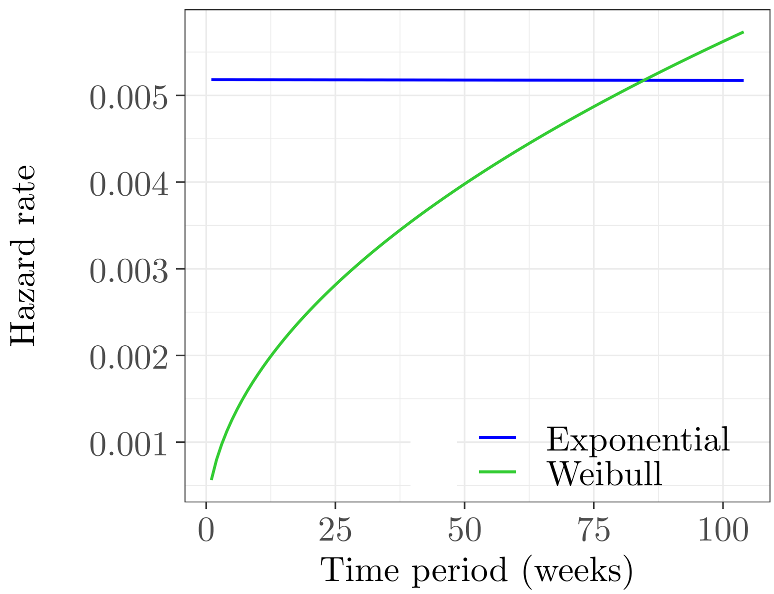

Figure EC.1 Hazard rate curves for an exponential failure distribution with average 192 weeks and for a Weibull distribution with shape and scale parameter equal to 192 and 1.5, respectively

.

EC.6. Results for an increasing hazard rate

Table EC.1 presents numerical results for the case where the failures follow a Weibull distribution and there is an increasing hazard rate like in Figure EC.1. All other parameters are equal to the simulation set-up reported in Section 5.1. In this case, devices will tend to fail later in their life-cycle which, in turn, delays the inventory peak compared to case when failures follow an exponential distribution. This hinders the performance of the YOF policy relative to the OOF and sampling policies since the average age of the devices which accumulate is higher than the case where failures follow the exponential distribution. The sampling policy's performance is again close to the lower bound.

Assignment policy	3% seed-stock	5% seed-stock	7% seed-stock
Random policy	2.56 (0.03)	2.68 (0.05)	2.67 (0.04)
YOF policy	0.56 (0.03)	0.65 (0.04)	0.68 (0.03)
OOF policy	0.39 (0.03)	0.38 (0.03)	0.36 (0.03)
Sampling	0.34 (0.03)	0.35 (0.03)	0.34 (0.03)
Lower bound	0.25 (0.02)	0.24 (0.03)	0.24 (0.02)

Table EC.1 Average and standard deviation (between brackets) of the uncovered time (in weeks) of each assignment policy for a 2% change in seed-stock. The failures follow a Weibull distribution with shape and scale parameter set to 192 and 1.5 respectively.

EC.7. Device hazard rate distributions for WSP

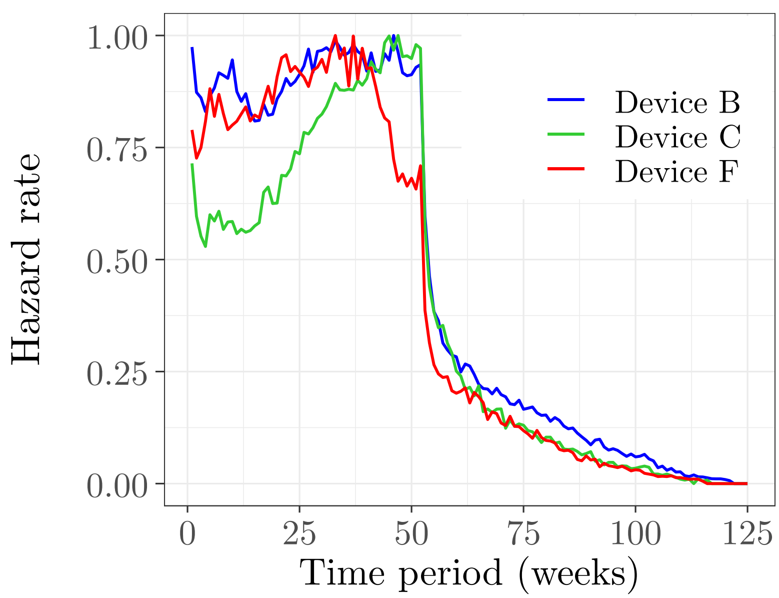

Figure EC.2 Three device hazard rate curves for WSP. The maximum of the curves was normalized to one to preserve data confidentiality.

EC.8. Sensitivity of the sampling policy to sales volume

We test the sensitivity of the sampling policy to the number of sales and depict the results in Table EC.2. The simulation is set up in a similar way as reported in Section 5.1, except that we change the number of sales (100,000 devices) as it will change the average total demand for warranty replacements in each period. We observe that the performance of the sampling policy deteriorates as the number of sales (on average) decreases. However, we find that the sampling policy outperforms the other policies for small but reasonable values for the number of sales. However, it does become challenging to identify the best performing policy in case of an extremely low number of sales.

Assignment Policy	25,000 devices sold (58 failures/period on avg.)	50,000 devices sold (116 failures/period on avg.)	200,000 devices sold (464 failures/period on avg.)
Random policy	2.99 (0.08)	2.91 (0.06)	2.83 (0.03)
YOF policy	0.56 (0.06)	0.45 (0.04)	0.37 (0.02)
OOF policy	0.48 (0.05)	0.37 (0.03)	0.28 (0.02)
Sampling	0.34 (0.04)	0.26 (0.02)	0.20 (0.02)

Table EC.2 Average and standard deviation (between brackets) of the uncovered time (in weeks) of each assignment policy for a change in the number of sales and the average number of failures per period

EC.9. Sales and failures for different device models sold by WSP

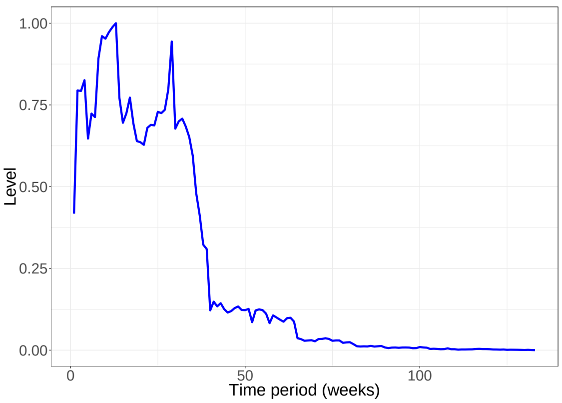

(a) Sales of model A

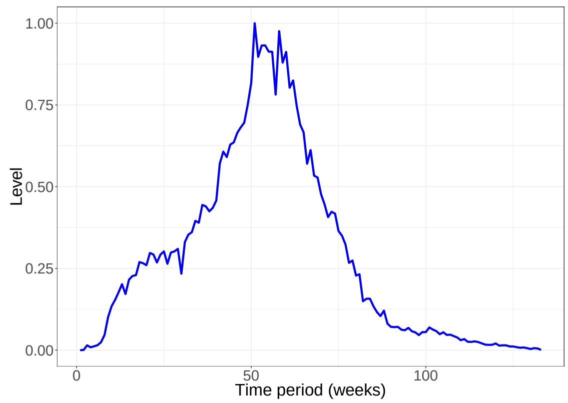

(b) Failures of model A

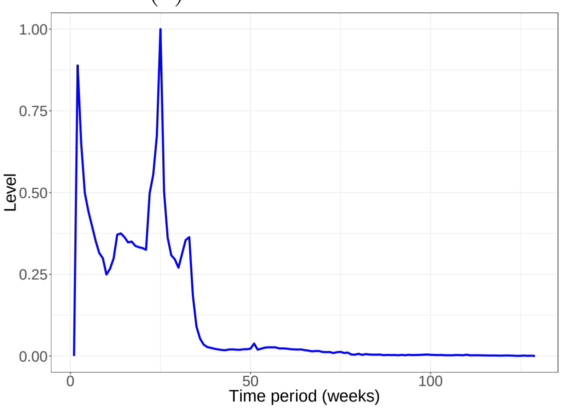

(c) Sales of model B

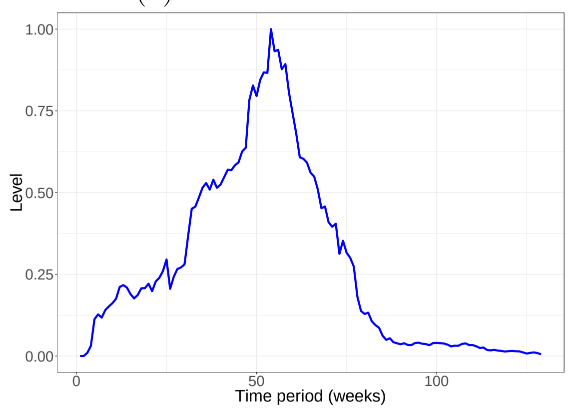

(d) Failures of model B

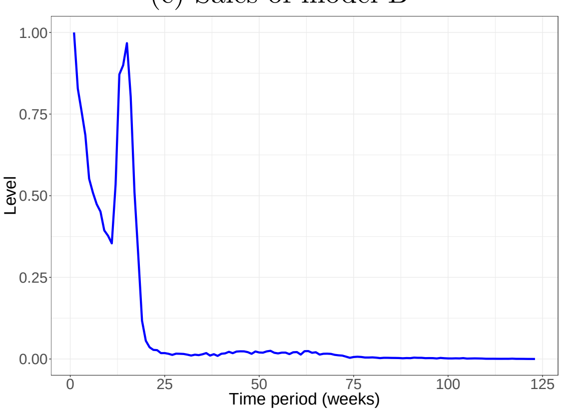

(e) Sales of model C

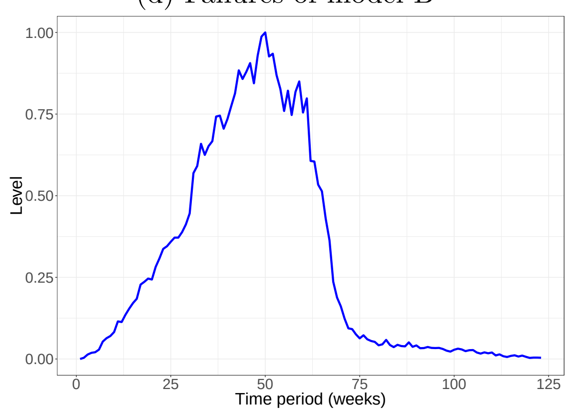

(f) Failures of model C

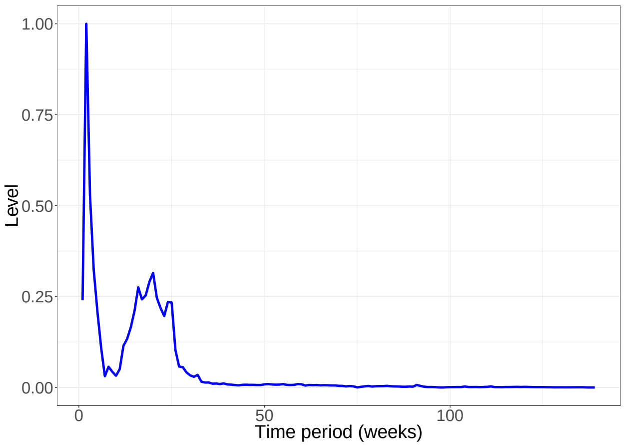

(g) Sales of model D

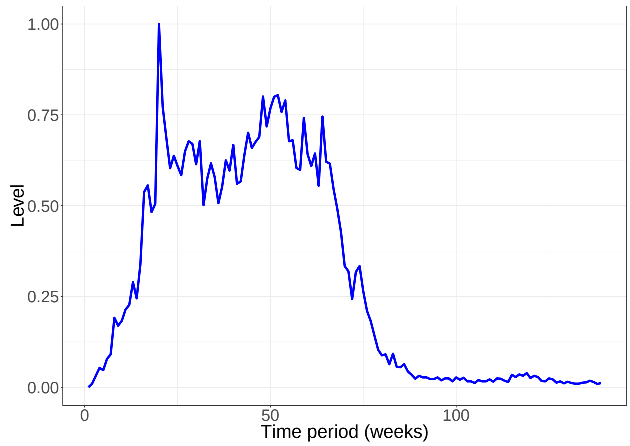

(h) Failures of model D

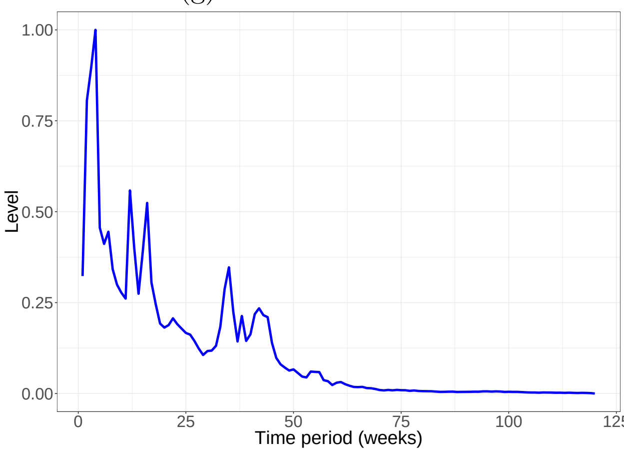

(i) Sales of model E

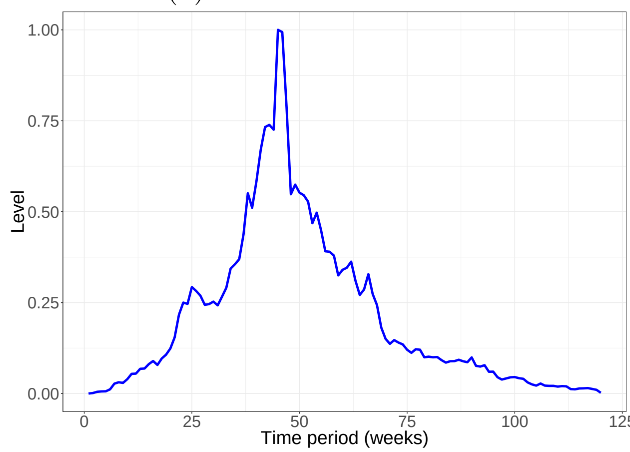

(j) Failures of model E

Figure EC.3 **Sales and failures for different device models. The maximum of the curves was normalized to one to preserve data confidentiality.**