



Delft University of Technology

Warranty Reserve Management Demand Learning and Funds Pooling

Wang, Xiao Lin; Zhong, Yuanguang; Li, Lishuai; Xie, Wei; Ye, Zhi Sheng

DOI

[10.1287/msom.2022.1086](https://doi.org/10.1287/msom.2022.1086)

Publication date

2022

Document Version

Final published version

Published in

Manufacturing and Service Operations Management

Citation (APA)

Wang, X. L., Zhong, Y., Li, L., Xie, W., & Ye, Z. S. (2022). Warranty Reserve Management: Demand Learning and Funds Pooling. *Manufacturing and Service Operations Management*, 24(4), 2221-2239. <https://doi.org/10.1287/msom.2022.1086>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



Manufacturing & Service Operations Management

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

Warranty Reserve Management: Demand Learning and Funds Pooling

Xiao-Lin Wang, Yuanguang Zhong, Lishuai Li, Wei Xie, Zhi-Sheng Ye

To cite this article:

Xiao-Lin Wang, Yuanguang Zhong, Lishuai Li, Wei Xie, Zhi-Sheng Ye (2022) Warranty Reserve Management: Demand Learning and Funds Pooling. *Manufacturing & Service Operations Management* 24(4):2221-2239. <https://doi.org/10.1287/msom.2022.1086>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2022, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Warranty Reserve Management: Demand Learning and Funds Pooling

Xiao-Lin Wang,^a Yuanguang Zhong,^b Lishuai Li,^c Wei Xie,^{b,*} Zhi-Sheng Ye^d

^aBusiness School, Sichuan University, Chengdu 610065, China; ^bSchool of Business Administration, South China University of Technology, Guangzhou 510641, China; ^cFaculty of Aerospace Engineering, Delft University of Technology, 2628 CD Delft, Netherlands; ^dDepartment of Industrial Systems Engineering and Management, National University of Singapore, Singapore 117576, Singapore

*Corresponding author

Contact: xlwang28-c@my.cityu.edu.hk,  <https://orcid.org/0000-0003-0100-8154> (X-LW); bmygzhong@scut.edu.cn,  <https://orcid.org/0000-0003-2805-8522> (YZ); lishuai.li@tudelft.nl,  <https://orcid.org/0000-0002-0990-5119> (LL); bmwxie@scut.edu.cn,  <https://orcid.org/0000-0002-5079-3315> (WX); yez@nus.edu.sg,  <https://orcid.org/0000-0001-5731-3911> (Z-SY)

Received: October 25, 2019

Revised: January 18, 2021; December 14, 2021

Accepted: December 24, 2021

Published Online in Articles in Advance: March 31, 2022

<https://doi.org/10.1287/msom.2022.1086>

Copyright: © 2022 INFORMS

Abstract. *Problem definition:* Warranty reserves are funds used to fulfill future warranty obligations for a product. In this paper, we investigate the warranty reserve planning problem faced by a manufacturing firm who manages warranties for multiple products. *Academic/practical relevance:* It is nontrivial to determine a proper amount of reserves to hold, because warranty expenditures are random in nature and reserving either excess or insufficient cash would incur losses. How can warranty reserve levels be optimized and promptly adjusted is a focal issue, especially for firms selling multiple products. *Methodology:* Inspired by the general pattern of empirical warranty claims data, we first develop an aggregate warranty cost (AWC) forecasting model for a single product by coupling stochastic product sales and failure processes, which can be used to plan for warranty reserves periodically. The reserve levels are then optimized via a distributionally robust approach, because the exact distribution of AWC is generally unknown. To reduce the losses generated from managing the funds, we further investigate two potential loss-reduction approaches: demand learning and funds pooling. *Results:* For the demand learning algorithm, we prove that, as the sales period grows, the optimal learning parameter asymptotically converges to a constant in a fairly fast rate; our simulation experiments show that the performance of demand learning is promising and robust under general warranty claim patterns. Moreover, we find that the benefits of funds pooling change over different stages of the warranty life cycle; in particular, the relative pooling benefit in terms of reserve losses is nonincreasing over time. *Managerial implications:* This study offers guidelines on how manufacturers should adaptively forecast and dynamically plan warranty reserves over the warranty life cycle.

Funding: This work was supported by the National Natural Science Foundation of China [Grants 71971085, 71871097, and 72071138] and the Natural Science Foundation of Guangdong Province [Grant 2020A1515011270].

Supplemental Material: The online supplement is available at <https://doi.org/10.1287/msom.2022.1086>.

Keywords: adaptive learning • distribution free • newsvendor • reserve management • risk pooling

1. Introduction

As one of the most important after-sales services, product warranty, which provides protection against premature failures for customers, has been adopted by most manufacturers to signal the quality and reliability of their products. In general, a warranty contract defines a specific protection period, in which free rectifications of or refund compensations for failed units will be offered by the manufacturer (Gallego et al. 2015, Pinçe et al. 2016). Although warranty services can address customers' concerns, managing a warranty program is costly from the manufacturer's perspective, which can account for as much as 15% of net sales. Essentially, warranty expenses are random

in nature, because they are closely related to two coupled stochastic processes—product sales and failure processes. For the manufacturer, servicing future warranty claims can incur liquidity risk, such as a shortage of cash, when unanticipated claims occur. A common solution to these issues adopted by manufacturers is to create an independent warranty reserve fund to cover contingent liabilities that arise from warranty obligations. As an industry practice, the main functions of a warranty reserve fund include: (1) signaling the firm's future performance to the stock market; (2) fulfilling contingent liabilities for future warranty service demands; and (3) playing as a tool of earnings management (Cohen et al. 2011).

However, an ever-increasing warranty reserve level is a major challenge for many manufacturers, especially the big ones. For example, Apple had a \$3.834 billion balance in its warranty reserve fund at the end of September 2017, and the figures for Ford and General Motors were \$5.031 billion and \$8.479 billion, respectively.¹ Many manufacturers, especially automakers, tend to keep much more money in their reserve funds. According to a report by *Warranty Week*,² U.S.-based manufacturers put aside 1.8% of product sales (on average) for future warranty claims, which is 17.2 times the amount they actually pay out in warranty claims per month. One of the main reasons is that many manufacturers are unable to accurately predict future warranty expenses, causing an over stock in warranty reserves. Conversely, it presents a bad image to customers if a manufacturer's actual warranty expense is beyond its expectation, which reduces its earnings and upsets its shareholders.³ Therefore, warranty reserve forecasting and management is an urgent yet challenging problem for manufacturers.


Typically, the *warranty life cycle*⁴ of a product follows the product life cycle to support the associated service claims (Murthy et al. 2004, Khawam et al. 2007). Servicing field warranty claims accounts for the biggest share of reserve consumptions, because the manufacturer is fully or partially responsible for the associated servicing expenses. Figure 1 shows the monthly warranty claims for three products of a leading electronics manufacturer over their warranty life cycles. One can observe that the warranty life cycle can be roughly partitioned into three stages: ramp-up, steady, and ramp-down. In the first stage, a new product is introduced to the market and its sales start growing. As a result, more and more potential failures will occur, resulting in an increasing number of warranty claims. After that, the warranty life cycle enters a mature stage in which both the number of units under warranty and

the number of claims remain roughly stable but with considerable variability. The reason is that new sales consecutively occur and warranties on earlier-sold units expire. Finally, the cycle reaches a ramp-down stage when the manufacturer discontinues the product, that is, no new sales are observed and the warranties on remaining units in service gradually expire. This pattern of warranty claims is quite general and has been reported by some existing studies (Khawam et al. 2007, Calmon and Graves 2017, Calmon et al. 2021). Despite the generality of this time-varying pattern, how to adaptively forecast future warranty claims and then dynamically schedule necessary warranty reserves, at the operational level, is underexplored.

This research is motivated by the warranty reserving issues faced by the aforementioned electronics manufacturer. The firm would like to determine optimal reserve levels throughout the warranty life cycles of its products to reduce the corresponding reserve losses. One should distinguish *warranty costs* from *reserve losses*: The former stands for the expenses resulting from the servicing of warranty claims, whereas the latter represents the costs or losses resulting from reserving excess and/or insufficient funds. In this work, we first develop an adaptive aggregate warranty cost (AWC) forecasting model to explicitly capture the underlying mechanism of warranty claim generation. In particular, to synthesize the uncertainties induced by product sales and failure processes, we derive two important statistics, the mean and (approximate) variance of AWC increment within an arbitrary time interval (Theorem 1), which are two key metrics for warranty reserve planning. The proposed model is a discrete variation of the AWC forecasting model in Xie and Ye (2016) and raises an obvious advantage compared with theirs (only focused on the AWC within $[0, t]$); that is, one can dynamically forecast the AWC increment and then plan the associated reserves within an arbitrary time interval (e.g., a fiscal quarter) to facilitate short-term and faster liquidity. Then, given the mean and variance of AWC increments, we derive optimal reserve levels periodically for a product throughout its warranty life cycle via a distributionally robust approach (Proposition 2). In addition, we compare the total reserve losses generated by the worst-case scenario with those by a benchmark—the simulated “true-distribution” scenario. The result shows that the worst-case reserving policy performs fairly well.

We further investigate two efficient approaches—demand learning and funds pooling—to explore potential reserve-loss-reduction opportunities. The demand learning model dynamically updates the reserving plan with field warranty data through an exponential smoothing-type mechanism. We prove that, as the sales period grows, the optimal learning parameter asymptotically converges to a constant in probability in a fairly fast rate (Propositions 3 and 4).

Figure 1. (Color online) Monthly Warranty Claims for Three Electronic Products



The performance of demand learning for general claim patterns is demonstrated via simulation experiments. Moreover, we analyze the benefits of funds pooling for a manufacturer who manages warranties for multiple products (e.g., Apple handles warranty claims for Mac, iPhone, and iPad). Under some mild conditions, we obtain an important finding, that is, the relative pooling benefit in terms of reserve losses decreases as the relative range of standard deviations increases (Proposition 6). In particular, we examine how pooling benefits change over different stages of the warranty life cycle. An interesting result is that when there are only two products, the relative pooling benefit in terms of reserve losses shows a nonincreasing trend (Corollary 2). Finally, a case study using a field data set is presented to demonstrate the effectiveness of the proposed reserve planning methodologies.

The remainder of this paper is structured as follows. Section 2 briefly reviews the relevant literature. Section 3 presents the adaptive AWC forecasting technique and the distributionally robust reserve planning model. The potential loss reduction approaches, that is, demand learning and funds pooling, are investigated and discussed in Sections 4 and 5, respectively. Section 6 presents a real-world case study. Section 7 concludes the paper. Some extended results and all proofs are relegated to the online supplement.

2. Literature Review

Warranty reserve management has been a crucial topic in the literature since the 1960s. Early efforts were mostly devoted to the estimation/prediction of warranty reserve demand (Menke 1969, Patankar and Worm 1981, Thomas 1989, Eliashberg et al. 1997, Ja et al. 2002). These studies focus predominately on estimating/predicting total warranty expenses over the warranty period or the entire product life cycle, which overlook the time-varying characteristics of warranty expenses and can only be used to determine warranty reserves from a long-term perspective. Nevertheless, the research related to (dynamic) cash management of warranty reserves is still scarce. Tapiero and Posner (1988) study a warranty reserving problem to specify the fraction of revenues added to the reserve fund after each sale. They define a compound Poisson process for warranty claims, which is generally not the case in reality. Buczkowski and Kulkarni (2006) aim at optimizing the initial reserve level and the contribution amount from each sale to ensure that the fund covers warranty liabilities over a given time period with a pre-specified probability. Gurgur (2011) tackles dynamic cash management of warranty reserves with the objective of minimizing total reserve losses. It is found that when there is no contribution after each sale, the optimal policy is a *base reserve* policy—resembling the base

stock policy in the inventory theory—that achieves a base reserve level at the beginning of each period by cash deposits or withdrawals. Fundamentally, our work is closely related to the papers that analogously treat cash flows as “commodity inventories” (i.e., modeling the cash management problem in a mathematically equivalent form to an inventory control problem). As an early attempt to exploit this analogy, Baumol (1952) applies the classical lot size model to solve a deterministic cash transactions problem. Then, Miller and Orr (1966) propose a two-parameter control limit policy for a stochastic cash balance problem. Following the two seminal works, cash management problems have received much attention (Eppen and Fama 1969, Neave 1970, Gormley and Meade 2007, Chen and Simchi-Levi 2009). In essence, warranty reserve management falls into the broad scope of corporate cash management that focuses on the control and planning of cash balance for business firms.

Our research differs from the previously mentioned ones in the following aspects: First, their models cannot characterize the dynamics in the process of generating warranty claims over the entire warranty life cycle, whereas the proposed AWC forecasting model overcomes this issue by fully coupling stochastic sales and failure processes. Second, because it is difficult to derive the exact distribution of warranty costs, they usually use the normal approximation technique to address this problem. In contrast, our research adopts a *distributionally robust* approach for reserve planning purposes, which is more conservative yet robust. Third, although our research simplifies the decision-making process involved in the dynamic reserve management problem to a static approximation, the exponential smoothing-type learning mechanism incorporates the “dynamics” into our model and improves its performance significantly. Finally, we take the same inventory view (precisely, in a “newsvendor” framework) as in the aforementioned cash management literature but deal with a distinct type of stochastic demand source—warranty reserve demand—that exhibits a general increasing-stable-decreasing pattern.

Another research stream—demand learning in inventory management (see Mišić and Perakis (2020) for a recent overview)—is also closely related to our work. Existing studies on this topic focus on either single-period or multiperiod inventory control settings. In the former (i.e., newsvendor) setting, demand distribution is generally assumed to be unknown while historical data are available (Huber et al. 2019, Besbes and Mouchtaki 2021, Keskin et al. 2021). In the latter setting, demands across periods are usually assumed to be independent and identically distributed (Chen and Chao 2019, Zhang et al. 2020). In general, the warranty reserving problem is a multiperiod problem and, from Figure 1, the reserve demands across

periods are essentially dependent and nonidentically distributed. This necessitates new data-driven methodologies that are capable to learn such demand from data. In this work, we propose an exponential smoothing-type learning mechanism, in which the optimal learning parameter is sequentially optimized by minimizing empirical total reserve losses directly. By leveraging the optimization problem structure (i.e., the objective function) in the learning process, we combine learning and optimization together to generate a better reserving plan. This also provides a response to Mišić and Perakis (2020) who call for new approaches to *predict-then-optimize* in operations management applications.

Finally, the paper contributes to the risk pooling literature in operations management as well. The idea of funds pooling is similar to that of inventory pooling, where demand variability is mitigated by aggregating demands across products or locations to reduce safety stocks and consequently inventory costs (Eppen 1979, Benjaafar et al. 2005, Berman et al. 2011). In recent years, many efforts have been devoted to studying the behaviors of risk pooling in various inventory settings, for example, heavy-tailed demands (Bimpikis and Markakis 2016), multilocation newsvendor (Govindarajan et al. 2021, Yang et al. 2021), and resource allocation (Karsten et al. 2015, Zhong et al. 2018). Nevertheless, most previous literature on inventory pooling deals only with the variability of product sales and considers relatively simple forms of demand distributions. The time-varying characteristics of warranty reserve demands—a new finite horizon demand source—over the warranty life cycle distinguish our problem from traditional inventory pooling issues and pose a challenge to modeling and analysis. In addition, the benefits of risk pooling in a multiperiod inventory setting have received limited attention (Tagaras and Cohen 1992, Bimpikis and Markakis 2016). In this research, we are particularly interested in an important perspective of risk pooling, *time-varying benefit behavior*, in multiperiod warranty reserve management.

3. Warranty Reserve Forecasting and Planning

In this section, we first develop a discrete-time reserve demand forecasting model that is adaptive to the uncertainties from product sales and failure processes. Then, we use a distributionally robust optimization model to periodically determine optimal reserve levels and evaluate the performance of the worst-case reserving model.

3.1. Warranty Reserve Demand Forecasting

Let L be the life cycle (sales period) of a product, that is, from its launch time to the date it is discontinued. Empirically, sales periods are finite and short for consumer electronics, due to fast technology obsolescence. For instance, Apple releases new generations of its

products every year while previous generations are soon discontinued. Denote $N(t)$ as the cumulative number of units sold within $[0, t]$, $t \in [0, L]$.

Assumption 1. *The cumulative product sales $N(t)$ follows a homogeneous Poisson process (HPP) with rate λ (i.e., $HPP(\lambda)$).*

This assumption is relatively strong, especially when the magnitude of sales volume is large. However, it is most natural to assume that new-product demand arrivals follow an HPP, when sales curves do not exhibit obvious patterns. In addition, this assumption might also be acceptable for the sales of parts, which is relatively stationary due to their long life cycles.⁵ Nevertheless, we acknowledge that this assumption is to make the problem analytically tractable. To relax the HPP assumption and generalize the adaptivity of our method to a larger scope of products with nonstationary sales, in Section 4.3, we will incorporate a nonhomogeneous Poisson process (NHPP)—which can capture a variety of time-varying patterns—to simulate real-world sales processes to further validate the robustness of the proposed methodologies.

Assumption 2. *The time to first failure of this product follows an exponential distribution $G(t; \theta)$ with a constant failure rate θ .*

The failure processes of most products exhibit three phases—infant mortality, useful life, and wear out, which correspond to a decreasing failure rate, a constant failure rate, and an increasing failure rate, respectively. In real applications, a constant failure rate has been demonstrated to be a good approximation for a unit that is under warranty (Blischke and Murthy 2000). This is because (i) the units must pass quality and reliability testing (e.g., burn-in and screening tests) before being released into the market, for which most of them have passed through the infant mortality period and entered the useful life period, and (ii) the warranty periods for most products are much shorter than their useful life periods. In particular, the failure rates for most electronics are, inherently, almost constant (Blischke and Murthy 2000).

Assumption 3. *The product is sold with a nonrenewing free replacement warranty policy of length W .*

Under this policy, if a unit is sold at time zero and fails at time $t \in [0, W]$, then the manufacturer will provide a free replacement that comes with a remaining warranty period of $W - t$. The nonrenewing free replacement warranty policy has been widely adopted for products that are nonrepairable or, economically, not worth repairing—for example, smartphones, battery packs, and some electric tools. Under this warranty policy, a failed unit or its critical component will be restored to a functioning state that is as good as new, through replacement by an identical new one.

Based on the aforementioned assumptions, we now introduce an adaptive model that can fully capture the interaction of product sales and under-warranty failures to forecast the AWC increments, which forms the foundation for short-term warranty reserve planning. Suppose that the warranty period is shorter than the sales period (i.e., $W \leq L$), which is quite common in practice. The warranty life cycle $[0, W + L]$ is of interest to us. Let T_i ($i \in \{1, 2, \dots, N(t)\}$) be the random purchase date of the i th sold unit. Let $C_i(\tau_i)$ ($\tau_i = \min\{t - T_i, W\}$) quantify the random warranty cost of the i th sold unit up to time $t \geq T_i$. Given T_i , the expected warranty cost of the i th sold unit up to time t can be calculated by the well-known renewal equation:

$$\begin{aligned}\mathbb{E}[C_i(\tau_i)] &= C_W m(\tau_i) = C_W \left[G(\tau_i; \theta) + \int_0^{\tau_i} m(\tau_i - x) dG(x; \theta) \right] \\ &= C_W \theta \tau_i, i \in \{1, 2, \dots, N(t)\},\end{aligned}\quad (1)$$

where C_W is the average replacement cost per failure and $m(\cdot)$ is the renewal function associated with $G(t; \theta)$. Servicing a warranty claim incurs various costs: administration, transportation, labor, and materials, among others. For the sake of convenience, we combine these costs into a single term C_W .

Define *warranted base* as the number of sold units that are still under warranty. Figure 2 illustrates the counting process of AWC during the warranty life cycle. It is important to point out that all units are eligible to generate warranty claims during their warranty periods, whereas the warranted base is changing over time. Because different units have different dates of purchase, the AWC up to time t can be formulated as

$$WC(t) = \begin{cases} \sum_{i=1}^{N(t)} C_i(\tau_i), & 0 \leq t \leq L, \\ \sum_{i=1}^{N(L)} C_i(\tau_i), & L < t \leq L + W. \end{cases}\quad (2)$$

Industry practice shows that warranty reserves are planned periodically, for example, quarterly, semiyearly, or yearly, for accounting purposes (Cohen et al. 2011). In accordance with this practice, we need to evaluate the AWC increments within a sequence of equal-length time intervals over the warranty life cycle. Suppose that the length of reserve planning cycle is Δ , say, one quarter. We assume that the ratios of warranty period W and sales period L to Δ are integers, which is a common practice. Without loss of generality, we rewrite $W = m\Delta$ ($m \geq 1$), $L = n\Delta$ ($n \geq m$ as $L \geq W$), and $L + W = (n + m)\Delta$. Hence, the total number of planning periods is $n + m$ over the warranty life cycle. Let $WC_k = WC(k\Delta) - WC((k-1)\Delta)$ denote the AWC increment in the k th planning period ($k \in \{1, 2, \dots, n + m\}$). In general, it is difficult, if not impossible, to obtain the exact distribution of WC_k because of the complexity of warranty costs (Ja et al. 2002, Buczkowski and Kulkarni 2006, Gurgur 2011). Fortunately, we can derive closed-form expressions of the mean and approximate variance of WC_k , which are two essential statistics for further analysis.


Theorem 1. Let μ_k and σ_k^2 be the mean and variance of WC_k , namely, $\mu_k = \mathbb{E}[WC_k]$ and $\sigma_k^2 = \mathbb{V}\text{ar}(WC_k)$. Then, one has

$$\mu_k = \begin{cases} C_W \lambda \theta \Delta^2 \left(k - \frac{1}{2}\right), & k = 1, \dots, m, \\ C_W \lambda \theta \Delta^2 m, & k = m + 1, \dots, n, \\ C_W \lambda \theta \Delta^2 \left(n + m - k + \frac{1}{2}\right), & k = n + 1, \dots, n + m, \end{cases}\quad (3)$$

and

$$\sigma_k^2 \approx \begin{cases} C_W^2 \lambda \theta \left[\Delta^2 \left(k - \frac{1}{2}\right) + \theta \Delta^3 \left(k^2 - k + \frac{1}{3}\right) \right], & k = 1, \dots, m, \\ C_W^2 \lambda \theta (1 + \theta m \Delta) \Delta^2 m, & k = m + 1, \dots, n, \\ C_W^2 \lambda \theta \left[\Delta^2 \left(n + m - k + \frac{1}{2}\right) + \theta \Delta^3 \left(m^2 - n^2 - n + k + 2nk - k^2 - \frac{1}{3}\right) \right], & k = n + 1, \dots, n + m. \end{cases}\quad (4)$$

Figure 2. (Color online) AWC Accumulation Process During the Warranty Life Cycle



It is worth pointing out that the asymptotic approximation in (4) is fairly accurate under practical settings.⁶ Hereafter, we will use the terms *approximate variance* and *variance* interchangeably. The mean μ_k and variance σ_k^2 are inherently correlated, because both are dependent on k , C_W , λ , and θ . In this scenario, we cannot change one while keeping the other unchanged. This property renders the warranty reserving problem different from traditional inventory problems, in which the mean and variance of product demands are mostly assumed to be independent.

In practice, it is not uncommon that the means and variances of AWC increments for different planning periods or products might have distinct magnitudes. As a consequence, the coefficient of variation, $CV_k = \sigma_k / \mu_k$, is a useful alternative to measuring the risks and comparing the uncertainties involved in AWC increments for different planning periods or products. The following lemma shows the role of CV_k .

Lemma 1. For any $\epsilon > 0$, we have

$$\Pr \left\{ \left| \frac{WC_k}{\mu_k} - 1 \right| > \epsilon \right\} \leq \frac{CV_k^2}{\epsilon^2}, k = 1, 2, \dots, n + m.$$

Lemma 1 tells us that given an acceptable error ϵ , the variability of WC_k relative to estimated μ_k can be well controlled by the bound CV_k^2 / ϵ^2 . This means that the risks of relative forecast errors follow the same pattern as that of CV_k for different periods and/or products. According to Equations (3) and (4), CV_k can be approximated by

$$CV_k \approx \begin{cases} \frac{\sqrt{(k-1/2) + \theta\Delta(k^2 - k + 1/3)}}{\sqrt{\lambda\theta\Delta(k-1/2)}}, & k = 1, \dots, m, \\ \frac{\sqrt{1 + \theta m\Delta}}{\sqrt{\lambda\theta m\Delta}}, & k = m + 1, \dots, n, \\ \frac{\sqrt{(n+m-k+1/2)}}{\sqrt{\lambda\theta\Delta(n+m-k+1/2)}}, & k = n + 1, \dots, n + m. \end{cases} \quad (5)$$


Proposition 1. The terms μ_k and σ_k^2 are increasing in $k \in \{1, \dots, m\}$, constant in $k \in \{m + 1, \dots, n\}$, and decreasing in $k \in \{n + 1, \dots, n + m\}$. In contrast, with respect to the same counting sets, CV_k shows an opposite decreasing-constant-increasing trend.

The tendencies of mean μ_k , variance σ_k^2 , and coefficient of variation CV_k over the warranty life cycle are illustrated in Figure 3. One can observe that the curves of both μ_k and σ_k^2 exhibit an upside-down bathtub shape, whereas that of CV_k is bathtub shaped. Basically, the trends of μ_k and σ_k^2 are analogous to the pattern in Figure 1, whereas the time-varying behavior of CV_k over the warranty life cycle is an interesting result. In the first stage of the warranty life cycle (i.e., $1 \leq k \leq m$), the variability indicated by CV_k shows a downward trend as time goes by. This is because the larger the warranted base, the lower the risk (Calmon and Graves 2017). Then, in the second stage (i.e., $m + 1 \leq k \leq n$), demand variability becomes stable because the warranted base moves into a relatively steady state. Finally, the warranted base shrinks in the third stage (i.e., $n + 1 \leq k \leq n + m$), which in turn increases the risk caused by uncertainties. That is to say, the AWC variability measured by the coefficient of variation decreases (respectively, increases) as the expected AWC increases (respectively, decreases). This result is consistent with our investigation of the behavior of warranted base over the warranty life cycle. It demonstrates the importance and necessity of adopting the coefficient of variation to capture demand variability in relation to mean μ_k .

3.2. Distributionally Robust Warranty Reserve Planning

Based on the previous reserve demand forecast, we develop an optimization model to periodically plan for warranty reserves over the warranty life cycle. Setting a proper amount of warranty reserves is nontrivial in

Figure 3. (Color online) Patterns of μ_k , σ_k^2 , and CV_k for Three Products



Notes. The parameters are arbitrarily set as $W = 2$ years, $L = 5$ years, $\Delta = 1/4$ year, $\lambda_1 = 210,000$, $\theta_1 = 0.024$, $C_{W1} = \$100$, $\lambda_2 = 120,000$, $\theta_2 = 0.015$, $C_{W2} = \$150$, $\lambda_3 = 135,000$, $\theta_3 = 0.010$, and $C_{W3} = \$75$, for illustrative purposes. In this setting, the warranty life cycle is $L + W = 7$ years, $m = W/\Delta = 8$, and $n = L/\Delta = 20$.

general, because the warranty expenditures are random—induced by stochastic product sales and failure processes—and both over reserving and under reserving may cause additional losses. More specifically, over reserving results in an opportunity cost in the form of lost interest, whereas under reserving might necessitate calling for emergency funds from other sources at an above-average interest rate (Gurgur 2011). Therefore, the risk of holding excess or insufficient reserves should be taken into consideration when optimizing the warranty reserve levels.

In the previous section, we have derived the mean and variance of the AWC increment in any reserve planning period. Under the assumption that there is no cash contribution to the reserve fund from each sale, the AWC increment in each period is exactly the reserve demand. As mentioned earlier, however, its exact distribution is generally unknown. To overcome this difficulty, we adopt the well-known distributionally robust approach in Scarf (1958) to determine the optimal warranty reserve level for each period. Let C_h be the unit over reserving cost and C_b be the unit under reserving cost, which can be estimated by the time values of money.⁷ Unlike the material property of commodity inventories, the monetary property of warranty reserves allows the manufacturer to calculate reserve levels for different periods independently (without leftover from previous periods). For a specific period $k \in \{1, 2, \dots, n + m\}$, the manufacturer should determine an optimal amount of warranty reserves to minimize the associated expected reserve loss $\pi_k(WR_k)$, which is expressed as

$$\begin{aligned}\pi_k(WR_k) &= \mathbb{E}[C_h(WR_k - WC_k)^+ + C_b(WC_k - WR_k)^+] \\ &= C_h \mathbb{E}[(WR_k - WC_k) + (WC_k - WR_k)^+] \\ &\quad + C_b \mathbb{E}[WC_k - WR_k]^+ \\ &= C_h WR_k - C_h \mu_k + (C_h + C_b) \mathbb{E}[WC_k - WR_k]^+, \end{aligned} \quad (6)$$

where $x^+ = \max\{0, x\}$.

Because deriving the exact distribution of WC_k is highly intractable, we minimize Equation (6) against the worst case. According to Gallego and Moon (1993), the following lemma holds.

Lemma 2. *The following inequality holds.*

$$\mathbb{E}[WC_k - WR_k]^+ \leq \frac{\sqrt{\sigma_k^2 + (WR_k - \mu_k)^2} - (WR_k - \mu_k)}{2}. \quad (7)$$

Moreover, for each WR_k , there exists a distribution of WC_k in which the bound (7) is tight.

Based on Lemma 2, the distributionally robust warranty reserving problem turns to minimize the following

upper bound:

$$\begin{aligned}\hat{\pi}_k(WR_k) &= C_h WR_k - C_h \mu_k + \frac{1}{2}(C_h + C_b) \\ &\quad \times \left[\sqrt{\sigma_k^2 + (WR_k - \mu_k)^2} - (WR_k - \mu_k) \right]. \end{aligned} \quad (8)$$

In general, it is reasonable to assume $C_h < C_b$ in real applications. Then, the following result can be obtained.

Proposition 2. *The distributionally robust optimal warranty reserve level for the k th planning period is given by*

$$WR_k^* = \mu_k + \frac{1}{2} A \sigma_k, \quad (9)$$

where $A = \sqrt{C_b/C_h} - \sqrt{C_h/C_b} > 0$.

Substituting Equation (9) into Equation (8), we can obtain the minimized expected reserve loss in the k th period as


$$\hat{\pi}_k^*(WR_k^*) = \frac{1}{2} B \sigma_k, \quad (10)$$

where $B = C_h(\sqrt{1 + A^2/4} + A/2) + C_b(\sqrt{1 + A^2/4} - A/2) > 0$.

Remark 1. The optimal warranty reserve level WR_k^* increases as C_b/C_h , μ_k , or σ_k increases, whereas the expected reserve loss $\hat{\pi}_k^*(WR_k^*)$ is independent of μ_k but increasing in σ_k .

An intuitive result from Remark 1 is that in any planning period, if the ratio of C_b/C_h or the mean and/or standard deviation of AWC increment becomes larger, the manufacturer should specify a higher reserve level, whereas the expected reserve loss resulting from over-reserving and/or under-reserving depends only on the standard deviation of AWC increment. Figure 4 shows the warranty reserve levels and the associated losses for the three products (with the same parameter setting as in Figure 3). As observed, the warranty reserve levels exhibit similar trends over the warranty life cycle, which can be well explained by the time-varying warranted base. In addition, the tendency of expected reserve losses is related to, yet slightly different from, that of warranty reserves.

Although the distributionally robust optimization technique can generate a robust reserving plan against the worst case, the plan tends to be conservative in most cases, possibly causing additional losses. To assess the performance of the proposed worst-case reserving model, we compare it with a benchmark—the simulated nonparametric reserving model. Specifically, based on Assumptions 1–3, we implement Monte Carlo simulation to generate an empirical distribution of the AWC increment in each period, which can be regarded as the “true distribution” and serves as a perfect benchmark for the worst-case performance

Figure 4. (Color online) Warranty Reserve Levels and Expected Reserve Losses of the Three Products ($C_h = \$0.02$ and $C_b = \$0.05$)

assessment. Our simulation studies show that, compared with the benchmark, the distributionally robust reserving model performs fairly well under a wide range of practically reasonable parameter settings. For more details, please refer to Section EC.1 in the online supplement.

In what follows, based on the worst-case reserving model, we investigate two possible reserve-loss-reduction approaches: demand learning and funds pooling. The results will be discussed in two separate sections for the sake of convenience.

4. Reserve Loss Reduction via Demand Learning

In real applications, warranty claims data are usually subject to a high degree of variability due to complex sales and failure patterns. For instance, if the sales record of a product presents a nonstationary pattern, then the reserving model under the HPP sales assumption cannot guarantee a satisfactory performance. As a result, although the proposed simple model in Section 3 delivers a fairly good performance under some specific assumptions, there is still a potential to improve its adaptivity and performance for general claims data. In this section, we aim to augment the performance of our simple reserving model through demand learning.

4.1. Exponential Smoothing-Type Learning

As discussed before, most of the existing demand learning models in the inventory literature cannot be directly applied in our context, because warranty reserve demands across periods are dependent and nonidentically distributed. Another practical hurdle is that one cannot repeatedly collect warranty claims data for a specific product; actually, only one single data point can be collected in each planning period. Therefore,

there is not enough data to learn the distributional properties of warranty reserve demands across periods.

Inspired by the demand-chasing heuristics in Bolton and Katok (2008) and Bostian et al. (2008), we implement a simple exponential smoothing-type mechanism to update the reserving plan derived from the distributionally robust model using observed warranty data, in which the smoothing factor is learnt dynamically to chase the time-varying reserve demands. Specifically, the adjusted reserve level $\widehat{WR}_k$ in the k th planning period is formulated as

$$\widehat{WR}_k = WR_k^* + \phi_k(\Delta Q_{k-1} - \mu_{k-1}), k = 2, 3, \dots, \quad (11)$$

where $\phi_k \geq 0$ is the smoothing/learning factor (which is dynamically updated from period to period), WR_k^* is the “optimal” reserve level given by Equation (9), and ΔQ_k represents actual warranty expenses (i.e., the AWC increment) in period k . In essence, the learning model anchors on the reserve level WR_k^* derived from the distributionally robust model and tends to pull the reserve level down when the forecasted mean μ_{k-1} in the previous period is high and vice versa.

It is worth noting that, unlike the smoothing factor in a classic exponential smoothing model, the learning parameter ϕ_k could take a value that is greater than one. The optimal value of ϕ_k^* in period $k \in \{3, 4, \dots\}$ is obtained by solving the following optimization problem:

$$\min_{\phi} \sum_{i=1}^{k-1} \{C_h[WR_i^* + \phi(\Delta Q_{i-1} - \mu_{i-1}) - \Delta Q_i]^+ + C_b[\Delta Q_i - WR_i^* - \phi(\Delta Q_{i-1} - \mu_{i-1})]^+\}. \quad (12)$$

In the first period, we set $\widehat{WR}_1 = WR_1^*$, whereas the learning parameter ϕ_2 in the second period could be arbitrarily chosen (say, $\phi_2 = 1$) for setup purposes. In this manner, we learn the parameter ϕ_k^* by minimizing

empirical total reserve losses directly, using the collected data from the first $k - 1$ periods. By leveraging the objective function in the learning process, we attempt to combine learning and optimization together to generate a better reserving plan. It is straightforward to verify that the optimization problem in (12) is convex; thus, an optimal ϕ_k^* exists.

4.2. Convergence of the Learning Algorithm

We are now in a position to explore some properties of the proposed learning mechanism. We first examine its asymptotic characteristic. To this end, define a convex function $\mathbb{H}(\phi) = \frac{1}{2}C_h A \sigma_{m+2} + (C_b + C_h)\mathbb{E}[\Delta Q_{m+2} - WR_{m+2}^* - \phi(\Delta Q_{m+1} - \mu_{m+1})]^+$, and let ϕ^* be its minimizer. Then, we have the following result.

Proposition 3. *Given the assumed model and planning cycle Δ , ϕ_k^* converges to ϕ^* in probability when the sales period $L = n\Delta \rightarrow \infty$.*

Proposition 3 shows that under the aforementioned assumptions, the optimal learning parameters obtained from successive periods asymptotically converge to a constant in probability if the sales period is long enough. This asymptotic characteristic guarantees a stable output of the learning mechanism in the long run. In practice, however, warranty reserve planning horizons (precisely, the warranty life cycles) are never “infinitely long,” especially for consumer electronics that are depreciating quickly. Consequently, the rate of learning, that is, the convergence rate of ϕ_k in our case, provides a more important indication that whether the learning mechanism would perform well for products with short warranty life cycles.

Proposition 4. *Given the assumed model and planning cycle Δ , suppose $\phi \mapsto \mathbb{H}(\phi)$ is second order differentiable at $\phi = \phi^*$. Then, $\phi_k^* - \phi^* = O_p(1/\sqrt{k})$ and $\sqrt{k}(\phi_k^* - \phi^*)$ converges weakly to a mean-zero Gaussian distribution with a finite asymptotic variance.*

Proposition 4 reveals that the rate of convergence of the learning parameter is $1/\sqrt{k}$ and is asymptotically normal. This is to be expected since ϕ_k^* can be regarded as an M -estimator. Our finite sample simulation in Section 4.3 also reveals that when k is greater than 10, ϕ_k^* is close enough to ϕ^* in most cases. This demonstrates that the proposed demand learning mechanism is expected to deliver an acceptable performance for products with short warranty life cycles.

In addition, Propositions 3 and 4, combined together, provide a meaningful guidance to algorithmic implementation in real applications. That is, when a manufacturer collects historical warranty claims data of a product for a relatively long period of time, it is expected that the optimal learning parameter ϕ_k^* will converge to a constant ϕ^* in probability as time goes by. In this situation, the optimization step in Equation (12) can be

skipped and the adjusted reserve level $\widehat{WR}_k$ can be updated by directly substituting ϕ^* into Equation (11), which simplifies the process.

4.3. Performance Evaluation with General Sales and Failure Patterns

The results in Propositions 3 and 4 hinge heavily on the underlying assumptions on product sales and failure processes (i.e., Assumptions 1 and 2). However, it is difficult, if not impossible, to explore the case when the assumptions are violated. Thus, we resort to simulation experiments to evaluate the performance of the demand learning mechanism. For this purpose, a general setting of sales and failure patterns is adopted to generate warranty claims.

In the simulation experiment, we first simulate sales data by synthesizing the NHPP and the well-known Bass model (Bass 1969). The Bass model is probably the most influential new-product diffusion model. It postulates that the trajectory of expected cumulative sales of a new product follows a deterministic function of time with three parameters: the potential market size κ , the innovator factor p , and the imitator factor q . For the NHPP-Bass model, the expected cumulative sales can be expressed as

$$\Lambda(t) = \kappa \frac{1 - e^{-(p+q)t}}{1 + \frac{q}{p}e^{-(p+q)t}}, \quad (13)$$

and the saturation point is $t^* = \ln(q/p)/(p+q)$, that is, the sales rate reaches its peak at t^* . Given the i th purchase date t_i ($i \in \{1, 2, \dots\}$), we can generate the i th interpurchase time x_i from the distribution function $Z_i(x) = 1 - \exp(-\Lambda(t_i + x) + \Lambda(t_i))$. Then, the $(i+1)$ th purchase date is simply $t_{i+1} = t_i + x_i$.

Moreover, to relax the product failure assumption, we assume that the product's time to first failure follows a Weibull distribution with $G(t; \eta, \beta) = 1 - \exp\{-(t/\eta)^\beta\}$, where η and β are scale and shape parameters, respectively. Simulation is conducted in the following steps. For each simulation run, following the illustration in Figure 2, we first generate all the purchase dates within $[0, L]$ according to the NHPP-Bass model. Given the purchase date of each sold unit, we then generate the associated under-warranty failure instant(s) based on the Weibull distribution $G(t; \eta, \beta)$. Finally, counting the total number of warranty claims and the corresponding costs within each planning period yields the simulated AWC increments over the entire warranty life cycle.

Let S be the total number of simulation runs and $\Delta Q_{i,k}$ ($i \in \{1, 2, \dots, S\}$, $k \in \{1, 2, \dots, m+n\}$) be the simulated AWC increment for planning period k in simulation run i . For each simulation run i , we need to estimate the parameters λ_i and θ_i through least squares (see Section 6 for details). With estimated $\hat{\lambda}_i$ and $\hat{\theta}_i$,

we can calculate the mean $\mu_{i,k}$ and variance $\sigma_{i,k}^2$ of the AWC increment in period k by Equations (3) and (4) and then determine the corresponding optimal reserve level $WR_{i,k}^*$ by Equation (9). Based on $\{\Delta Q_{i,k}\}_{k \in \{1,2,\dots,m+n\}}$ and $\{WR_{i,k}^*\}_{k \in \{1,2,\dots,m+n\}}$, the adjusted reserve level $\widehat{WR}_{i,k}$ can be obtained by Equation (11). Then, substituting $\widehat{WR}_{i,k}$ and $WR_{i,k}^*$ into $\tilde{\pi}_{i,k}(WR_{i,k}) = C_h(WR_{i,k} - \Delta Q_{i,k})^+ + C_b(\Delta Q_{i,k} - WR_{i,k})^+$ yields the simulated reserve losses, with and without learning, respectively. To evaluate the learning benefit, we define $\Omega_i = (\sum_{k=1}^{m+n} \tilde{\pi}_{i,k}(WR_{i,k}^*) - \sum_{k=1}^{m+n} \tilde{\pi}_{i,k}(\widehat{WR}_{i,k})) / \sum_{k=1}^{m+n} \tilde{\pi}_{i,k}(WR_{i,k}^*)$ to represent the percentage of total reserve loss reduction resulting from demand learning, for each run i .

In this simulation experiment, we propose to examine the algorithm performance with respect to two critical factors—the location of sales peak (t^*) and the shape of failure rate function ($\beta t^{\beta-1}/\eta^\beta$). We consider three types of sales peak, that is, early ($t^* < L/2$), middle ($t^* \approx L/2$), and late ($t^* > L/2$), as well as three shapes of failure rate function, that is, constant ($\beta = 1$), concave increasing ($1 < \beta < 2$), and convex increasing ($\beta > 2$). To this end, we fix the values of κ , q , and η at $\kappa = 200,000$, $q = 0.5$, and $\eta = 5$, respectively, and alter the values of p and β to satisfy the aforementioned settings. More specifically, we consider $p = 0.238$ ($t^* \approx 1$), $p = 0.109$ ($t^* \approx 2.5$), and $p = 0.054$ ($t^* \approx 4$) for the three types of sales peak, and $\beta = 1.0$, $\beta = 1.5$, and $\beta = 3.5$ for the three shapes of failure rate function, respectively (see Figure 5 for illustrations). Thus, we have nine ($= 3 \times 3$) different yet comprehensive scenarios in total. Other parameters are set as $W = 2$ years, $L = 5$ years, $\Delta = 1/4$ year, $C_h = \$0.02$, $C_b = \$0.05$, and $C_W = \$100$.

To illustrate the performance of demand learning, we arbitrarily simulate one sample for each of the nine scenarios. Figure 6 shows the corresponding

simulated warranty costs and planned reserve levels (with and without learning) over the warranty life cycle. It is clear that the curves of simulated warranty costs are generally right-skewed (respectively, left-skewed) when the sales peak occurs early (respectively, late), whereas they are relatively symmetric when the sales peak locates at the middle of the sales period. More interestingly, we can see that the simulated cost curves in panels (b) and (c) are more symmetric than that in panel (a). This phenomenon is driven by the influence of failure rate function: When the shape parameter β is greater than one, the failure rate function increases faster over time (Figure 5), implying that more warranty claims tend to occur later. With a similar logic, although the simulated cost curves in the second row of Figure 6 are relatively symmetric, those in panels (e) and (f) tend to be left-skewed due to the impact of increasing failure rates. Furthermore, all the simulated cost curves in the third row are left-skewed, whereas the skewness of the curve in panel (g) is the smallest.

We note that the curves of mean AWC increments in Equation (3) are symmetric with respect to the midpoint of the warranty life cycle (Figures 3 and 6). As a result, when simulated cost curves are more symmetric, the proposed simple model (with simplified sales and failure processes) performs better for curve fitting. When the curves of simulated warranty costs are either left- or right-skewed, the distributionally robust optimal reserve levels deviate from simulated costs significantly. Nevertheless, the reserve levels adjusted by the learning algorithm are quite close to the simulated costs in all scenarios. In addition, Figure 7 illustrates the evolution of optimal learning parameters over the warranty life cycle for the nine scenarios in Figure 6. One can observe that in general sales and failure settings, the optimal learning parameters tend to converge as well. Moreover, the rate of convergence

Figure 5. (Color online) Sales Rates (a) and Failure Rates (b) Under Various Parameter Settings





Figure 6. (Color online) Simulated Warranty Costs and Planned Reserve Levels (Without and with Learning) for Nine Scenarios



is fairly fast: when k is greater than 10, ϕ_k^* is close enough to ϕ^* in most scenarios. This implies that the exponential smoothing-type mechanism is capable to learn warranty reserve demands for products with both long and short life cycles.

We further examine the cost performance of demand learning through a large number of simulation runs. Figure 8 shows the box plots of Ω_i over 1,000 simulation runs for the nine scenarios. It is clear from this figure that demand learning can lead to a

significant reduction in total reserve losses. The value of Ω_i can be as high as 83.4% (the maximum value in scenario 9). An interesting observation is that the smaller (respectively, larger) the discrepancy between analytical model and simulated curve, the smaller (respectively, larger) the learning benefit Ω_i . As discussed previously, the discrepancies between analytical model and simulated curve in scenarios 3, 4, and 7 (corresponding to panels (c), (d), and (g) in Figure 6) are comparatively small in their respective subgroups.

Figure 7. (Color online) Convergence of Optimal Learning Parameters ϕ_k^* in the Corresponding Nine Scenarios in Figure 6





Figure 8. (Color online) Box Plots of Ω_i in the Nine Scenarios



As a result, the benefits of demand learning in these scenarios are relatively small. More specifically, the learning benefit in scenario 3 (respectively, 4 and 7) is the smallest among the first (respectively, middle and last) three scenarios. This observation indicates that the proposed demand learning algorithm performs better when the AWC forecasting model fails to fully capture the dynamics of actual warranty costs.

Another interesting phenomenon is regarding the dispersion—reflected by the interquartile range of each box—of learning benefits Ω_i over the 1,000 simulation runs. One can observe that scenarios 4, 5, and 6 (corresponding to middle sales peak) have a larger dispersion, relative to those for early and late sales peaks. The reason is that under this case, the analytical model matches the simulated curve better, for which the learning benefits are subject to a higher degree of randomness. In addition, when $\beta = 3.5$ (corresponding to scenarios 3, 6, and 9), the associated dispersion is relatively larger as well. This is because when $\beta = 3.5$, the failure rate function over the warranty period (i.e., $0 \leq t \leq 2$ years) is lower (Figure 5). Thus, the expected number of warranty claims per unit sold would be fewer. As a result, the AWC variability would enlarge, which in turn translates into a larger dispersion of Ω_i .

5. Reserve Loss Reduction via Funds Pooling

In practice, a manufacturing firm usually produces and sells multiple products at the same time, and the firm is thus responsible to maintain enough reserves to fulfill warranty obligations for these products. Because warranty expenditures are random and both over reserving and under reserving risks are involved in warranty reserve planning, it is reasonable to expect, just like with physical inventories, that pooling warranty reserves for multiple products could reduce reserve levels and reserve losses. In particular, given the nature of reserves—cash that can be freely used for any product sold by the firm and easily added to or withdrawn from a fund—pooling is especially simple and attractive. However, it is surprising that a considerable number of firms have not yet pooled their warranty reserves.⁸ In this section, we aim to investigate the warranty reserve pooling scheme to quantify the corresponding (time-varying) pooling benefits, which is able to reveal the value of funds pooling.

5.1. Funds Pooling Scheme

We consider a manufacturing firm that manages the warranties for M ($M > 1$) different products. The following assumption is then introduced.

Assumption 4. *The M products are released onto the market at the same time and share the same warranty period W and sales period L .*

Although this assumption is adopted mainly for the sake of simplicity and tractability, it is not uncommon in reality. For example, Apple often releases products (e.g., iPhone, iPod, iPad, Mac) of different generations simultaneously. The warranty periods of these products are typically one year, and their sales periods are roughly the same.

Hereafter, we use the following notation to facilitate our presentation. Subscript j is attached to each parameter to indicate the j th product ($j \in \{1, 2, \dots, M\}$). For instance, we use λ_j , θ_j , and C_{Wj} to denote the sales rate, failure rate, and replacement cost of the j th product, respectively. In this way, the optimal warranty reserve level and the corresponding expected reserve loss, for the j th product in the k th planning period, become $WR_{j,k}^* = \mu_{j,k} + \frac{1}{2}A\sigma_{j,k}$ and $\hat{\pi}_{j,k}^*(WR_{j,k}^*) = \frac{1}{2}B\sigma_{j,k}$, respectively (according to Equations (9) and (10)). Hence, without pooling, the total amount of warranty reserves planned for the M products in the k th period can be summed up as

$$TWR_k^* = \sum_{j=1}^M WR_{j,k}^* = \sum_{j=1}^M \mu_{j,k} + \frac{A}{2} \sum_{j=1}^M \sigma_{j,k}. \quad (14)$$

Conversely, the firm can run a combined fund for the M products to reduce reserve levels and pool risks. In the pooled case, the distributionally robust optimal warranty reserve level PWR_k^* in the k th period can be determined by minimizing

$$\begin{aligned} \pi_k^c(PWR_k) &= \mathbb{E} \left[C_h \left(PWR_k - \sum_{j=1}^M WC_{j,k} \right)^+ + C_b \left(\sum_{j=1}^M WC_{j,k} - PWR_k \right)^+ \right] \\ &= C_h PWR_k - C_h \sum_{j=1}^M \mu_{j,k} \\ &\quad + (C_h + C_b) \mathbb{E} \left[\sum_{j=1}^M WC_{j,k} - PWR_k \right]^+. \end{aligned} \quad (15)$$

Based on the result in Lemma 2, we can obtain the following proposition.

Proposition 5. *The distributionally robust optimal level of combined warranty reserves for the k th planning period is given by*

$$PWR_k^* = \sum_{j=1}^M \mu_{j,k} + \frac{A}{2} \sqrt{\sum_{j=1}^M \sigma_{j,k}^2}. \quad (16)$$

The corresponding expected total reserve losses for the pooled case can be obtained as

$$\hat{\pi}_k^c(PWR_k^*) = \frac{B}{2} \sqrt{\sum_{j=1}^M \sigma_{j,k}^2}. \quad (17)$$

5.2. Benefits of Funds Pooling

We investigate the benefits of funds pooling via four efficiency measures. In particular, we are interested in the time-varying properties of pooling benefits over the entire warranty life cycle. The first two measures are related to the levels of warranty reserves, that is, absolute benefit Υ_k^* and relative benefit Ψ_k^* ($k \in \{1, 2, \dots, n + m\}$), which can be expressed as

$$\Upsilon_k^* = TWR_k^* - PWR_k^* = \frac{A}{2} \left(\sum_{j=1}^M \sigma_{j,k} - \sqrt{\sum_{j=1}^M \sigma_{j,k}^2} \right), \quad (18)$$

and

$$\Psi_k^* = \frac{TWR_k^* - PWR_k^*}{TWR_k^*} = \frac{\frac{A}{2} \left(\sum_{j=1}^M \sigma_{j,k} - \sqrt{\sum_{j=1}^M \sigma_{j,k}^2} \right)}{\sum_{j=1}^M \mu_{j,k} + \frac{A}{2} \sum_{j=1}^M \sigma_{j,k}}. \quad (19)$$

We also examine the absolute benefit Υ_k^o and relative benefit Ψ_k^o in terms of expected total reserve losses:

$$\begin{aligned} \Upsilon_k^o &= \sum_{j=1}^M \hat{\pi}_{j,k}^*(WR_{j,k}^*) - \hat{\pi}_k^{c*}(PWR_k^*) \\ &= \frac{B}{2} \left(\sum_{j=1}^M \sigma_{j,k} - \sqrt{\sum_{j=1}^M \sigma_{j,k}^2} \right), \end{aligned} \quad (20)$$

and

$$\begin{aligned} \Psi_k^o &= \frac{\sum_{j=1}^M \hat{\pi}_{j,k}^*(WR_{j,k}^*) - \hat{\pi}_k^{c*}(PWR_k^*)}{\sum_{j=1}^M \hat{\pi}_{j,k}^*(WR_{j,k}^*)} \\ &= 1 - \frac{\sqrt{\sum_{j=1}^M \sigma_{j,k}^2}}{\sum_{j=1}^M \sigma_{j,k}}. \end{aligned} \quad (21)$$

It is obvious that all the four efficiency measures are positive, even though the products are released at different dates (i.e., Assumption 4 does not hold). This implies that funds pooling is indeed beneficial. In terms of time-varying behaviors, it is straightforward to verify that absolute benefits Υ_k^* and Υ_k^o exhibit an upside-down bathtub shape (refer to Section EC.2.1 in the online supplement for details). In practice, the relative benefits are of more interest to manufacturing firms. In what follows, we focus only on the time-varying properties of relative benefits Ψ_k^* and Ψ_k^o .

Define $\sigma_{\max,k} = \max_{j=1,\dots,M} \{\sigma_{j,k}\}$ and $\sigma_{\min,k} = \min_{j=1,\dots,M} \{\sigma_{j,k}\}$. The following lemma examines the monotonicity of relative benefit Ψ_k^* with respect to the mean and standard deviation of the AWC increment.

Lemma 3. *The relative benefit Ψ_k^* decreases as $\mu_{j,k}$ increases and as $\sigma_{\min,k}$ decreases.*

Lemma 3 presents an interesting finding that somewhat explains the relationship between the pooling benefit Ψ_k^* and the range of standard deviations of different products. If we decrease the value of $\sigma_{\min,k}$ while

keeping other standard deviations unchanged—that is, the left end-point of interval $[\sigma_{\min,k}, \sigma_{\max,k}]$ becomes smaller (the range is partially enlarged)—then the relative benefit Ψ_k^* becomes smaller. It is difficult to obtain the general monotonicity of Ψ_k^* with respect to k , because of its complex formulation. Nevertheless, when the AWC increments of all products share the same variance, we have the following result.

Corollary 1. *If the AWC increments of the M products share the same variance, then Ψ_k^* would follow the same time-varying monotonicity of $CV_{j,k}$ ($j \in \{1, 2, \dots, M\}$) over the warranty life cycle.*

Corollary 1 implies that Ψ_k^* would present a bathtub shape that is consistent with the pattern of coefficients of variation (Figures 9(a) and 10(a)). This supports the use of coefficient of variation to describe AWC variability, as discussed earlier. It is worth pointing out that when the variances of AWC increments for different products are identical, their means are not necessarily the same.⁹

Next, we investigate the time-varying property of relative benefit Ψ_k^o in terms of reserve losses. Define $\delta_{\max,k} := \sigma_{\max,k}/\sigma_{\min,k}$ as the “relative range” of the standard deviations. Then, we have the following key result.

Proposition 6. *The relative benefit Ψ_k^o decreases as the ratio $\delta_{\max,k}$ increases.*


Proposition 6 shows that when all other standard deviations (i.e., $\sigma_{j,k} \neq \sigma_{\min,k}, \sigma_{\max,k}$) are fixed, enlarging the relative range of AWC standard deviations will weaken the pooling benefit. This phenomenon can be explained by the so-called *resonance* effect. That is, when the magnitudes of AWC variabilities of different products become larger, the manufacturer must allocate more reserves to those with higher variabilities, which clearly reduces the relative benefit of funds pooling. This result shows that the relative range $\delta_{\max,k}$ plays a critical role in comparing relative benefits Ψ_k^o in different planning periods. By defining $\Gamma_k^{(j,l)} := \sigma_{j,k}/\sigma_{l,k}$, $\forall j \neq l$, the following lemma is useful for investigating the time-varying behavior of the pooling benefit.

Lemma 4. *If $\Gamma_1^{(j,l)} > 1$ and $\theta_j > \theta_l$ hold in the first planning period, then (1) $\Gamma_k^{(j,l)} > 1$ holds for $k \in \{2, 3, \dots, n + m\}$, and (2) $\Gamma_k^{(j,l)}$ is a nondecreasing function of k .*

Combining Proposition 6 and Lemma 4 (notice that $\delta_{\max,k} = \Gamma_k^{(\max,\min)}$), the time-varying characteristic of Ψ_k^o can be captured by the following corollary.

Corollary 2. *Define $\theta_{\max}$ and $\theta_{\min}$ as the failure rates associated with $\sigma_{\max,1}$ and $\sigma_{\min,1}$, respectively. For $M = 2$, if $\theta_{\max} > \theta_{\min}$, then Ψ_k^o is a nonincreasing function of $k \in \{1, 2, \dots, n + m\}$.*

This finding is quite insightful. Although the relative benefit Ψ_k^o is also varying over time, its tendency


Figure 9. (Color online) Relative Benefits Ψ_k^* and Ψ_k^o with Respect to k When Pooling Two of the Three Products

differs from that of Ψ_k^* . Figure 9 shows the relative benefits Ψ_k^o when pooling any two of the three products, with the parameter settings in Figures 3 and 4. It can be observed that Ψ_k^o shows an opposite tendency to $\delta_{\max,k}$. Moreover, in a specific period k , a larger (respectively, smaller) $\delta_{\max,k}$ corresponds to a smaller (respectively, larger) Ψ_k^o . For example, the relative benefit Ψ_k^* of pooling warranty reserves of products 1 and 2 (black circle in Figure 9) is the largest one among the three cases, because the associated ratio $\delta_{\max,k}$ is the smallest. These observations are consistent with the analytical results in Proposition 6 and Corollary 2.

It is important to note that Corollary 2 strictly holds only for $M = 2$; that is, there are only two products. This is because only when $M = 2$, Ψ_k^o decreases exclusively with $\delta_{\max,k}$, can the chain rule be applied to

obtain Corollary 2. For $M > 2$, it is difficult, if not impossible, to obtain the analytical relationship between Ψ_k^o and k . Nevertheless, Figure 10 illustrates relative benefits Ψ_k^* and Ψ_k^o when the warranty reserves of all three products are pooled. It shows that Ψ_k^* has a bathtub shape and Ψ_k^o exhibits a nonincreasing pattern over the entire warranty life cycle. Moreover, it is noteworthy that when pooling three products, both Ψ_k^* and Ψ_k^o are higher than those in the case of pooling any two products. This demonstrates that pooling benefits would be more significant when the warranty reserves of more products are pooled.

Another interesting observation from Figures 9 and 10 is that the relative benefit Ψ_k^o varies within a small range. In particular, Ψ_k^o is almost constant in some cases. This is because when sales rate λ and

Figure 10. Relative Benefits Ψ_k^* and Ψ_k^o with Respect to k When Pooling the Three Products

replacement cost C_W are large and failure rate θ is small (which is quite common for electronics), we know from Equation (4) that the ratio $\Gamma_k^{(j,l)} = \sigma_{j,k}/\sigma_{l,k}$ between any two products j and l would be almost constant over time, which in turn leads to an almost constant Ψ_k^o .

The previous discussions of time-varying pooling benefits rely heavily on the simultaneous release assumption (i.e., Assumption 4). In practice, however, manufacturing firms often introduce their products in a one-by-one or batch-by-batch manner. We thus extend our discussion to the case of periodic release (which is common for consumer electronics); please refer to Section EC.2.2 in the online supplement. In addition, we provide a brief discussion in Section EC.2.3 of the online supplement on optimal release scheduling from a reserve management perspective.

6. Case Study with Field Data

In this section, we demonstrate the performance of the proposed reserve planning model and the reserve-loss-reduction strategies through a real warranty data set provided by the aforementioned electronics manufacturer. The firm manufactures various kinds of electronic products and components. We select the sales data and associated warranty claims data of three products as examples (data are scaled for confidentiality). The warranty periods of the three products are identical ($W = 12$ months), whereas their release dates and sales periods are different (Table 1). Monthly sales data are plotted in Figure 11(a), which shows that the sales volumes fluctuate over time but do not exhibit obvious patterns. We apply the Augmented Dickey-Fuller test¹⁰ to examine the sales data and find that we are at least 70% confident that the data are stationary. Thus, it is “safe” to adopt the HPP sales assumption in our problem. In particular, the monthly sales rates are estimated as 61,316, 48,187, and 59,103, respectively. Moreover, the associated monthly warranty claims are shown in Figure 1. In the case study, warranty period W and sales period L are directly obtained from the firm and sales rate λ is fitted from the sales data, whereas failure rate θ is the only parameter that must be estimated from the claims data. For this purpose, we adopt the following least squares method:

$$\hat{\theta} = \arg \min \sum_{k=1}^{n+m} (\mu_k - \Delta Q_k)^2,$$


where μ_k is given by Equation (3) and ΔQ_k is observed warranty expenses over month k . By solving this optimization problem, the estimated monthly failure rates are summarized in Table 1. Figure 11(b) further plots actual and estimated monthly warranty expenses. One can observe that, although the stylized AWC model can roughly capture real warranty claim patterns, the curve-fitting performance presents certain deficiency; that is, a systematic error exists. More specifically, the fitted curves tend to overestimate the actual data at early and late stages while underestimating the actual data at the middle stage.

Based on the product sales, reliability, and cost information in Table 1, the mean $\mu_{j,k}$ and variance $\sigma_{j,k}^2$ of the AWC increment for each product j in each planning month k can be computed by Equations (3) and (4), respectively. Then, the distributionally robust optimal reserve level $WR_{j,k}^*$ is determined by Equation (9). Finally, the corresponding actual warranty reserve loss can be calculated through $\pi_{j,k}(WR_{j,k}^*) = C_h(WR_{j,k}^* - \Delta Q_{j,k})^+ + C_b(\Delta Q_{j,k} - WR_{j,k}^*)^+$. Figure 12 shows actual warranty costs and planned reserve levels for the three products. It can be observed that for each product, the planned reserves are excessive at first, insufficient in the middle, and over-stocked again at the last, which results from the aforementioned systematic error in curve fitting.

The unsatisfactory performance of the model-driven reserving policy necessitates demand learning. In this scenario, the adjusted warranty reserve levels $\widehat{WR}_{j,k}$ for the three products are dynamically updated by Equation (11), which are depicted in Figure 12. We can observe that the reserve levels with learning are much closer to actual warranty costs than those without learning. Moreover, funds pooling is also applied to improve warranty reserve planning performance. By pooling warranty reserves of the three products (with different release dates and sales periods), the combined reserve levels and corresponding reserve losses are evaluated by Equations (16) and (17), respectively (see Figure 12(d) for pooled reserve levels). Not surprisingly, the planned reserve level in the pooled case is closer to actual combined warranty expenses due to the risk pooling effect. One thing noteworthy is that it is feasible to apply demand learning to the pooled case, but the funds adjusted by learning cannot be used to evaluate pooling benefits.

Table 1. Product Information Used in the Case Study

Product	Sales period	W	L	$\hat{\lambda}$	$\hat{\theta}$	C_W	C_h	C_b
1	September 2009 to August 2011	12	24	61,316	0.00126	\$100	\$0.01	\$0.025
2	September 2010 to February 2012	12	18	48,187	0.00794	\$60	\$0.01	\$0.025
3	May 2010 to April 2011	12	12	59,103	0.00541	\$45	\$0.01	\$0.025

Figure 11. (Color online) Monthly Sales Volumes (a) and Aggregate Warranty Costs (b) of the Three Products

To compare the scenarios with and without learning, Table 2 presents total warranty reserve levels and total reserve losses for each product over its entire warranty life cycle. Compared with the no-learning scenario, the total reserve levels and total reserve losses in the learning scenario are indeed lower. In particular, the reductions in total reserve losses are significant in terms of percentage (27.64%–62.76%). This shows that adopting the demand learning mechanism in warranty reserve planning can save a large proportion of money for the firm, either in the pooled or nonpooled case. Moreover, when examining the pooling benefit in the learning scenario, we find that the pooled scheme results in even a slightly higher reserve level (increased by \$48,979, 0.55%) than the nonpooled case. However, the corresponding total reserve losses are reduced significantly (decreased by \$8496, 47.25%). Nevertheless, a higher total reserve level in the pooled case is not surprising. This is because, in the nonpooled cases, the reserve levels in

some periods are insufficient, which results in large shortage expenses. In contrast, the pooled case puts more reserves in these periods to mitigate the stock-out effect. This further demonstrates the effectiveness of demand learning and funds pooling. Furthermore, the magnitudes of reductions in total reserve losses due to learning and pooling are generally higher than those of total reserve levels, which implies that the two reserve-loss-reduction approaches have more significant impacts on reserve losses than reserve levels.

7. Conclusions and Implications

The goal of this paper is to investigate a warranty reserve planning problem during the entire warranty life cycle from a manufacturing firm's perspective. For this purpose, we first develop an adaptive AWC model by fully characterizing the mechanism that governs the generation of warranty claims, which allows the firm to plan for warranty reserves periodically.


Figure 12. (Color online) Actual Monthly Warranty Costs and Planned Reserve Levels (Without and with Learning) for the Three Products and the Pooled Case

Table 2. Reserve Levels and Reserve Losses over the Warranty Life Cycle in Learning and No-Learning Scenarios

	Total reserve levels (\$)			Pooled vs. nonpooled		
	Product 1	Product 2	Product 3	Pooled case	Reduction	Percentage of reduction
No-learning case	2,265,826	5,003,016	2,093,834	9,327,879	34,779	0.37%
Learning case	2,097,786	4,787,992	2,029,187	8,963,944	−48,979	−0.55%
Reduction	168,039	215,024	64,647	363,952		
Percentage of reduction	7.42%	4.30%	3.09%	3.90%		

	Total reserve losses (\$)			Pooled vs. nonpooled		
	Product 1	Product 2	Product 3	Pooled case	Reduction	Percentage of reduction
No-learning case	12,506	19,569	5,258	14,629	22,704	60.82%
Learning case	4,657	9,520	3,805	9,485	8,496	47.25%
Reduction	7,849	10,049	1,453	5,144		
Percentage of reduction	62.76%	51.35%	27.64%	35.16%		

Because the exact distribution of AWC increments is generally unknown, the optimal warranty reserve level in each planning period is determined via a distributionally robust approach. On this basis, two reserve-loss-reduction strategies—demand learning and funds pooling—are further discussed. The demand learning algorithm can dynamically update the reserving plan by combining real warranty data, whereas the funds pooling scheme combines reserve funds of multiple products into a single fund to mitigate risks. Analytical results and simulation experiments show that demand learning and funds pooling are indeed beneficial for the firm in terms of reducing total reserve losses.

The main conclusions and implications of this work are summarized as follows:

(i). For the exponential-smoothing-type demand learning algorithm, an important finding is that when the sales period is long enough, the learning parameter will converge to a constant in probability, with a convergence rate of $1/\sqrt{k}$. This not only guarantees a stable output of the algorithm, but also simplifies the algorithm implementation. In addition, simulation experiments show that by incorporating the objective function in the learning process, the algorithm can significantly reduce total reserve losses under general sales and failure patterns, for which the proposed AWC model fails to capture the dynamics of warranty costs. We believe that this simple yet effective “predict-then-optimize” approach would be also useful for other operations management applications.

(ii). For the funds pooling scheme, we find that the pooling benefits change over different stages of the warranty life cycle, due to the dynamics of warranty expenses. More importantly, because reserve loss reduction is the primary objective of funds pooling, we prove an important result that the relative pooling benefit in terms of reserve losses exhibits a nonincreasing pattern (note that this result analytically holds only when there are two products). Another compelling insight is that the relative pooling benefit in terms of reserve losses

decreases as the relative range of demand standard deviations enlarges, meaning that it is more attractive to pool warranty reserves for products whose AWC uncertainties are of similar magnitudes. This insight is quite general in the sense that it is applicable to traditional risk pooling problems in operations management.

(iii). Finally, our simulation experiments provide guidelines on the implementation of the proposed methodologies. Specifically, when the product sales and failure processes governing warranty claims generation obey those described in Assumptions 1 and 2, then implementing the distributionally robust reserving model, combined with funds pooling, can deliver a fairly good reserving plan; if this is not the case, then the demand learning mechanism can be further used to compensate the model’s inefficiency by leveraging real claims data.

However, this paper presents several limitations that deserve future research efforts. First, because of the simplified assumptions on product sales and failure processes, the proposed AWC model is somewhat stylized, although it is still capable to roughly capture any inverted bathtub-like warranty claim patterns. Because deriving the mean and variance of warranty expenses in general settings is highly intractable, a promising future research topic is to develop an approximate, closed-form expression of the mean and variance to facilitate warranty reserve planning, without inducing much computational burden. Moreover, an implicit assumption adopted here is that the sales of different products are not correlated. In reality, however, products could be substitutes or complements. It is interesting to investigate the impact of such correlations on warranty reserve planning. One possible way is to apply diffusion-choice demand models (Li 2020) that integrate the Bass diffusion model and discrete choice models. Furthermore, only a specific form of warranty policy, that is, free replacement warranty, is considered in this work. Future research could incorporate more warranty policies, for example, pro rata and combination warranties, into the reserve planning framework.

Acknowledgments

The authors thank the department editor, associate editor, and two anonymous referees for constructive comments and suggestions that improved the paper significantly. The authors also thank Chung-Piaw Teo and Hanlin Liu for valuable discussions on an original version of this paper.

Endnotes

¹ Apple's warranties and service contracts: <http://www.warrantyweek.com/archive/www20171109.html>.

² Warranty reserves: <http://www.warrantyweek.com/archive/ww20060509.html>.

³ Warranty adjustments: <http://www.warrantyweek.com/archive/ww20151105.html>.

⁴ One should distinguish *product life cycle* from *warranty life cycle*. The former starts from the time a product is first introduced to the market and extends the instant it is discontinued from the market, while the latter is the period from the beginning of a warranty program to the end of the program (when the warranties for final sales expire). In general, let L denote product life cycle and W denote warranty period; then the warranty life cycle should be $L + W$.

⁵ Once all units of a product are sold to end consumers, the *installed base* (i.e., the total number of sold units that are currently in use) will remain almost constant for a relatively long period of time. As a result, the demand of parts (resulting from repair/replacement) gradually turns to be stationary, as the units have statistically identical quality and reliability.

⁶ Our simulation studies reveal that the approximation accuracy is relatively insensitive to λ but exhibits an obvious increasing trend as θ decreases. When failure rate θ is small enough (say, $\theta < 0.1$), the approximation error vanishes.

⁷ Suppose that the warranty reserve fund accrues interest at a constant rate ζ_0 , the opportunistic investment return rate is ζ_1 , and the interest rate for obtaining emergency funds is ζ_2 (in general, $\zeta_0 < \zeta_1 < \zeta_2$). Then, C_h and C_b can be roughly estimated by $C_h = \zeta_1 - \zeta_0$ and $C_b = \zeta_2 - \zeta_0$, respectively. Nevertheless, what matters to warranty reserve planning, as can be seen from Equation (9), is not the absolute values of C_h and C_b , but their relative value C_b/C_h or C_h/C_b .

⁸ Based on our consultation with some firms, including a leading engine manufacturer, an electronics manufacturer, and a credit card company, all of them do not implement warranty reserve pooling for their products.

⁹ For example, based on Equations (3) and (4), given two products j and l , combining the conditions $\theta_j = \theta_l$, $C_{Wj}^2\lambda_j = C_{Wl}^2\lambda_l$, and $C_{Wj}\lambda_j \neq C_{Wl}\lambda_l$ can guarantee that $\sigma_{j,k} = \sigma_{l,k}$ and $\mu_{j,k} \neq \mu_{l,k}$ for any $k \in \{1, 2, \dots, n + m\}$. The following parameter setting satisfies the previous three conditions: $\theta_j = \theta_l = 0.05$, $C_{Wj} = 30$, $C_{Wl} = \sqrt{150}$, $\lambda_j = 500$, and $\lambda_l = 3,000$.

¹⁰ The augmented Dickey-Fuller test is a popular method that assesses the presence of a unit root in an autoregressive model $y_t = \rho y_{t-1} + \sum_{j=1}^p \nu_j \Delta y_{t-j} + \epsilon_t$, where $\Delta y_t = y_t - y_{t-1}$, p is the lag order, and ϵ_t is a white noise. In this model, testing for a unit root is equivalent to testing $\rho = 1$, which implies that the time series y_t is nonstationary. The test statistic is $(\hat{\rho} - 1)/\widehat{Std}(\hat{\rho})$, where $\hat{\rho}$ is the ordinary least squares estimator of ρ and $\widehat{Std}(\hat{\rho})$ represents an estimator of the standard deviation of $\hat{\rho}$. If the null hypothesis $H_0: \rho = 1$ is rejected at the significance level α , then we are $100(1 - \alpha)\%$ confident that the time series sample is stationary. For more details, please refer to Fuller (1996, chapter 10).

References

- Bass FM (1969) A new product growth for model consumer durables. *Management Sci.* 15(5):215–227.
- Baumol WJ (1952) The transactions demand for cash: An inventory theoretic approach. *Quart. J. Econom.* 66(4):545–556.
- Benjaafar S, Cooper WL, Kim JS (2005) On the benefits of pooling in production-inventory systems. *Management Sci.* 51(4):548–565.
- Berman O, Krass D, Mahdi Tajbakhsh M (2011) On the benefits of risk pooling in inventory management. *Production Oper. Management* 20(1):57–71.
- Besbes O, Mouchtiki O (2021) How big should your data really be? Data-driven newsvendor and the transient of learning. Working paper, Columbia Business School, New York.
- Bimpikis K, Markakis MG (2016) Inventory pooling under heavy-tailed demand. *Management Sci.* 62(6):1800–1813.
- Blischke WR, Murthy DNP (2000) *Reliability: Modeling, Prediction, and Optimization* (John Wiley & Sons, New York).
- Bolton GE, Katok E (2008) Learning by doing in the newsvendor problem: A laboratory investigation of the role of experience and feedback. *Manufacturing Service Oper. Management* 10(3): 519–538.
- Bostian AJA, Holt CA, Smith AM (2008) Newsvendor “pull-to-center” effect: Adaptive learning in a laboratory experiment. *Manufacturing Service Oper. Management* 10(4):590–608.
- Buczowski PS, Kulkarni VG (2006) Funding a warranty reserve with contributions after each sale. *Probab. Engrg. Inform. Sci.* 20(3):497–515.
- Calmon AP, Graves SC (2017) Inventory management in a consumer electronics closed-loop supply chain. *Manufacturing Service Oper. Management* 19(4):568–585.
- Calmon AP, Graves SC, Lemmens S (2021) Warranty matching in a consumer electronics closed-loop supply chain. *Manufacturing Service Oper. Management* 23(5):1314–1331.
- Chen B, Chao X (2019) Parametric demand learning with limited price explorations in a backlog stochastic inventory system. *IIE Trans.* 51(6):605–613.
- Chen X, Simchi-Levi D (2009) A new approach for the stochastic cash balance problem with fixed costs. *Probab. Engrg. Inform. Sci.* 23(4):545–562.
- Cohen D, Darrough MN, Huang R, Zach T (2011) Warranty reserve: Contingent liability, information signal, or earnings management tool? *Accounting Rev.* 86(2):569–604.
- Eliashberg J, Singpurwalla ND, Wilson SP (1997) Calculating the reserve for a time and usage indexed warranty. *Management Sci.* 43(7):966–975.
- Eppen GD (1979) Note—Effects of centralization on expected costs in a multi-location newsboy problem. *Management Sci.* 25(5): 498–501.
- Eppen GD, Fama EF (1969) Cash balance and simple dynamic portfolio problems with proportional costs. *Internat. Econom. Rev.* 10(2):119–133.
- Fuller WA (1996) *Introduction to Statistical Time Series*, 2nd ed. (John Wiley & Sons, New York).
- Gallego G, Moon I (1993) The distribution free newsboy problem: Review and extensions. *J. Oper. Res. Soc.* 44(8):825–834.
- Gallego G, Wang R, Hu M, Ward J, Beltran JL (2015) No claim? Your gain: Design of residual value extended warranties under risk aversion and strategic claim behavior. *Manufacturing Service Oper. Management* 17(1):87–100.
- Gormley FM, Meade N (2007) The utility of cash flow forecasts in the management of corporate cash balances. *Eur. J. Oper. Res.* 182(2):923–935.
- Govindarajan A, Sinha A, Uichanco J (2021) Distribution-free inventory risk pooling in a multilocation newsvendor. *Management Sci.* 67(4):2272–2291.

- Gurgur CZ (2011) Dynamic cash management of warranty reserves. *Engrg. Econom.* 56(1):1–27.
- Huber J, Müller S, Fleischmann M, Stuckenschmidt H (2019) A data-driven newsvendor problem: From data to decision. *Eur. J. Oper. Res.* 278(3):904–915.
- Ja SS, Kulkarni VG, Mitra A, Patankar JG (2002) Warranty reserves for nonstationary sales processes. *Naval Res. Logist.* 49(5):499–513.
- Karsten F, Slikker M, Van Houtum GJ (2015) Resource pooling and cost allocation among independent service providers. *Oper. Res.* 63(2):476–488.
- Keskin NB, Min X, Song JSJ (2021) The nonstationary newsvendor: Data-driven nonparametric learning. Working paper, Fuqua School of Business, Duke University, Durham, NC.
- Khawam J, Hausman WH, Cheng DW (2007) Warranty inventory optimization for Hitachi Global Storage Technologies, Inc. *Interfaces* 37(5):455–471.
- Li H (2020) Optimal pricing under diffusion-choice models. *Oper. Res.* 68(1):115–133.
- Menke WW (1969) Determination of warranty reserves. *Management Sci.* 15(10):542.
- Miller MH, Orr D (1966) A model of the demand for money by firms. *Quart. J. Econom.* 80(3):413–435.
- Mišić VV, Perakis G (2020) Data analytics in operations management: A review. *Manufacturing Service Oper. Management* 22(1):158–169.
- Murthy DNP, Solem O, Roren T (2004) Product warranty logistics: Issues and challenges. *Eur. J. Oper. Res.* 156(1):110–126.
- Neave EH (1970) The stochastic cash balance problem with fixed costs for increases and decreases. *Management Sci.* 16(7):472–490.
- Patankar JG, Worm GH (1981) Note: Prediction intervals for warranty reserves and cash flows. *Management Sci.* 27(2):237–241.
- Pinçe Ç, Ferguson M, Toktay B (2016) Extracting maximum value from consumer returns: Allocating between remarketing and refurbishing for warranty claims. *Manufacturing Service Oper. Management* 18(4):475–492.
- Scarf H (1958) A min max solution of an inventory problem. Arrow K, Karlin S, Scarf H, eds. *Studies in the Mathematical Theory of Inventory and Production* (Stanford University Press, Stanford, CA), 201–209.
- Tagaras G, Cohen MA (1992) Pooling in two-location inventory systems with non-negligible replenishment lead times. *Management Sci.* 38(8):1067–1083.
- Tapiero CS, Posner MJ (1988) Warranty reserving. *Naval Res. Logist.* 35(4):473–479.
- Thomas MU (1989) Note: A prediction model for manufacturer warranty reserves. *Management Sci.* 35(12):1515–1519.
- Xie W, Ye ZS (2016) Aggregate discounted warranty cost forecast for a new product considering stochastic sales. *IEEE Trans. Reliability* 65(1):486–497.
- Yang C, Hu Z, Zhou SX (2021) Multilocation newsvendor problem: Centralization and inventory pooling. *Management Sci.* 67(1):185–200.
- Zhang H, Chao X, Shi C (2020) Closing the gap: A learning algorithm for lost-sales inventory systems with lead times. *Management Sci.* 66(5):1962–1980.
- Zhong Y, Zheng Z, Chou MC, Teo CP (2018) Resource pooling and allocation policies to deliver differentiated service. *Management Sci.* 64(4):1555–1573.