

Managing Rentals with Usage-Based Loss

Vincent W. Slauch

Tepper School of Business, Carnegie Mellon University

Bahar Biller

General Electric, Global Research Center

Sridhar R. Tayur

Tepper School of Business, Carnegie Mellon University

Draft: January 7, 2015

Motivated by new and innovative rental business models, this paper develops a novel discrete-time model of a rental operation with random loss of inventory due to customer use. The inventory level is chosen before the start of a finite rental season, and customers not immediately served are lost. Our analysis framework uses stochastic comparisons of sample paths to derive structural results that hold under good generality for demands, rental durations, and rental unit lifetimes. Considering different “recirculation” rules — i.e., which rental unit to choose to meet each demand — we prove the concavity of the expected profit function and identify the optimal recirculation rule. A numerical study clarifies when considering rental unit loss and recirculation rules matters most for the inventory decision: Accounting for rental unit loss can increase the expected profit by 7% for a single season and becomes even more important as the time horizon lengthens. We also observe that the optimal inventory level in response to increasing loss probability is non-monotonic. Finally, we show that choosing the optimal recirculation rule over another simple policy allows more rental units to be profitably added, and the profit-maximizing service level increases by up to 6 percentage points.

Key words: Service Operations; Capacity Planning and Investment; Inventory Theory and Control; Supply Chain Management; Stochastic Methods

1. Introduction

Advances in online commercial models have produced a new generation of innovative businesses built upon renting goods. The flexibility and affordability promised by renting a wide array of products have led to rental businesses specializing in just about every aspect of our business and

personal lives. Besides the traditional rental products such as movies, cars, and hotel rooms, less common goods available to rent range from bicycles to jets, cribs to coffins, and furniture to camping gear. According to IBISWorld industry analysts, the annual revenue of fifteen different rental industries in the United States each exceeded \$1 billion in 2013, while the annual revenue of each of the car, heavy equipment, and industrial equipment rental industries surpassed \$25 billion.

Luxury goods have received particular attention as fertile ground for rental businesses that make those goods available to new customer classes. For example, Rent the Runway is a company that allows customers to rent high-fashion dresses for either four or eight days at approximately 10% of the retail price of a dress (Wortham 2009). Customers can view the selection of dresses and their availability through a website, and receive style and fit advice from Rent the Runway consultants and customer reviews. Dresses are shipped to customers and returned by mail. However, the critical decision about the number of dresses that will comprise Rent the Runway's seasonal rental inventory must be made shortly after pre-season fashion shows, which are several months in advance of the rental season (Binkley 2011).

Choosing the number of rental units to procure before the start of a rental season without the possibility of replenishment during the season is an important problem that many rental businesses face. Despite the seemingly fundamental nature of this problem, operations management literature offers very little analytical support when lost sales and discrete time periods — natural assumptions for many rental systems — are considered. In this paper, we analyze a single-product rental system using a discrete-time framework. We focus on the usage-based loss of rental units over a finite rental horizon during which no additional rental units may be ordered, e.g., when long procurement lead times prohibit in-season reordering. In particular, we consider each rental unit to have a random lifetime, which is characterized by a general probability distribution on the number of times the unit can be rented before its retirement from the rental inventory. Our goal is to understand the role of this uncertainty arising from the usage-based loss of rental units on the management of rental inventory.

In addition to Rent the Runway, whose dresses are susceptible to both destructive incidents and wearing out over time, other rental systems face the challenge of losing inventory that can be difficult to replace in the middle of the rental season. For example, a Paris-based bicycle sharing program that began with 20,600 bicycles in 2007 had more than 8,000 bikes stolen and another 8,000 bikes severely damaged and in need of replacement within two years (Erlanger and De La Baume 2009). Inventory loss can also occur when customers exercise an option to purchase a product. Users of Redbox, an automated movie and game rental kiosk, rent a DVD for \$1.20 a day. If the DVD is not returned in 20 days, then the customer pays \$24 for the accrued daily rental charge and owns the DVD. Another example is Rent-A-Center, a company with over \$3 billion in revenue in 2012 and which rents furniture, appliances and electronics to customers who can own the item if it is rented beyond a certain duration. In its 2012 annual report, Rent-A-Center states that approximately 25% of its rental agreements result in customer ownership.

Existing work supporting capacity planning for rental businesses relies primarily on queueing models. Although Poisson or compound Poisson arrival processes may adequately represent demands for some rental businesses, better choices may exist for modeling demand in rental systems characterized by discrete rental time slots. For example, business travelers occupy a hotel room for a discrete number of days and are more likely to begin renting a hotel room on Monday night than a Saturday night. At Rent the Runway, for example, whose customers primarily rent dresses for events on Fridays and Saturdays, a discrete-time demand model with a period of one week more accurately represents a customer demand pattern than a Poisson arrival process. Therefore, extending the discrete-time inventory theory to include loss of rental inventory offers an advantage for a rental system like Rent the Runway. We develop a model that makes no distributional assumptions and captures (a) operational details such as random rental unit lifetimes (with constant, increasing or decreasing failure rates) and random rental durations, (b) very general demand models with features such as seasonality or auto-correlations and (c) recirculation rules that are used in practice for choosing among available rental units to satisfy demands.

We make the following contributions regarding the inventory management of rental systems:

1. **Model and Framework.** To the best of our knowledge, we are the first to consider the loss of rental units according to distributions over the number of times that each unit can be rented before loss. Thus, our model includes a state variable that represents the number of times that a rental unit has been rented out (i.e., a “count-based” model) or a state variable that represents a rental unit’s condition (i.e., a “condition-based” model). It also accommodates an arbitrary demand process and general distributions for lifetime and duration of each rental unit.

2. Structural Results:

(a) We establish the concavity of the expected profit function in the initial inventory of rental units for geometric lifetime distributions. Not surprisingly, this structural property holds independent of the rental unit recirculation rules as the loss probability is constant over time.

(b) For general lifetime distributions, it becomes necessary to consider the rules that allocate rental units to satisfy customer demand for both count-based or condition-based models.

(c) **Count-Based Model:** We establish the concavity of the expected profit function for the “static priority” recirculation rule; i.e., the units to be rented are prioritized according to a list that does not change over the rental horizon. We show that the concavity of the expected profit function also holds for a policy that spreads the rental load evenly over all units, allocating the rental unit that has been rented out the fewest number of times. Referring to this recirculation rule as the “even spread” policy, we prove its optimality when rental unit loss probabilities are non-decreasing in the number of times that the unit has been rented.

(d) **Condition-Based Model.** We demonstrate analogous results for the condition-based model, showing the concavity of the expected profit function for the “best-first” policy in which the rental unit in the best condition receives the highest allocation priority. Also, we prove that the best-first policy is optimal when the state transition probability matrix is totally positive of order 2, a condition that implies that the rental unit failure rate is increasing as its condition worsens.

3. Managerial Insights from Numerical Study

(a) Failing to account for usage-based loss of rental inventory leads to a significant reduction in the expected profit. For a 5% probability of loss each time a unit is rented, we find that ignoring

the loss of rental units reduces the expected profit by 7.3% and 33.0% for a half-year and a full-year rental horizon, respectively.

(b) The optimal response to the increasing loss probability is to first increase the number of rental units, then decrease the number of rental units and finally stock zero rental units.

(c) For a rental unit lifetime distribution with increasing loss probability, the rental unit recirculation rule plays an important role according to the rate at which the loss probability increases. We focus on the count-based model, as similar results apply for the condition-based model, and compare the even spread policy to the static priority recirculation rule. Choosing the even spread policy increases the optimal initial inventory level with a corresponding increase of up to 6 percentage points in the profit-maximizing service level.

The remainder of the paper is organized as follows. Section 2 reviews the rental inventory management literature. Section 3 introduces our rental inventory model. We establish the structural properties of this model for geometric lifetime distributions in Section 4 and for general lifetime distributions in Section 5, where we further identify the optimal rental unit recirculation rule under certain conditions. The numerical analysis follows in Section 6. We conclude with a summary of findings and future research directions in Section 7.

2. Literature Review

Early research on rental inventory management exclusively uses queueing models as a foundation for analysis. The initial advances in queueing theory by Takács (1962) and Riordan (1962) for the telephone trunking problem — finding the stationary probabilities of a multi-server pure loss system — have sparked two seminal papers on the problem of sizing a fleet of rental equipment. Tainiter (1964) formulates an optimization problem for $M/G/c/c$ and $G/M/c/c$ rental systems based on the limiting distributions of the system states derived by Takács (1962). The decision variable is the capacity of the rental system and the problem is studied both asymptotically and over a finite horizon. Whisler (1967), on the other hand, shows that the optimal policy structure for a rental system with lost sales, periodic reordering, and nonstationary state transition probabilities

— as in Riordan (1962) — has upper critical values above which inventory should be discarded and lower critical values below which inventory should be ordered. Our work differs from these studies by its focus on the inventory decision prior to the rental season, the challenge of handling random usage-based loss of rental units and stochastic rental duration, and the use of a discrete-time model for demand representation.

The early research on rental inventory management with lost sales is followed by an extensive study of the $M/M/c$ queueing model with backlogged demands. Specifically, the problem is posed as finding the optimal number of servers to employ in a multi-server queueing system, where servers represent rental units and service time corresponds to the rental duration; see Huang et al. (1977), Jung and Lee (1989), Green et al. (2001), and Zhang et al. (2012). Motivated by the time-specific nature of customers' rentals, however, we restrict our focus to lost sales models in this paper. Table 1 compares our rental inventory model to the other rental inventory models that also make the assumption of lost sales. In addition to the continuous-time rental inventory models of Tainiter (1964) and Whisler (1967) tabulated here, Papier and Thonemann (2008) build on the $M/M/c/c$ queueing model in Harel (1988), where approximations, as well as lower and upper bounds, are developed for the lost sales rate as a function of the system capacity. Extending this model to account for a compound Poisson arrival process, Papier and Thonemann (2008) conduct a stationary queueing analysis to obtain structural results for a fleet sizing problem and provide an approximation suitable for implementation. The use of the $M/M/c/c$ or $M/G/c/c$ queueing model as a basis for studying capacity management for rental systems further follows in Savin et al. (2005), Gans and Savin (2007), Adelman (2008), Hampshire et al. (2009), and Levi and Shi (2011). However, our work is different from this stream of research by our consideration of a discrete-time rental model with a finite rental season, random usage-based inventory loss, and an arbitrary demand model possessing the ability to capture any distributional characteristic.

In contrast to the continuous-time queueing models, Cohen et al. (1980) use a discrete-time model to represent a return process to a blood bank with the goal of determining an optimal order-up-to

Table 1 Comparison of lost sales rental inventory models.

	CONTINUOUS TIME			DISCRETE TIME		
	Taniter (1964)	Whisler (1967)	Papier and Thonemann (2008)	Cohen (1980)	Baron et al. (2011)	Our Paper
Inventory Decision	One Time	Repeated	One Time	Repeated	One Time	One Time
Time Horizon	Finite	Finite	Infinite	Finite	Finite	Finite
Demand Process	IID Interarrival Times	IID Interarrival Times	Compound Poisson Stationary; Also Nonstationary	General IID	Arbitrary	Arbitrary
Rental Duration	General IID	General IID	General IID	Deterministic	General IID with a Restricted Return Process	General IID
Inventory Loss	N/A	N/A	N/A	Constant Decay	N/A	Usage-Based Random Loss

level in every period. Reflecting hospitals' tendency to order significantly more units of blood than needed, a constant percentage of the quantity rented by hospitals is returned to the blood bank and the rest is consumed after a rental duration of a fixed number of periods. A constant percentage of the inventory leftover at the blood bank is, on the other hand, considered to have decayed. The problem of finding the optimal inventory level under a periodic review policy is formulated as a dynamic program and an approximate solution is provided. In comparison, we examine the one-time pre-seasonal ordering problem and consider the loss of inventory as random, instead of being a constant proportion. Furthermore, we do not require the assumption of an independent and identically distributed demand process, and we allow randomness in the rental duration.

Closest to our model is Baron et al. (2011), who determine the optimal pre-season order quantity for a video rental store with lost sales but no inventory loss. In particular, Baron et al. (2011) consider a return process that is monotone; i.e., the percentage of the rental units rented in period t and returned by period n is always greater than or equal to the percentage of the units rented in period $t + 1$ and returned by the same period n . The key result is the concavity of the expected number of rentals in the number of rental units procured. We are, on the other hand, the first to establish this result for a rental system with random usage-based inventory loss. We also address the issue of rental unit inventory allocation, which arises only in our rental inventory model as a result of accounting for random lifetimes of the rental units.

To analyze models in which rental unit lifetimes do not follow a geometric distribution, we use sample path analysis in a very general setting to prove the two main results of our work: the concavity of the expected profit function and the optimal rental unit recirculation rule. This approach has been used in various settings to model complexities of production and inventory systems. Examples include the number of customers and their utilities for a model with dynamic substitution by Mahajan and van Ryzin (2001) and the processing times for multi-station production lines by Muth (1979) and Tayur (1993). Also, our proofs of concavity bear similarities to that of Shanthikumar and Yao (1987) in their study of systems with multi-server stations.

3. Rental Model: A Sample Path Approach

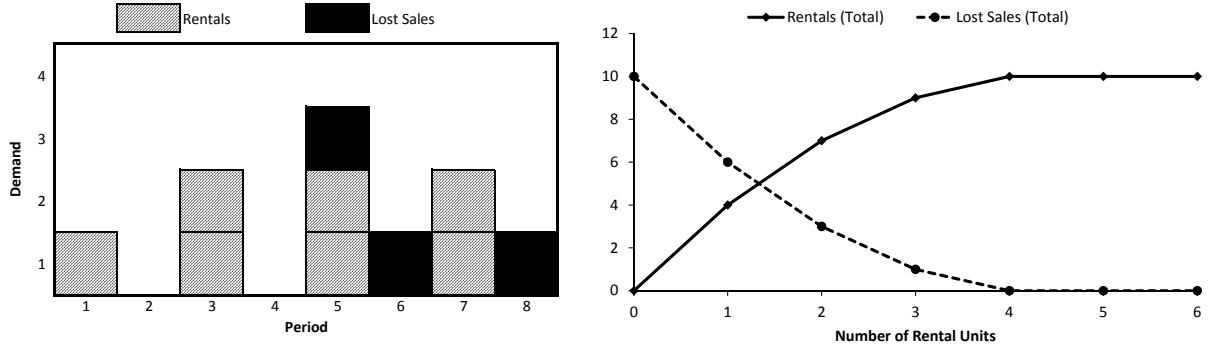
In this section, we describe a sample path approach to modeling a rental inventory system that allows us to analyze the system — including rules for recirculating rental units — under general assumptions about the demand process. Motivated by the problem of selecting the number of rental units to procure before the start of a finite rental season, we begin our analysis with the following two-stage model of a single-product, discrete-time rental inventory system with lost sales. In the first stage, the size of the rental inventory is chosen to be y . Each rental unit is procured before the start of the season and has a salvage value at the end of the rental season that depends on whether the rental unit retires from the inventory before the end of the season. Hence, the unit procurement cost accounts for not just the purchase price, but is adjusted to also include the salvage value for a dress in “good” condition and the cost of holding the item for the duration of the rental season. In the second stage, demands occur over N periods and the units purchased in the first stage are rented to satisfy the customer demands. Each customer is assumed to rent a single unit, and for simplicity we begin by considering the case in which each rental lasts for a deterministic duration of A periods. Thus, fulfilling one unit of demand requires that one unit of the inventory is withdrawn for the period in which the demand is received and for the $A - 1$ succeeding periods, resulting in the rental having a deterministic duration of A periods.

A critical aspect of rental inventory planning is to account for the loss of rental units. Misuse by customers, customer options to purchase rented items or simply the deterioration of the rental

unit's quality over time present reasons for why a unit would be retired from the rental inventory. To account for random usage-based loss of inventory in our model, we assume that each rental unit $m \in \{1, 2, \dots, y\}$ fails after a random number of rentals, \mathbf{l}_m , characterized by a general probability mass function. More precisely, upon completion of its \mathbf{l}_m th rental, unit m satisfies no further demands, although it does have an expected salvage value that is earned at the end of the horizon. In addition, the demand d_n is received in period $n \in \{1, 2, \dots, N\}$. Taken together, the demands d_1, d_2, \dots, d_N and the rental unit lifetimes $\mathbf{l}_1, \mathbf{l}_2, \dots, \mathbf{l}_y$ comprise a sample path, which we denote by $\boldsymbol{\xi}$; i.e., $\boldsymbol{\xi} = \{d_1, d_2, \dots, d_N, \mathbf{l}_1, \mathbf{l}_2, \dots, \mathbf{l}_y\}$. When rental unit loss probabilities change based on the number of times rented, we must also specify the recirculation rule γ to fully characterize the system's operation.

We use $R_n^\gamma(y, \boldsymbol{\xi})$ for the number of units rented and $L_n^\gamma(y, \boldsymbol{\xi})$ for the number of sales lost in period n as a function of the initial inventory of y rental units and the sample path $\boldsymbol{\xi}$ of demands and rental unit lifetimes for a recirculation rule γ . For convenience, the total number of rentals and lost sales over the entire horizon are defined as $\mathcal{R}^\gamma(y, \boldsymbol{\xi}) := \sum_{n=1}^N R_n^\gamma(y, \boldsymbol{\xi})$ and $\mathcal{L}^\gamma(y, \boldsymbol{\xi}) := \sum_{n=1}^N L_n^\gamma(y, \boldsymbol{\xi})$, respectively. We also let $W_n^\gamma(y, \boldsymbol{\xi})$ denote the number of units that are successfully returned to the system in the beginning of period n and available to be rented again in that period, and define $Z_n^\gamma(y, \boldsymbol{\xi}) := R_{n-A}^\gamma(y, \boldsymbol{\xi}) - W_n^\gamma(y, \boldsymbol{\xi})$ as the number of rental units that would have been returned in period n but were lost. A reward r is earned every time a unit is rented, and c is the unit cost of a lost sale.

The rental system operates for period n of the second stage as follows: (1) Of all the items rented in period $n - A$, $W_n^\gamma(y, \boldsymbol{\xi})$ units are returned while $Z_n^\gamma(y, \boldsymbol{\xi})$ retire from the rental inventory. After returns are received but before rentals are made, the total inventory available to rent out in period n is $I_n^\gamma(y, \boldsymbol{\xi}) := y - \sum_{t=1}^{n-1} R_t^\gamma(y, \boldsymbol{\xi}) + \sum_{t=1}^n W_t^\gamma(y, \boldsymbol{\xi})$. (2) The demand D_n is realized as d_n . If $d_n \leq I_n^\gamma(y, \boldsymbol{\xi})$, then d_n units are rented out. Otherwise, $I_n^\gamma(y, \boldsymbol{\xi})$ units are rented out. More succinctly, $R_t^\gamma(y, \boldsymbol{\xi}) := d_n \wedge I_n^\gamma(y, \boldsymbol{\xi})$, where $a \wedge b$ denotes the minimum of a and b . The rental unit recirculation rule determines which rental unit is allocated to satisfy each unit of demand, and consequently

(a) Illustration with two rental units ($y = 2$).(b) Effect of number of rental units y .**Figure 1** Number of rentals and lost sales for Example 1 with a rental duration of two periods ($A = 2$).

determines $W_n^\gamma(y, \xi)$ and $Z_n^\gamma(y, \xi)$. (3) Excess demand $L_n^\gamma(y, \xi) := [I_n^\gamma(y, \xi) - d_n]^-$, which can be alternatively written as $d_n^\gamma - R_n^\gamma(y, \xi)$, is lost. Therefore, given the sample path ξ , the dynamics of the rental system's operation can be represented recursively as follows, where $I_0^\gamma(y, \xi) = y$ and $R_t^\gamma(y, \xi) = 0$ for $t \leq 0$:

$$I_{n+1}^\gamma(y, \xi) = I_n^\gamma(y, \xi) - R_n^\gamma(y, \xi) + W_{n+1}^\gamma(y, \xi).$$

$$R_{n+1}^\gamma(y, \xi) = d_{n+1} \wedge I_{n+1}^\gamma(y, \xi). \quad (1)$$

$$L_{n+1}^\gamma(y, \xi) = d_{n+1} - R_{n+1}^\gamma(y, \xi).$$

EXAMPLE 1. Figure 1 illustrates this rental system with a demand sequence of $\{d_1, \dots, d_8\} = \{1, 0, 2, 0, 3, 1, 2, 1\}$ for eight periods ($N = 8$). Each rental lasts for two periods ($A = 2$); i.e., a unit that is rented in period n will next be available to be rented again in period $n + 2$. If the system would operate with only one rental unit (i.e., $y = 1$), then that unit would be rented in periods 1, 3, 5, and 7 for a total of four rentals, while six units of the demand would be lost. Figure 1a shows how the demand is divided into rentals and lost sales for a system with $y = 2$ rental units. Thus, the addition of the second rental unit allows an additional unit of demand to be satisfied in periods 3, 5, and 7, so that there are now 7 units of fulfilled demand and 3 units of lost sales. Figure 1b shows how the number of rentals and lost sales change with the number of rental units y . We observe that the number of rentals is concave in y and that the number of lost sales is convex in y on this

sample path. In other words, the number of additional rentals produced by one additional rental unit (i.e., the slope of the rentals curve) is decreasing in y .

As the return process depends on the specific recirculation rule γ , we will describe $W_{n+1}^\gamma(y, \xi)$ and $Z_{n+1}^\gamma(y, \xi)$ as needed when referring to specific rules. To account for the return of rental units that are rented in periods $N - A + 1, N - A + 2, \dots, N$, we define $W_{N+1}^\gamma(y, \xi), W_{N+2}^\gamma(y, \xi), \dots, W_{N+A}^\gamma(y, \xi)$ as the returns and $Z_{N+1}^\gamma(y, \xi), Z_{N+2}^\gamma(y, \xi), \dots, Z_{N+A}^\gamma(y, \xi)$ as the lost units in each of the corresponding periods. The total number of lost rental units is denoted by $Z^\gamma(y, \xi) := \sum_{n=A+1}^{N+A} Z_n^\gamma(y, \xi)$.

One way to model rental unit loss is to consider geometrically distributed rental unit lifetimes. The memorylessness of the geometric distribution leads to a constant probability of rental unit loss over time. However, if a rental unit does indeed have a higher probability of wearing out over time, then a rental unit lifetime distribution with an increasing failure rate (i.e., a loss probability increasing with the number of times the unit has been rented) would be a suitable choice. Bikes, cars and large equipment are examples of assets for which an increasing loss probability as a function of the number of rentals could be used to model the rental unit lifetime. Furthermore, lifetimes that are deterministic — when enforced by safety regulations that require their disposal after a certain number of uses — can be analyzed as a special case of an increasing loss probability.

Next, we discuss how to incorporate the salvage value of a rental unit into our rental inventory model. This is an important issue because the salvage value of a rental unit that retires from the rental inventory during the rental season may differ from the salvage value of a unit that is still functional at the end of the season. In that case, we separately define the procurement cost s^g for the unit that can be still rented at the end of the rental season and the procurement cost s^b for the unit that has already retired from inventory. The relation $s^g \leq s^b$ indicates that the unit retiring from the rental inventory has been damaged. Hence, it has lost a portion of its value. The relation $s^g \geq s^b$ may, on the other hand, represent the purchase of the rental unit by the customer who is renting it as discussed in Section 1 for the rental companies Redbox and Rent-A-Center.

To account for the cost of inventory loss in the objective function of our rental inventory model, the reduction in the salvage value of a lost rental unit ($s^b - s^g$) is multiplied by the number of lost

rental units and subtracted from the revenue as part of the profit function, which we denote by $\Pi^\gamma(y, \boldsymbol{\xi})$. Consequently, we obtain the profit function as follows:

$$\Pi^\gamma(y, \boldsymbol{\xi}) = r \sum_{n=1}^N D_n - (r + c) \mathcal{L}^\gamma(y, \boldsymbol{\xi}) - s^g y - (s^b - s^g) \mathcal{Z}^\gamma(y, \boldsymbol{\xi}).$$

We are now ready to formulate the rental inventory optimization problem as the maximization of the expected profit function $\pi^\gamma(y) := \mathbb{E}[\Pi^\gamma(y, \boldsymbol{\xi})]$ subject to $y \geq 0$. We investigate the concavity of this expected profit function in the initial inventory of y rental units for geometric lifetime distributions in Section 4 and for general lifetime distributions in Section 5.

We can also extend this basic model to include a random duration for each rental. We define $A_{m,i}$ as the random variable denoting the rental duration for the i th demand served by the m th rental unit for $i \geq 1$ and $m = 1, 2, \dots, y$. Each rental lasts for any number of periods between a minimum of A_{min} and a maximum of A_{max} ; i.e., $A_{m,i} \in \{A_{min}, A_{min} + 1, \dots, A_{max}\}$. We consider $A_{m,i}$ to be independent and identically distributed according to a general probability mass function characterized by $h(a) := \mathbb{P}\{A_{m,i} = a\}$, $a = A_{min}, A_{min} + 1, \dots, A_{max}$. For any sample path, we let $D_{a,n}$ represent the number of units of demand in period n , $n = 1, 2, \dots, N$, that have a duration of a periods. Thus, our model of stochastic rental duration differs from that of Cohen et al. (1980), who model unit demands as having different durations with respect to constant proportions. It also differs from that of Baron et al. (2011), who require the following condition on the rental duration.

We let $R_{a,n}^\gamma(y, \boldsymbol{\xi})$ denote of the number of rentals of duration a that begin in period n , and $\mathcal{R}_a^\gamma(y, \boldsymbol{\xi})$ denote the number of rentals of duration a that occur over the entire rental horizon. Similarly, $\mathcal{L}_a^\gamma(y, \boldsymbol{\xi})$ represents the number of rentals of duration a that are lost over the entire rental horizon. We allow $W_{a,n}^\gamma(y, \boldsymbol{\xi})$ and $Z_{a,n}^\gamma(y, \boldsymbol{\xi})$ to represent the number of rental units returned and lost, respectively, in period n after a rental duration of a periods with $\mathcal{Z}_a^\gamma(y, \boldsymbol{\xi}) := \sum_{n=A+1}^{N+A} Z_{a,n}^\gamma(y, \boldsymbol{\xi})$. Furthermore, we define $a_{m,i}$ to denote the realized rental duration of the i th demand served by the rental unit m . It is important to note that the sample path $\boldsymbol{\xi}$ now consists of not only the demand realizations d_n , $n = 1, 2, \dots, N$, and the rental unit lifetimes \mathbf{l}_m , $m = 1, 2, \dots, y$, but also the rental durations $a_{m,i}$, $i \geq 1$ and $m = 1, 2, \dots, y$. Also, we use r_a for denoting the reward earned with a

rental that has a duration of a periods. With this notation, the profit on any sample path can be expressed as:

$$\Pi^\gamma(y, \xi) = \sum_{n=1}^N \sum_{a=A_{min}}^{A_{max}} r_a D_{a,n} - \sum_{a=A_{min}}^{A_{max}} (r_a + c) \mathcal{L}_a^\gamma(y, \xi) - s^g y - (s^b - s^g) \mathcal{Z}^\gamma(y, \xi).$$

When rental duration is random, the convexity of the number of lost sales $\mathcal{L}(y, \xi)$ in y and thus, the concavity of the number of rentals $\mathcal{R}(y, \xi)$ in y , might not hold for every sample path ξ . As an example, we consider the addition of two rental units to our inventory system, where the first additional unit fulfills one customer demand with a very long duration and the second additional unit fulfills several customer demands with short rental durations. In this case, the number of additional customer demands satisfied by one extra rental unit is not necessarily non-increasing in y . Therefore, we proceed by analyzing the structural properties of the expected number of lost sales and the expected number of rentals. We are the first to consider this modeling aspect simultaneously with random loss of rental inventory in the following section. It is worth noting that the random rental duration accounts for each customer's decision to keep the rental unit for a different number of periods, but it can also include the random service time needed to repair the rental unit depending on its condition upon return.

4. Rental Inventory Loss with Geometric Lifetime Distributions

This section considers a model in which each rental unit m has a loss probability of p with each rental. More specifically, the random variable \mathbf{l}_m , which denotes the number of times the unit $m \in \{1, 2, \dots, y\}$ is rented before retiring from the rental inventory, follows a geometric distribution with an expected value of $1/p$. We assume that $r + c \geq p(s^b - s^g)$. This condition implies that the benefit, $r + c$, of converting a lost sale into a rental is greater than or equal to the expected cost, $p(s^g - s^b)$, of the rental unit loss. Due to the constant failure rate, the recirculation rule has no effect on $\mathbb{E}[\mathcal{L}_n^\gamma(y, \xi)]$, $\mathbb{E}[\mathcal{R}_n^\gamma(y, \xi)]$, or $\mathbb{E}[\mathcal{W}_n^\gamma(y, \xi)]$ for any n . Therefore, we omit the superscript in the notation used in this section.

We establish the concavity of the expected profit function by first presenting a condition related to the rental return process for which the expectation of the number of lost sales $\mathcal{L}(y, \xi) :=$

$\sum_{n=1}^N L_n(y, \xi)$ is convex. Correspondingly, the expectation of the number of rentals $\mathcal{R}(y, \xi) := \sum_{n=1}^N R_n(y, \xi)$ is concave in the initial inventory of y rental units for this condition.

LEMMA 1. *If the expected number of rental units returned $\mathbb{E}[\sum_{t=1}^n W_t(y, \xi)]$ is concave and non-decreasing in y for $n = 1, 2, \dots, N$, then the expected number of lost sales $\mathbb{E}[\mathcal{L}(y, \xi)]$ is convex and non-increasing while the expected number of rentals $\mathbb{E}[\mathcal{R}(y, \xi)]$ is concave and non-decreasing in y .*

All proofs appear in the appendix.

PROPOSITION 1. *When rental unit lifetimes are geometrically distributed, the expected profit $\pi(y)$ is concave in y for any rental unit recirculation rule.*

5. Rental Inventory Loss with General Lifetime Distributions

When the lifetimes of the rental units follow a general distribution, the number of rental units returned in period n may depend on the policy used to choose among available rental units to satisfy the demand in previous periods. Also, the number of rentals $\mathcal{R}(y, \xi)$ and the number of lost sales $\mathcal{L}(y, \xi)$ might not necessarily be concave and convex, respectively, in y due to a rental unit that has a particularly long or short lifetime. Therefore, we investigate whether it is possible to establish the concavity of the expected number of rentals $\mathbb{E}[\mathcal{R}(y, \xi)]$ as well as the convexity of the expected number of lost sales $\mathbb{E}[\mathcal{L}(y, \xi)]$ in the initial inventory of y rental units. We modify the assumption of $r + c \geq p(s^b - s^g)$ for deterministic rental duration to be $\sum_{a=A_{min}}^{A_{max}} r_a h(a) + c \geq p(s^b - s^g)$ for random rental duration.

Because we have not yet found a direct algebraic proof, we compare sample paths via coupling, as described in Chapter 4 of Lindvall (1992). A coupling approach allows us to compare the value of an additional rental unit in two systems that differ only in the number of rental units. Due to the rental unit lifetime distributions and the recirculation rule, analysis of the change in the expected number of rentals would otherwise be extremely difficult. Our approach uses the following steps:

1. Establish demand values d_1, d_2, \dots, d_N , which do not require any distributional assumptions.
2. Operate the system with y rental units, each of which has a lifetime \mathbf{l}_m , $m = 1, \dots, y$. The i th demand, $i \geq 1$, served by rental unit m has duration $a_{m,i}$.

3. Add an additional rental unit — the $(y+1)$ st unit to the system — that has a lifetime \mathbf{l}' and serves demands with durations $\{a'_1, a'_2, \dots\}$. To be clear, the system has rental units with lifetimes $\mathbf{l}_1, \mathbf{l}_2, \dots, \mathbf{l}_y, \mathbf{l}'$.

4. To the system described in Step 2 (i.e., ignoring Step 3), add a $(y+1)$ st unit that has a lifetime \mathbf{l}_{y+1} and serves demands with durations $\{a_{y+1,1}, a_{y+1,2}, \dots\}$.

5. To the system described in Step 4, add an additional rental unit — the $(y+2)$ nd unit — so that the system has rental units with lifetimes $\mathbf{l}_1, \mathbf{l}_2, \dots, \mathbf{l}_y, \mathbf{l}_{y+1}, \mathbf{l}'$. This additional rental unit has the same lifetime \mathbf{l}' and serves demands with same durations $\{a'_1, a'_2, \dots\}$ as the additional unit added to the system in Step 3.

For notational convenience, we define $\boldsymbol{\xi}(y)$ as the sample path consisting of the demands d_1, d_2, \dots, d_N of all N periods, the rental unit lifetimes $\mathbf{l}_1, \mathbf{l}_2, \dots, \mathbf{l}_y$, and the demand durations $\{a_{m,1}, a_{m,2}, \dots\}$ for $m = 1, 2, \dots, y$, as well as the lifetime \mathbf{l}' and rental durations $\{a'_1, a'_2, \dots\}$ for an additional rental unit. For example, $\boldsymbol{\xi}(y)$ and $\boldsymbol{\xi}(y+1)$ contain all of the sample path information necessary to analyze the systems described in Steps 3 and 5, respectively.

We consider two types of decisions for rental unit allocations. First, we examine a “count-based” rental unit state in which the allocation decision is based on the number of times that each unit has been rented. Then, we study a “condition-based” rental unit state in which the allocation decision is based on the current state of each rental unit. Each of these models may be relevant for Rent the Runway. Specifically, the dress’s physical condition may not be observed — requiring a count-based model — if it is not carefully inspected or if the cause of a dress failure is difficult to observe as the dress’s condition degrades. For example, a zipper may be more likely to fail over time even if indications of impending failure may not be observed. On the other hand, a dress’s physical condition may be observed if it relates to the condition of the fabric. Satin dresses are susceptible to developing minor damage to individual threads due to their loose weaves, and the repeated ironing of silk taffeta dresses may cause them to lose their ideal appearance around pleats and seams. A condition-based model would then be more appropriate for this setting.

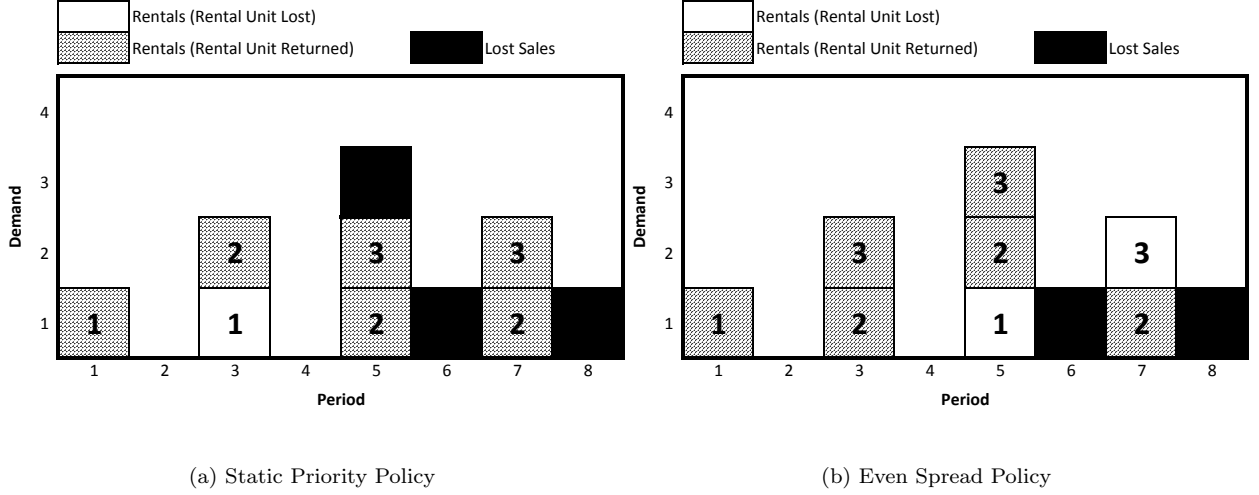


Figure 2 Rental unit recirculation schemes for Example 2 with $y = 3$. The number inside a box identifies the rental unit that satisfies a demand.

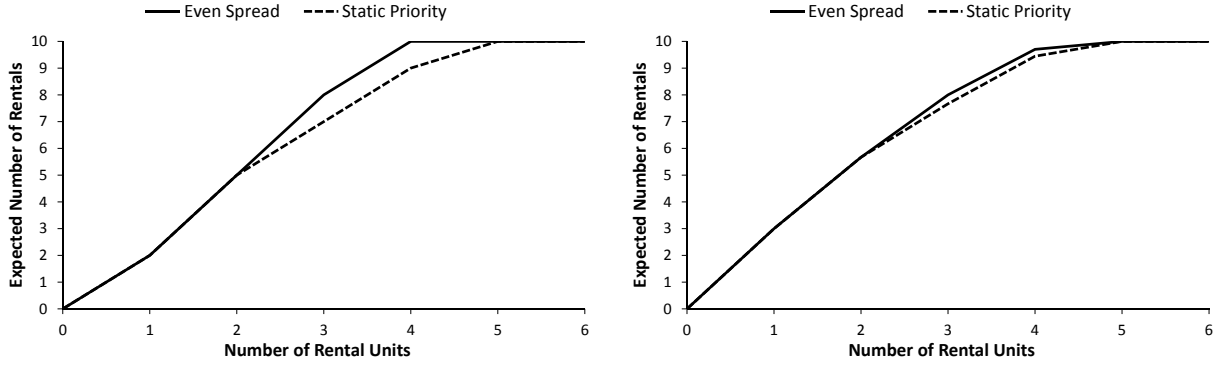
5.1. Count-Based Rental Unit State

For the analysis in this section, we let ℓ_i denote the probability that a rental unit has a lifetime of i rentals; i.e., it is the probability that a rental unit retires from the rental inventory after its i th rental, or $Pr(\mathbf{l}_m = i)$. Under the assumption of deterministic rental duration in which each rental has a duration of A periods, we assume that $r + c \geq (s^b - s^g)\ell_i$ for $1 \leq i \leq N/A$; i.e., the benefit of an additional rental to a customer (i.e., $r + c$) is greater than or equal to the expected reduction in the salvage value due to rental inventory loss (i.e., $(s^b - s^g)\ell_i$, $i \geq 1$). Similar to the presentation in Section 4 for geometrically distributed lifetimes, this condition takes the form $\sum_{a=A_{min}}^{A_{max}} r_a h(a) + c \geq (s^b - s^g)\ell_i$, $i \geq N/A_{min}$, for random rental duration so that the expected profit of serving a customer is not negative.

In this section, we examine two recirculation rules: the even spread policy and the static priority policy. In the even spread policy, each demand is served by an available rental unit that has been rented out the fewest number of times among all available rental units. We note that the priority is assigned to rental units in the order of increasing hazard rate under the even spread policy. The static priority policy, on the other hand, allocates rental units according to a priority list that does not change over the course of the rental horizon.

EXAMPLE 2. Figure 2 shows how different rental unit recirculation schemes can affect the number of rentals and lost sales for an example sample path ξ . Using the same demand values as in Example 1, we now state the lifetimes of available rental units as $\{\mathbf{l}_1, \mathbf{l}_2, \dots, \mathbf{l}_5\} = \{2, 4, 3, 4, 2\}$. For $y = 1$ and $y = 2$ with $\mathbf{l}_1 = 2$ and $\mathbf{l}_2 = 4$, both the even spread and static priority policies satisfy the same number of demands; i.e., $\mathcal{R}(1, \xi) = 2$ and $\mathcal{R}(2, \xi) = 5$. However, when $y = 3$, the even spread recirculation rule enables one more rental over the rental horizon than the static priority rule. Under the static priority rule, rental unit 1 is lost after serving a demand in period 3, while it is lost after serving a demand in period 5 under the even spread rule. This allows one extra demand to be served in period 5 for the even spread rule because it has one more rental unit available than the static priority rule. Figure 3a shows that the even spread rule also serves one more demand than the static priority rule when $y = 4$ and that both policies serve all ten units of demand when $y \geq 5$. Figure 3a also demonstrates that the number of rentals is not necessarily concave in y ; i.e., the addition of rental unit 1 with lifetime $\mathbf{l}_1 = 2$ satisfies fewer additional units of demand than the addition of rental unit 2 with lifetime $\mathbf{l}_2 = 4$.

In Figure 3b, we use the demand values from Example 1 but instead let the lifetime of each rental unit be a discrete uniform random variable between 2 and 4 (i.e., $\ell_2 = \ell_3 = \ell_4 = 1/3$), and estimate the expected number of rentals with a simulation executed for a sufficiently large number of replications so that the standard error of the experiment is negligible. Even though concavity is violated on individual sample paths, the expected number of rentals is revealed to be a concave function of the number of rental units. The even spread and static priority policies result in the same number of rentals regardless of the sample path for $y \leq 2$ and $y \geq 5$. However, the expected number of rentals for the even spread policy exceeds that of the static priority policy by 0.33 when $y = 3$ and by 0.26 when $y = 4$. When $y = 3$, the even spread policy results in at least one more rental than the static priority policy on 44.1% of the sample paths, and the static priority policy exceeds the even spread policy on 10.9% of all sample paths.



(a) Example sample path.

(b) Expected number of rentals.

Figure 3 Effect of the number of rental units and rental unit recirculation rule for Example 2.

5.1.1. The Static Priority Recirculation Rule The static priority recirculation rule, denoted by the superscript SP , selects the units to be rented according to a priority list, which is the same for every time period. That is, when a rental unit is needed to satisfy demand, the one with the highest priority among the set of available rental units is chosen. Therefore, when the rental duration is deterministic, the static priority recirculation rule is the same as the policy selecting the rental unit that has been rented the most.

PROPOSITION 2. *In a rental system with random rental duration and general rental unit lifetime distributions, $\pi(y)$ is concave and non-decreasing in y for the static priority recirculation rule.*

5.1.2. The Even Spread Recirculation Rule We now consider the even spread recirculation policy — denoted by the superscript ES — which satisfies a demand with the rental unit that has been rented the fewest number of times. Defining $R_{n,m}^\gamma(y, \xi)$ as the number of times that rental unit m is rented in period n under some policy γ , an even spread compliant policy selects a rental unit m to satisfy a demand based on available rental units that minimize $\sum_{t=1}^{n-1} R_{t,m}^\gamma(y, \xi)$. When the loss probability for rental units is non-decreasing in the number of times rented, recirculation priority corresponds to a hazard rate ordering. We assume that ties are broken by some static priority list for allocating rental units. We first investigate the relation between the expected number of rentals and the initial rental inventory level (Proposition 3) and discuss the concavity of the

expected profit function in the initial rental inventory (Proposition 4). We then demonstrate the optimality of the even spread policy to maximize the expected profit when the loss probability of each rental unit increases with the number of times that the unit has been rented (Proposition 5).

PROPOSITION 3. *For a rental system with random rental duration that follows the even spread recirculation rule, the expected number of rentals $\mathbb{E}[\mathcal{R}^{ES}(y, \xi)]$ is concave and non-decreasing in the initial inventory of y units.*

For the special case of $s^g = s^b$, the expected profit can be expressed in terms of the linear procurement cost and the expected number of rentals. Therefore, Proposition 3 also implies the concavity of the expected profit function $\pi^{ES}(y)$ in y when $s^g = s^b$. However, when $s^g \neq s^b$, we must show that the expected profit function $\pi^{ES}(y)$ remains concave in the initial inventory of y rental units.

PROPOSITION 4. *When the rental duration is random and the loss probability of each rental unit increases with the number of times that the unit has been rented, the expected profit $\pi^{ES}(y)$ is concave in the initial inventory of y rental units under the even spread recirculation rule.*

When the rental unit loss probability is increasing in the number of times that the unit has been rented, we identify the even spread policy as the optimal rental unit recirculation rule to maximize the expected profit. Our key argument is a pairwise interchange argument in which switching an allocation that violates the even spread policy to conform to the even spread policy increases the expected number of rentals. We require additional notation to compare sample paths and make this argument, which we describe along with an overview of the steps of the proof:

1. Find the first allocation decision over the rental horizon that violates the even spread policy. We denote this existing policy with the superscript V for “violating.” Assume that this violating decision occurs in some period n . Specifically, a rental unit j is allocated to demand when some other rental unit i is available and $\sum_{t=1}^{n-1} R_{t,j}^V(y, \xi) > \sum_{t=1}^{n-1} R_{t,i}^V(y, \xi)$. The availability of rental units i and j implies that $\sum_{t=1}^{n-1} R_{t,j}^V(y, \xi) < \mathbf{1}_j$ and $\sum_{t=1}^{n-1} R_{t,i}^V(y, \xi) < \mathbf{1}_i$.

2. Consider a switched forward allocation path of units i and j in periods $n, n+1, \dots, N$ so that rental unit i is allocated instead of rental unit j . We refer to this allocation with the superscript S for “switched.”

3. Change values in ξ related to the lifetimes and rental durations after period n for units i and j with two new partial sample path vectors $\xi_{(1)}$ and $\xi_{(2)}$. Specifically, we generate two sets of random durations $(a_{(1),1}, a_{(1),2}, \dots)$ and $(a_{(2),1}, a_{(2),2}, \dots)$ for demands after period n served by two different rental units and inverse probability mass function values $\eta_{(1)}$ and $\eta_{(2)}$ for the conditional lifetime distributions of the two rental units. The latter information allows the determination of the lifetimes $\mathbf{l}_{(1)}$ and $\mathbf{l}_{(2)}$.

4. Calculate the number of rentals over the entire horizon under four scenarios (with corresponding notation for the total number of rentals used for convenience): (1) $\mathcal{R}^V(\xi_{(1)}, \xi_{(2)})$ for the violating allocation with $\xi_{(1)}$ applied to rental unit i and $\xi_{(2)}$ to rental unit j ; (2) $\mathcal{R}^V(\xi_{(2)}, \xi_{(1)})$ for the violating allocation with $\xi_{(2)}$ applied to unit i and $\xi_{(1)}$ to unit j ; (3) $\mathcal{R}^S(\xi_{(1)}, \xi_{(2)})$ for the switched allocation with $\xi_{(1)}$ applied to unit i and $\xi_{(2)}$ to unit j ; and (4) $\mathcal{R}^S(\xi_{(2)}, \xi_{(1)})$ for the switched allocation with $\xi_{(2)}$ applied to unit i and $\xi_{(1)}$ to unit j .

5. Compare scenarios to observe that $\mathbb{E}[\mathcal{R}^S(y)] \geq \mathbb{E}[\mathcal{R}^V(y)]$, which implies that $\mathbb{E}[\Pi^S(y)] \geq \mathbb{E}[\Pi^V(y)]$ under certain assumptions on the cost parameters.

6. Go to Step 1 and repeat until the switched allocation is equivalent to the even spread allocation.

PROPOSITION 5. *If the loss probability of each rental unit increases with the number of times that the unit has been rented and the rental duration is random, the even spread recirculation rule maximizes the expected profit.*

5.2. Condition-Based Rental Unit State

We now study a different model of rental units in which each rental unit has a known state in the set $\{1, 2, \dots, S\}$ that may change after each time that the unit is rented. On a sample path ξ , we define s_{mi} as the state of rental unit m after it is rented for the i th time, $m \in \{1, 2, \dots, y\}$

and $i \in \{1, 2, \dots, l_m\}$. The initial state of each rental unit is defined as $s_{m0} = 1$ and a rental unit's retirement from recirculation corresponds to $s_{m, l_m} = S$. A transition probability matrix P governs the evolution of each rental unit's state upon each instance in which the unit is rented. We define $P(i, j)$ as the probability that a rental unit transitions from state i to state j after each rental with $i, j \in \{1, 2, \dots, S\}$. We also assume that $\sum_{a=A_{min}}^{A_{max}} r_a h(a) + c \geq (s^b - s^g)P(i, S)$ for $i = 1, 2, \dots, S - 1$ so that the expected value of offering a rental is never negative.

One simple recirculation policy based on the observed rental unit state is to allocate the rental units in increasing order of their state. In other words, the rental unit that is in the best condition is given the highest allocation priority. We label this policy as the “best-first” policy. Similarly, the “worst-first” policy gives the highest priority to the rental unit in the worst condition for which it can still be rented out. For either policy, we show that the expected number of rentals is concave in the initial inventory level.

PROPOSITION 6. *The expected profit $\pi(y)$ is concave in the initial inventory of y rental units under either the best-first or worst-first policy with random rental duration.*

We next consider the optimal rental unit recirculation policy when rental unit selection decisions are based on the rental unit condition. We assume that the transition matrix is totally positive of order 2; i.e., that $P(i, j)P(i', j') \geq P(i, j')P(i', j)$ for all $i < i', j < j'$. Brown and Chaganty (1983) show that this property implies that the first passage time from state 1 to some state $C_j = \{i : i > j\}$ has an increasing failure rate for $j = 1, \dots, S - 1$.

PROPOSITION 7. *If the transition matrix P is totally positive of order 2, then the best-first policy maximizes the expected profit for a system with random rental duration.*

6. An Industrial Numerical Analysis: Rent the Runway

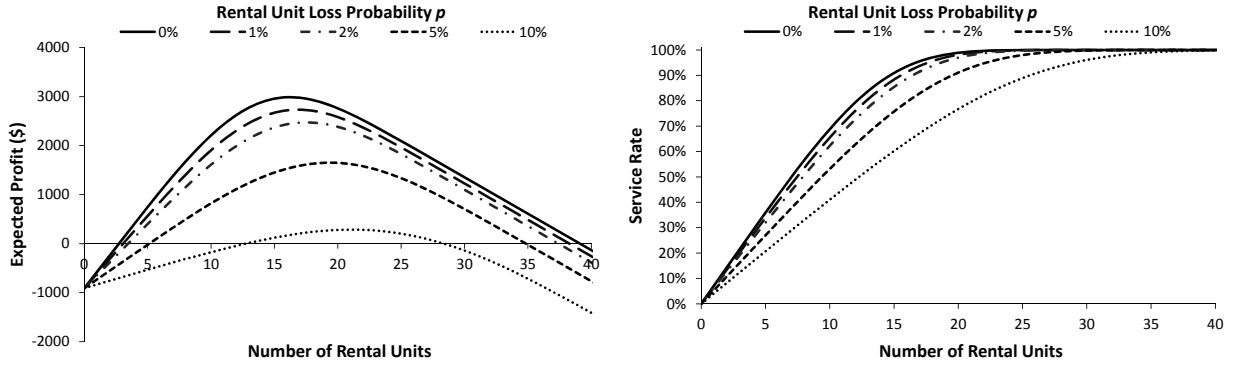
Motivated by the high-fashion dress rental business Rent the Runway, we introduce the model parameters representing a rental system with usage-based loss of inventory in Section 6.1. We discuss the impact of the rental inventory loss on the optimal procurement decision in Section 6.2

and the effect of the rental unit recirculation rule on rental inventory management in Section 6.3. All numerical testing is performed via sample average approximation, as described in Kleywegt et al. (2002) and Shapiro (2003).

6.1. Rental Model Parameters

The product we consider is a “middle-tier” dress as described in Eisenmann and Winig (2012); i.e., a full-price rental provides a net revenue of \$59, which is the difference between \$90 in revenue and \$31 in costs of cleaning, shipping, packaging and credit card processing. However, customers are allowed to rent a second style for \$25 and a second size for free; thus, a unit may not achieve \$59 in net revenue every time it is rented. We assume that these three scenarios for a rental — renting as the primary dress with net revenue of \$59, renting as the secondary dress with net revenue of \$20, and renting as the free second size with net cost of \$5 — occur with probabilities 50%, 20%, and 30%, resulting in an expected net revenue of $r = \$32$ per rental.

Eisenmann and Winig (2012) report that Rent the Runway purchases a middle-tier dress with a retail price of \$750 for \$226. We assume an annual unit holding cost that is equal to 20% of the purchase price of the dress to account for the cost of storage and the cost of capital. At the end of a fashion season, dresses in a variety of conditions are sold in New York City at what is known as a “sample sale.” Based on websites such as Yannetta (2013) that report on these sales, we let a dress in good condition sell for 80% – 85% off of the \$750 retail price and a dress in bad condition (i.e., a dress that retires from the rental inventory) to sell for 95% off of the retail price. Adjusting these sample sale prices for staging and transaction costs, we assume a dress that does not retire from the rental inventory by the end of the season to have a salvage value of \$100 and a dress that retires from the rental inventory to have a salvage value of \$30. We also calculate procurement costs separately for these two types of dresses by combining their purchase prices, holding costs and salvage values. For a 26-week horizon, the cost of procuring a dress is $s^g := \$149$, which consists of a purchase price of \$226, a holding cost of \$23 and a salvage value of \$100. A dress retiring from the rental inventory incurs an additional penalty of \$70, resulting in a procurement cost of



(a) Expected profit.

(b) Service rate.

Figure 4 The optimal inventory level increases with the loss probability for a 26-week system, and ignoring inventory loss significantly reduces expected profit compared to the optimal inventory level.

$s^b := \$219$. Finally, we choose $c := \$5$ as a customer goodwill penalty for the loss of a sale. With these parameters, a dress must be rented five times (seven times, on average, if the possibility of loss exists) to break even based on the ratio s^g/c (s^b/c).

With each period corresponding to a week, we consider Poisson distributed demand with a mean of $\lambda = 7$ per week and a rental horizon of $N = 26$ weeks, which corresponds to one of two major fashion seasons each year. We will also consider a longer rental horizon of $N = 52$ for a dress that could be in style for two consecutive seasons. We model each rental duration as lasting for a constant of $A = 2$ periods; i.e., the rented dress will be unavailable during the weekend for which it is rented and the weekend either preceding or following that weekend, depending on the day of the week on which the rental begins. A more granular representation of the rental duration in terms of the individual days is certainly possible. However, we believe that weekly periods adequately represent the system under the assumption that customers of Rent the Runway rent dresses primarily for weekend events.

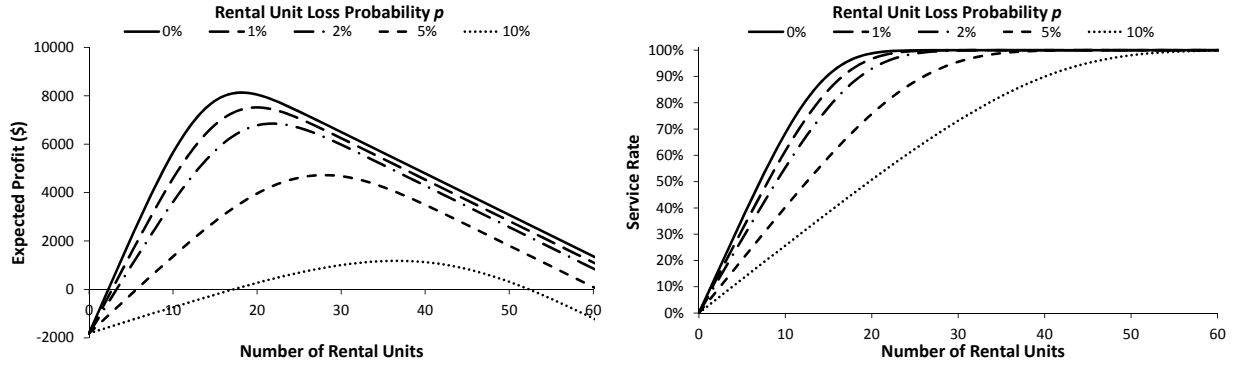
6.2. Rental Inventory Loss

The objective is to investigate the importance of accounting for the possibility of usage-based loss when choosing the initial inventory of rental units. We first analyze a rental system with a simple demand model in which demand is independent and identically distributed in each period

(Section 6.2.1). Then, we consider an advanced demand model that incorporates autocorrelated demand as an important trait for the sale of fashion items (Section 6.2.2).

6.2.1. Independent and Identically Distributed Demand Allowing the lifetime of each rental unit to follow a geometric distribution with a loss probability of $p \in \{0, 0.01, 0.02, 0.05, 0.10\}$, we illustrate the expected profit as a function of the initial inventory of rental units for the short rental horizon of $N := 26$ periods in Figure 4, and compare it to the “state of the art” in rental inventory management represented by Baron et al. (2011), which does not include inventory loss. Consistent with Proposition 1, we observe the expected profit function to be concave in the number of rental units to procure in the beginning of the rental season. In the system with no inventory loss ($p = 0$), we identify the optimal solution as 16 units with a corresponding service rate — i.e., the percentage of customers that are served — of 93.5%. However, when there is the possibility of inventory loss (i.e., $p > 0$), we find the optimal number of rental units to increase in the rental unit loss probability. Specifically, for a 5% loss probability, the optimal policy is to add three rental units to the initial inventory. Hence, ignoring inventory loss and using 16 rental units instead of the optimal 19 rental units results in a reduction of 7.3% in the expected profit. Furthermore, the service rate would only be 79.4% instead of the 88.7% corresponding to the optimal number of rental units for the system with $p = 5\%$.

Figure 5 shows that the impact of ignoring inventory loss is more dramatic for the longer rental horizon covering 52 weeks than for the shorter rental horizon with 26 weeks. This can be explained by the availability of fewer rental units to rent towards the end of the longer rental horizon. The comparison of Figure 5 to Figure 4 reveals more asymmetry in the expected profit as a function of the initial inventory of rental units for the longer horizon. More specifically, the slope of the expected profit function for a lower value of the number of rental units is steeper because each rental unit averts more lost sales in a long horizon than in a short horizon. Furthermore, the higher optimal service rate for the system with the longer horizon than the system with the shorter horizon reflects the higher value of a marginal rental unit. In other words, the consequence of having too few rental units is more severe in the longer horizon.



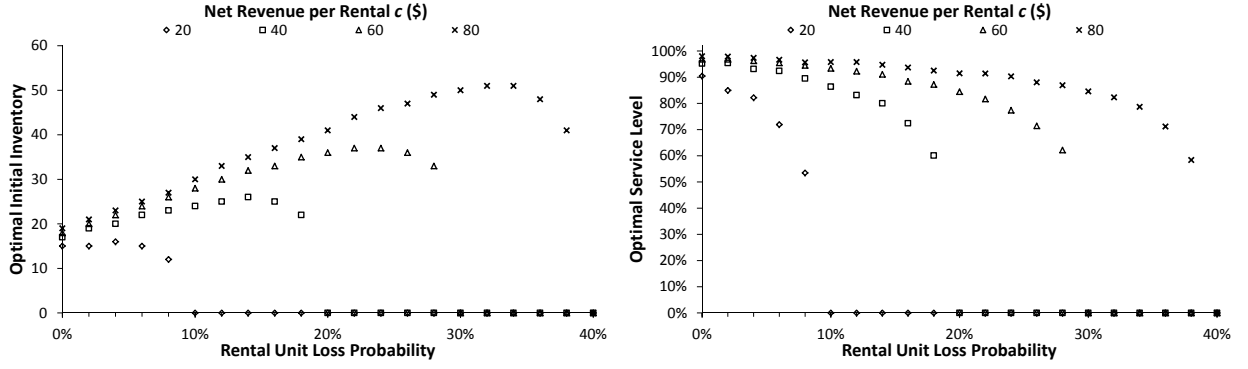
(a) Expected profit.

(b) Service rate.

Figure 5 Accounting for inventory loss is more important in terms of effect on expected profit for a 52-week system.

For rental systems considered in Figure 4 and Figure 5, the optimal policy is to always add more rental inventory to account for the loss of rental units; i.e., the profit-maximizing inventory level is increasing in p for $p \in \{0, 0.01, 0.02, 0.05, 0.10\}$. However, if the loss probability is sufficiently high, then the units will not be rented enough to justify having any stock at all, which means that the optimal policy is to not stock any rental units. Figure 6 illustrates such a policy by considering $c \in \{20, 40, 60, 80\}$ for the net revenue per rental to represent varying levels of profitability per rental, $N = 26$ weeks for a rental horizon, and $\lambda = 7$ for the mean demand. We observe that the optimal response to an increasing inventory loss probability is to initially increase the inventory of rental units until we reach a certain value of the loss probability p associated with the optimal number of rental units y^* to procure in the beginning of the rental horizon. As p continues to increase, the optimal number of rental units decreases and the optimal service level also appears to be non-increasing in the loss probability. Eventually, a loss probability $\hat{p}(c)$ is reached such that $y^* = 0$ for all $p \geq \hat{p}(c)$. Naturally, the optimal service level and $\hat{p}(c)$ increase with c because a dress with a higher net revenue per rental requires fewer rentals to be profitable.

6.2.2. Autocorrelated Demand Next, we relax the assumption of independent and identically distributed demand when generating random demand to demonstrate how advanced demand models can be used to represent the nuances of demand for fashion products. We model demand



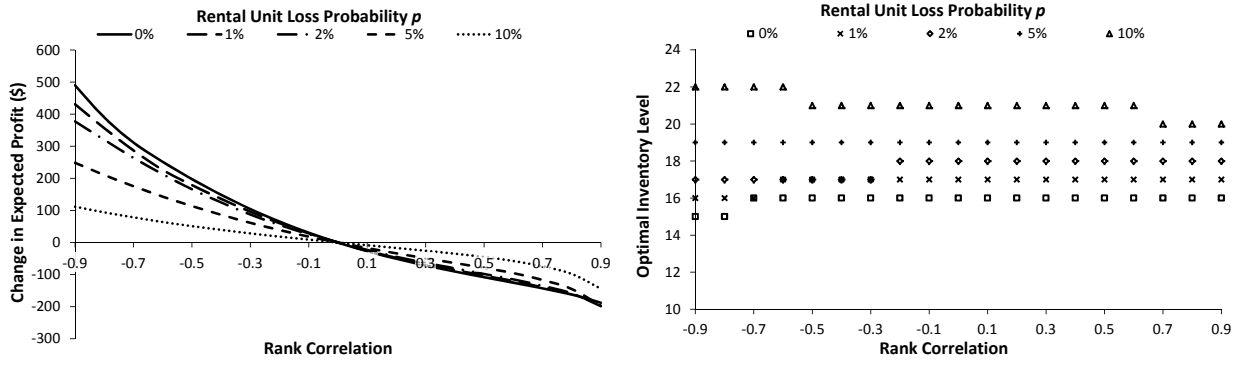
(a) Optimal initial inventory.

(b) Optimal service rate.

Figure 6 The optimal inventory level is non-monotonic in the rental unit loss probability p .

using a stationary first-order autocorrelated process in which the marginal distribution is Poisson with mean 7, and generate random demand by specifying the rank correlation as described in Biller and Ghosh (2006).

Figure 7a demonstrates that negative first-order autocorrelation improves upon the expected profit compared to the case of no rank autocorrelation, while positive first-order autocorrelation causes it to decrease. When demand exhibits negative first-order autocorrelation, the periodicity of demand and the constant two-period rental duration work together to improve the service rate for a given number of rental units. On the other hand, positive first-order autocorrelation reduces the ability of each rental unit to serve demand. Higher rental inventory loss dampens the effect of autocorrelation, as rental units that do not last for the entire horizon do not experience autocorrelation's effects as severely. For a system with no inventory loss ($p = 0\%$), rank autocorrelation of -0.5 increases the expected profit by \$197 (6.6%), and rank autocorrelation of 0.5 decreases the expected profit by \$108 (3.6%). For a system with $p = 5\%$, rank autocorrelation of -0.5 increases the expected profit by \$114 (6.8%), and rank autocorrelation of 0.5 decreases the expected profit by \$79 (4.8%). As shown in Figure 7b, the effect of autocorrelation on the optimal initial inventory level varies with the loss probability. When the loss probability is low (i.e., $p \leq 2\%$), the optimal policy is to add rental units as the rank correlation increases. However, the optimal initial inventory level decreases with the rank correlation for high loss probability (i.e., $p = 10\%$), reflecting the



(a) Optimal expected profit.

(b) Optimal inventory level.

Figure 7 Positive first-order rank autocorrelation reduces the expected profit, while negative autocorrelation increases the expected profit.

increased difficulty for each rental unit to profitable due to the combination of the loss probability and rank correlation.

6.3. Rental Unit Recirculation Rules

Different recirculation rules employed during the rental horizon may result in different numbers of units available near the horizon's end. We expect that the importance of the rental unit recirculation policy varies according to factors such as the horizon length, rental unit lifetime distribution, and demand characteristics. Of concern to us is a horizon that is short enough that some rental units are still functional by the end of the last time period but long enough that some rental units have already retired from the rental inventory during the season. In this section, we compare the even spread and static priority policy for the count-based model, omitting similar managerial insights and results for the best-first and worst-first rules of the condition-based model.

Executives at Rent the Runway indicate that the policy used in practice more closely resembles the static priority policy than the even spread policy. Out of convenience, dresses that have just returned from cleaning after a rental may be selected to satisfy the next rental. However, because individual units are not tracked, there may be an element of randomness in dress selection as workers select a dress to rent out. The goal of this section is to quantify the effect of using the even spread policy for rental unit recirculation over the static priority recirculation rule. For an adequate

representation of the role of the rate at which the loss probability is increasing, we consider the lifetime of a rental unit to be a discrete uniform random variable that takes values between 1 and $A_{max} \in \{10, 11, \dots, 20\}$ rentals. As before, we consider a rental horizon of 26 periods and a mean demand of 7 units, with all other parameters remaining the same.

The rental system illustrated in Figure 8 is only profitable when $A_{max} \geq 13$ for the even spread policy and when $A_{max} \geq 14$ for the static priority policy due to the costs incurred when rental units are lost. Consistent with Proposition 5, the even spread policy achieves a higher expected profit than the static priority policies. This performance difference can be explained by the nature of the even spread policy to delay the failure of rental units until later periods; thus, the even spread policy satisfies higher demand in later periods compared to the static priority recirculation policy. On the other hand, the static priority policy causes failures to occur earlier in the rental horizon, limiting the system's ability to meet higher quantities of demand in later periods.

The optimal number of rental units for the even spread, as well as the corresponding service rate, exceeds that of the static priority policy. Choosing the even spread policy instead of the static priority policy allows for rental units to be profitably added, thereby increasing the service rate. For example, the optimal initial inventory level is two units higher for even spread policy than the static priority policy when $A_{max} = 14$, and the service rate is 6.0 percentage points higher for the even spread policy.

7. Conclusion

As rental industries continue to grow in size and the scope of products rented, inventory management techniques that account for the complexities of rental systems become critical for achieving profitability and service goals. We develop a discrete-time rental model with random usage-based loss of inventory that also includes arbitrarily distributed customer demands and random rental durations, and identify structural properties for this model. The concavity of the expected profit function in the initial inventory of rental units is shown to hold for geometrically distributed rental unit lifetimes regardless of the rental unit recirculation rule. When rental unit lifetimes are generally distributed, we also show the concavity of the expected profit function in the initial inventory

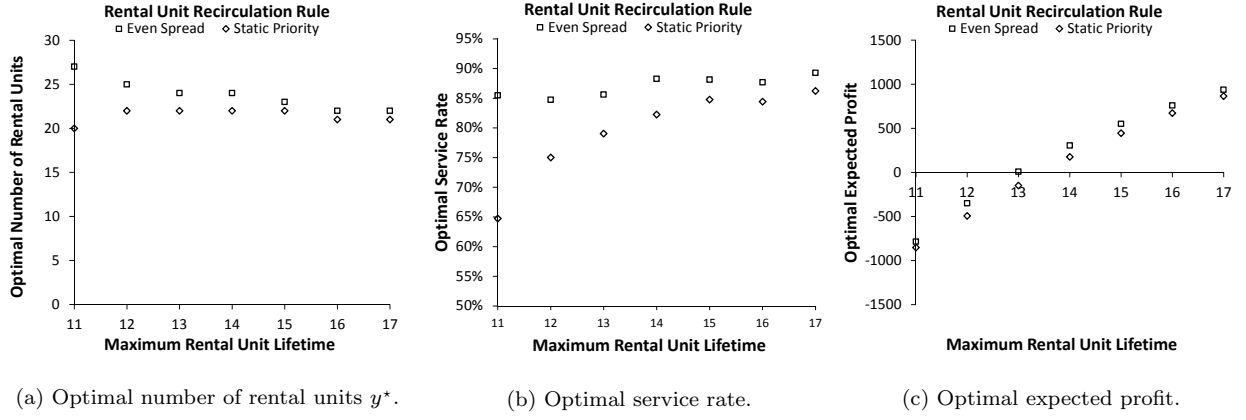


Figure 8 Choosing the even spread policy over the static priority policy increases the optimal initial inventory level, service rate, and expected profit.

of rental units for simple rental unit recirculation rules which are count-based or condition-based. We further demonstrate the optimality of the even spread policy in the count-based setting and the best-first policy in the condition-based setting to maximize the expected profit when the loss probability of each rental unit increases with the number of times it is rented.

Several important insights emerge from a numerical analysis of our rental inventory management solutions for a high-fashion dress rental business. First, we find that the possibility of inventory loss during the rental season can significantly affect profitability, even with a small probability of loss each time that a unit is rented. Choosing the number of rental units to procure in the beginning of the rental season by ignoring the effect of rental inventory loss can reduce the expected profit by 7%. Second, we examine how the optimal inventory policy responds to the increasing loss probability. We show that the optimal policy is to first procure additional rental units, then decrease the number of rental units to be procured and eventually procure zero rental units. Finally, we consider rental unit lifetime distributions with loss probabilities that are increasing in the number of rentals. For horizon lengths and lifetime distributions in which the recirculation rule affects the expected profit, we show that choosing the even spread policy allows for more inventory to be profitably obtained and can increase the service level by up to 6 percentage points.

Many categories of products available to rent, including dresses at Rent the Runway, present customers with the option to substitute if their first choice is unavailable. Thus, a future research direction in the study of rental inventory management is the case of multiple products with stock-out

based substitution. Potential future work also includes advance reservations, in-season reordering, and rental pricing.

Acknowledgments

The authors wish to thank Vijay Subramanian, Chief Analytics Officer of Rent the Runway, for his insights into Rent the Runway's operations. They also thank participants of David Simchi-Levi's MIT operations research seminar and Alan Scheller-Wolf for their helpful feedback.

References

- Adelman, D. 2008. A simple algebraic approximation to the Erlang loss system. *Operations Research Letters* **36**(4) 484 – 491.
- Baron, O., I. Hajizadeh, J. Milner. 2011. Now playing: DVD purchasing for a multilocation rental firm. *Manufacturing and Service Operations Management* **13**(2) 209 – 226.
- Biller, B., S. Ghosh. 2006. Multivariate input processes. *Handbooks in Operations Research and Management Science*, vol. 13. 123 – 153.
- Binkley, C. 2011. Fashion 101: Rent the Runway targets students. *Wall Street Journal* URL <http://online.wsj.com/article/SB10001424052748703806304576244952860660370.html>.
- Brown, M., N. R. Chaganty. 1983. On the first passage time distribution for a class of markov chains. *The Annals of Probability* **11**(4) 1000 – 1008.
- Cohen, M. A., W. P. Pierskalla, S. Nahmias. 1980. A dynamic inventory system with recycling. *Naval Research Logistics Quarterly* **27**(2) 289 – 296.
- Eisenmann, T. R., L. Winig. 2012. *Rent the Runway*. Harvard Business Publishing, Boston.
- Erlanger, S., M. De La Baume. 2009. French ideal of bicycle-sharing meets reality. *The New York Times* URL <http://www.nytimes.com/2009/10/31/world/europe/31bikes.html>.
- Gans, N., S. Savin. 2007. Pricing and capacity rationing for rentals with uncertain durations. *Management Science* **53**(3) 390 – 407.
- Green, L. V., P. J. Kolesar, J. Soares. 2001. Improving the SIPP approach for staffing service systems that have cyclic demands. *Operations Research* **49**(4) 549 – 564.

- Hampshire, R., W. Massey, Q. Wang. 2009. Dynamic pricing to control loss systems with quality of service targets. *Probability in the Engineering and Informational Sciences* **23**(2) 357 – 383.
- Harel, A. 1988. Sharp bounds and simple approximations for the erlang delay and loss formulas. *Management Science* **34**(8) 959 – 972.
- Huang, C. C., S. L. Brumelle, K. Sawaki, I. Vertinsky. 1977. Optimal control for multi-servers queueing systems under periodic review. *Naval Research Logistics Quarterly* **24**(1) 127 – 135.
- Jung, M., E.S. Lee. 1989. Numerical optimization of a queueing system by dynamic programming. *Journal of Mathematical Analysis and Applications* **141**(1) 84 – 93.
- Kleywegt, A. J., A. Shapiro, T. Homem-de Mello. 2002. The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization* **12**(2) 479 – 502.
- Levi, R., C. Shi. 2011. Revenue management of reusable resources with advanced reservations. Working Paper.
- Lindvall, T. 1992. *Lectures on the Coupling Method*. Wiley, New York.
- Mahajan, S., G. van Ryzin. 2001. Stocking retail assortments under dynamic consumer substitution. *Operations Research* **49**(3) 334 – 351.
- Muth, E. J. 1979. The reversibility property of production lines. *Management science* **25**(2) 152 – 158.
- Papier, F., U. W. Thonemann. 2008. Queuing models for sizing and structuring rental fleets. *Transportation Science* **42**(3) 302 – 317.
- Riordan, J. 1962. *Stochastic Service Systems*. Wiley, New York.
- Savin, S. V., M. A. Cohen, N. Gans, Z. Katalan. 2005. Capacity management in rental businesses with two customer bases. *Operations Research* **53**(4) 617 – 631.
- Shanthikumar, J. G., D. D. Yao. 1987. Optimal server allocation in a system of multi-server stations. *Management Science* **33**(9) 1173 – 1180.
- Shapiro, A. 2003. Monte carlo sampling methods. *Handbooks in Operations Research and Management Science*, vol. 10. 353 – 425.

- Tainiter, M. 1964. Some stochastic inventory models for rental situations. *Management Science* **11**(2) 316 – 326.
- Takács, L. 1962. *Introduction to the Theory of Queues*. Oxford University Press, New York.
- Tayur, S. 1993. Structural properties and a heuristic for kanban-controlled serial lines. *Management Science* **39**(11) 1347 – 1368.
- Whisler, W. D. 1967. A stochastic inventory model for rented equipment. *Management Science* **13**(9) 640 – 647.
- Wortham, J. 2009. Rent the Runway offers designer dresses in the Netflix model. *New York Times* URL <http://www.nytimes.com/2009/11/09/technology/09runway.html>.
- Yannetta, T. 2013. No kidding: Rent the Runway’s sample sale returns April 1st. URL http://ny.racked.com/archives/2013/03/26/rent_the_runway_4.php.
- Zhang, Y., M. L. Puterman, M. Nelson, D. Atkins. 2012. A simulation optimization approach to long-term care capacity planning. *Operations Research* **60**(2) 249 – 261.

Appendix. Proofs

Proof of Lemma 1. For notational convenience, we denote the first forward difference $L_n(y + 1, \xi) - L_n(y, \xi)$ by $\Delta L_n(y)$, indexing this difference to the time period n and the initial inventory of y rental units, with the objective of comparing the inventory system with y rental units to the system with $y + 1$ rental units. We also define $\mathbf{1}(\cdot)$ as the indicator function that takes the value of one if the argument is true and zero otherwise. In addition, $\Delta I_n(y)$, $\Delta R_n(y)$, and $\Delta W_n(y)$ indicate the effect of increasing the value of y by one on the available inventory to rent, the number of rentals, and the number of rental units returned, respectively; i.e., $\Delta I_n(y) = I_n(y + 1, \xi) - I_n(y, \xi)$, $\Delta R_n(y) = R_n(y + 1, \xi) - R_n(y, \xi)$ and $\Delta W_n(y) = W_n(y + 1, \xi) - W_n(y, \xi)$.

In this proof, we represent the effect of a unit increase in y on the total number of rentals in terms of the aggregate effect over all N periods; i.e., $\Delta \mathcal{R}(y, \xi) = \sum_{n=1}^N \Delta R_n(y) = \sum_{n=1}^N \Delta I_n(y) \mathbf{1}(d_n > I_n(y, \xi))$. We separate the summation term into two components — one for period n and another for periods $1, 2, \dots, n - 1$. We further substitute $\Delta I_n(y)$ with its counterpart from the state equations in (1) and $W_{n+1}(y, \xi) = \sum_{a=A_{min}}^{A_{max}} W_{a,n+1}(y, \xi)$ to obtain the following characterization for $\sum_{n=1}^N \Delta R_n(y)$:

$$\sum_{n=1}^N \Delta R_n(y) = \Delta I_N(y) \mathbf{1}(d_N > I_N(y, \xi)) + \sum_{n=1}^{N-1} \Delta R_n(y)$$

$$\begin{aligned}
&= \left(\Delta I_{N-1}(y) + \sum_{a=A_{\min}}^{A_{\max}} \Delta W_{a,N}(y) - \Delta R_{N-1}(y) \right) \mathbb{1}(d_N > I_N(y, \xi)) + \sum_{n=1}^{N-1} \Delta R_n(y) \\
&= \left(1 + \sum_{n=1}^N \sum_{a=A_{\min}}^{A_{\max}} \Delta W_{a,n}(y) \right) \mathbb{1}(d_N > I_N(y, \xi)) + \sum_{n=1}^{N-1} \Delta R_n(y) \mathbb{1}(d_N \leq I_N(y, \xi)) \\
&= \sum_{n=1}^N \left(1 + \sum_{t=1}^n \sum_{a=A_{\min}}^{A_{\max}} \Delta W_{a,n}(y) \right) \mathbb{1}(d_n > I_n(y, \xi)) \prod_{v=n+1}^N \mathbb{1}(d_v \leq I_v(y, \xi)).
\end{aligned}$$

In the case of experiencing at least one lost sale, this expression reduces to an equivalence between $\sum_{n=1}^N \Delta R_n(y)$ and $1 + \sum_{t=1}^u \Delta W_t(y)$ with $u = \max\{n \in \{1, 2, \dots, N\} : d_n > I_n(y, \xi)\}$. This expression further reduces to $1 + \sum_{t=1}^u \sum_{a=A_{\min}}^{A_{\max}} \Delta W_{a,u}(y)$ with $u = \max\{n \in \{1, 2, \dots, N\} : d_n > I_n(y, \xi)\}$ when the system experiences at least one lost sale. Thus, that $\mathbb{E}[\sum_{t=1}^n \sum_{a=A_{\min}}^{A_{\max}} W_{a,t}(y, \xi)]$ is concave and non-decreasing in y for $n = 1, 2, \dots, N$ is a sufficient condition for the expected number of rentals $\mathbb{E}[\mathcal{R}(y, \xi)]$ to be concave and non-decreasing in y , and for the expected number of lost sales $\mathbb{E}[\mathcal{L}(y, \xi)]$ to be convex and non-increasing in y . \square

Proof of Proposition 1. We use induction to show the satisfaction of the sufficiency condition in Lemma 1 for the concavity of $\mathbb{E}[\mathcal{R}(y, \xi)]$ in the initial inventory of y rental units. First, we note that there are no rental units returned in period 1. Therefore, the sufficiency condition is trivially satisfied for $n = 1$. We next assume that $\mathbb{E}[\sum_{i=1}^t R_i(y, \xi)]$ is concave in y for $t = 1, 2, \dots, n-1$ with $n \geq 2$. What is important to recognize here is that the expected number of returns by period n , $\mathbb{E}[\sum_{t=1}^n \sum_{a=A_{\min}}^{A_{\max}} W_{a,t}(y, \xi)]$ can be written as $(1-p) \sum_{a=A_{\min}}^{A_{\max}} h(a) \sum_{t=1}^{n-a} \mathbb{E}[R_t(y, \xi)]$. Therefore, $\mathbb{E}[\sum_{t=1}^n \sum_{a=A_{\min}}^{A_{\max}} W_{a,t}(y, \xi)]$ is a concave function of y , and it follows from Lemma 1 that $\mathbb{E}[\mathcal{R}(y, \xi)]$ is concave and non-decreasing in y .

Building on this structural property of the expected number of rentals, we show the concavity of the expected profit function in two steps: (1) The revenue acquired from all rentals is given by $\sum_{a=A_{\min}}^{A_{\max}} r_a \sum_{n=1}^N R_{a,n}(y, \xi)$. Because the duration of a rental that begins in period n is independent of the durations of any rentals that begin in periods $1, 2, \dots, n-1$, it holds that $\mathbb{E}[R_{a,n}(y, \xi)] = h(a) \mathbb{E}[R_n(y, \xi)]$. Consequently, we obtain the expected rental revenue as $\sum_{a=A_{\min}}^{A_{\max}} r_a h(a) \mathbb{E}[\mathcal{R}(y, \xi)]$. (2) To account for different salvage values of lost rental units, we consider the expectation of the difference $\sum_{n=1}^N R_n(y, \xi) - \sum_{n=1}^{N+A_{\max}} \sum_{a=A_{\min}}^{A_{\max}} W_{a,n}(y, \xi)$. Because $\sum_{a=A_{\min}}^{A_{\max}} h(a) = 1$, the expected total number of returns $\mathbb{E}[\sum_{t=1}^{N+A_{\max}} \sum_{a=A_{\min}}^{A_{\max}} W_{a,t}(y, \xi)] := (1-p) \sum_{a=A_{\min}}^{A_{\max}} h(a) \sum_{t=1}^{n-a} \mathbb{E}[R_t(y, \xi)]$ can be rewritten as $(1-p) \sum_{t=1}^N \mathbb{E}[R_t(y, \xi)]$. Consequently, we conclude the expected profit function $\pi(y) := -s^g y - c \mathbb{E}[\sum_{n=1}^N D_n] +$

$(\sum_{a=A_{min}}^{A_{max}} r_a h(a) + c - p(s^b - s^g))\mathbb{E}[\mathcal{R}(y, \xi)]$ to be concave in the initial inventory of y rental units for any rental unit recirculation rule. \square

Proof of Proposition 2. We first compare forward differences of systems with y rental units and $y + 1$ rental units to show that $\mathbb{E}[\Delta \mathcal{R}^{SP}(y, \xi)] \geq \mathbb{E}[\Delta \mathcal{R}^{SP}(y + 1, \xi)] \geq 0$ for $y \geq 0$. Since all rental units are unconditionally stochastically equivalent in terms of their lifetimes, without loss of generality we focus on a marginal unit of inventory that has the lowest priority for when rental units are assigned to demands. By the definition of the static priority policy, the additional lowest-priority unit will not change the allocation of any other rental units, implying that $\Delta I_n^{SP}(y, \xi(y)) \geq 0$. Thus, the forward difference of the state equations (1) for the number of rentals in a period $n + 1$ is

$$\Delta R_{n+1}^{SP}(y, \xi(y)) = \begin{cases} 0 & \text{if } d_{n+1} \leq I_{n+1}^{SP}(y, \xi(y)) \\ \Delta I_{n+1}^{SP}(y, \xi(y)) & \text{if } d_{n+1} > I_{n+1}^{SP}(y, \xi(y)). \end{cases}$$

Next, we show that $\sum_{t=1}^n \Delta R_t^{SP}(y, \xi(y)) \geq \sum_{t=1}^n \Delta R_t^{SP}(y + 1, \xi(y + 1))$ by induction. For period 1, we observe that $\Delta R_1^{SP}(y, \xi(y)) \geq \Delta R_1^{SP}(y + 1, \xi(y + 1))$ due to the forward difference of the state equations with $I_1^\gamma(y, \xi(y)) = y$ and $\Delta I_1^\gamma(y, \xi(y)) = 1$.

Now, we assume that $\sum_{t=1}^{n-1} \Delta R_t^{SP}(y, \xi(y)) \geq \sum_{t=1}^{n-1} \Delta R_t^{SP}(y + 1, \xi(y + 1))$ for some period n . We account for the following cases:

1. $\sum_{t=1}^{n-1} \Delta R_t^{SP}(y, \xi(y)) = I'$ and $I' > \sum_{t=1}^{n-1} \Delta R_t^{SP}(y + 1, \xi(y + 1))$: In this case, the additional rental unit is lost before period n for the system with y rental units but is not lost for the system with $y + 1$ rental units. Because $\Delta R_n^{SP}(y + 1, \xi(y + 1)) \leq 1$ by the state equations, $\sum_{t=1}^n \Delta R_t^{SP}(y, \xi(y)) \geq \sum_{t=1}^n \Delta R_t^{SP}(y + 1, \xi(y + 1))$.
2. $I' > \sum_{t=1}^{n-1} \Delta R_t^{SP}(y, \xi(y))$ and $\sum_{t=1}^{n-1} \Delta R_t^{SP}(y, \xi(y)) > \sum_{t=1}^{n-1} \Delta R_t^{SP}(y + 1, \xi(y + 1))$: As in the previous case, $\sum_{t=1}^n \Delta R_t^{SP}(y, \xi(y)) \geq \sum_{t=1}^n \Delta R_t^{SP}(y + 1, \xi(y + 1))$ because $\Delta R_n^{SP}(y + 1, \xi(y + 1)) \leq 1$.
3. $\sum_{t=1}^{n-1} \Delta R_t^{SP}(y, \xi(y)) = \sum_{t=1}^{n-1} \Delta R_t^{SP}(y + 1, \xi(y + 1))$: If $\sum_{t=1}^{n-1} \Delta R_t^{SP}(y, \xi(y)) = I'$ and $\sum_{t=1}^{n-1} \Delta R_t^{SP}(y + 1, \xi(y + 1)) = I'$, then the additional unit is unavailable for either system, and $\sum_{t=1}^n \Delta R_t^{SP}(y, \xi(y)) = \sum_{t=1}^n \Delta R_t^{SP}(y + 1, \xi(y + 1)) = I'$. Otherwise, it suffices to show that $\Delta I_n^{SP}(y, \xi(y)) \geq \Delta I_n^{SP}(y + 1, \xi(y + 1))$. Let $i = \sum_{t=1}^{n-1} \Delta R_t^{SP}(y, \xi(y))$. By the inductive hypothesis, the i th rental of the additional unit occurred no later for the system with $y + 1$ units than the system with $y + 2$ units. Thus, after a rental duration of A'_i periods, the additional unit returns to become available in an earlier period for the system with $y + 1$ units than with $y + 2$ units, which implies that $\Delta I_n^{SP}(y, \xi(y)) \geq \Delta I_n^{SP}(y + 1, \xi(y + 1))$.

Having shown that the change in the total number of rentals up to period n from one additional rental unit is non-increasing in y on coupled sample paths (i.e., $\sum_{t=1}^n \Delta R_t^{SP}(y, \xi(y)) \geq \sum_{t=1}^n \Delta R_t^{SP}(y+1, \xi(y+1))$), the property that $\mathbb{E}[\sum_{t=1}^n \Delta R_t^{SP}(y, \xi(y))] \geq \mathbb{E}[\sum_{t=1}^n \Delta R_t^{SP}(y+1, \xi(y+1))]$ follows because the lifetime and rental durations of the additional unit for the two systems being compared are independent and identically distributed. Naturally, this property implies that $\mathbb{E}[\Delta \mathcal{R}_t^{SP}(y, \xi(y))] \geq \mathbb{E}[\Delta \mathcal{R}_t^{SP}(y+1, \xi(y+1))]$.

Finally, we show that $\pi^{SP}(y)$ is concave and non-decreasing in y ; i.e., $\Delta \pi^{SP}(y) \geq \Delta \pi^{SP}(y+1)$ for $y \geq 0$. To determine when the concavity of the expected number of rentals in the initial inventory level implies the concavity of the expected profit, we must analyze whether the $(y+1)$ st unit for a system with y units and the $(y+2)$ nd unit for a system with $y+1$ units are lost. To do so, we compare $\Delta \Pi^{SP}(y, \xi(y))$ to $\Delta \Pi^{SP}(y+1, \xi(y+1))$ with the profit on a sample path written as

$$\Delta \Pi^{SP}(y, \xi(y)) = \sum_{a=A_{min}}^{A_{max}} (r_a + c) \Delta \mathcal{R}_a^{SP}(y, \xi(y)) - s^g - (s^b - s^g) \mathbf{1} \left\{ \sum_{t=1}^n \Delta \mathcal{R}_t^{SP}(y, \xi(y)) \geq \mathbf{I}' \right\}. \quad (2)$$

We immediately observe that $\Delta \Pi^{SP}(y, \xi(y)) \geq \Delta \Pi^{SP}(y+1, \xi(y+1))$ if $\Delta \mathcal{R}^{SP}(y, \xi(y)) = \Delta \mathcal{R}^{SP}(y+1, \xi(y+1))$. If $\Delta \mathcal{R}^{SP}(y, \xi(y)) > \Delta \mathcal{R}^{SP}(y+1, \xi(y+1))$, then the expected effect for the system with y rental units of the $\Delta \mathcal{R}^{SP}(y, \xi(y)) - \Delta \mathcal{R}^{SP}(y+1, \xi(y+1))$ extra rentals of an additional unit must be non-negative. Considering the loss probability for the extra $\Delta \mathcal{R}^{SP}(y, \xi(y)) - \Delta \mathcal{R}^{SP}(y+1, \xi(y+1))$ of the system with y rental units to which the $(y+1)$ st unit is being added, it suffices that $\sum_{a=A_{min}}^{A_{max}} r_a h(a) + c \geq (s^b - s^g) \ell_i$, $i \geq N/A_{min}$ for $\Delta \pi^{SP}(y) \geq \Delta \pi^{SP}(y+1)$ to hold, completing the proof. \square

Proof of Proposition 3. The goal is to show that the forward difference of the number of rentals is decreasing in y when the rental system follows the even spread recirculation rule. To do so, we define a restricted allocation of rental units to a reduced number of demands $d'_n(y, \xi(y)) = d_n - R_{n,y+1}^{ES}(y+1, \xi(y+1))$, $n = 1, 2, \dots, N$; i.e., any unit demand served by the $(y+1)$ st rental unit for the system with $y+1$ units is not allowed to be satisfied by any rental unit when the system has only y units. We denote the number of units rented in period n with this restricted allocation by $R_n^{ESr}(y, \xi(y))$. By definition of the restricted allocation, we immediately recognize the equivalence between $R_{n,m}^{ESr}(y, \xi(y))$ and $R_{n,m}^{ES}(y+1, \xi(y+1))$ for $n = 1, 2, \dots, N$ and $m = 1, 2, \dots, y$. In other words, the satisfaction of the i th demand by the rental unit $m \in \{1, 2, \dots, y\}$ occurs in the same period for the restricted system with y units and the unrestricted system with $y+1$ rental units. Thus, the addition of one rental unit with lifetime \mathbf{I}' and random durations $\{a'_1, a'_2, \dots\}$ to the system with y rental units has the same effect for the system with $y+1$ rental units; i.e., $\Delta R_n^{ESr}(y, \xi(y)) = \Delta R_n^{ES}(y+1, \xi(y+1))$. Hence, we also have $\sum_{n=1}^N \Delta R_n^{ESr}(y, \xi(y)) = \sum_{n=1}^N \Delta R_n^{ES}(y+1, \xi(y+1))$.

Next, we show that $\sum_{n=1}^N \Delta R_n^{ES}(y, \xi(y)) \geq \sum_{n=1}^N \Delta R_n^{ESr}(y, \xi(y))$ by removing the allocation restrictions so that the recirculation rule obeys the even spread policy and that state equation $R_n^\gamma(y, \xi) = \min\{d_n, I_n^\gamma(y, \xi)\}$ in each period n . Specifically, we show that $\sum_{t=1}^n \Delta R_t^{ES}(y, \xi(y))$ is non-decreasing in $\{d_1, d_2, \dots, d_n\}$, which implies that $\sum_{t=1}^n \Delta R_t^{ES}(y, \xi(y)) \geq \sum_{t=1}^n \Delta R_t^{ESr}(y, \xi(y))$ holds because $d_t \geq d'_t(y, \xi(y))$ for $t = 1, 2, \dots, n$. In period 1, the inductive hypothesis is true by the state equations because $\Delta R_1^{ES}(y, \xi(y)) = \mathbf{1}\{d_1 > y\}$ is non-decreasing in d_1 . For some period n with $\sum_{t=1}^{n-1} \Delta R_t^{ES}(y, \xi(y)) \geq \sum_{t=1}^{n-1} \Delta R_t^{ESr}(y, \xi(y))$, we consider the following cases:

1. $\Delta R_n^{ES}(y, \xi(y)) \geq \Delta R_n^{ESr}(y, \xi(y))$. The result follows immediately.
2. $\Delta R_n^{ES}(y, \xi(y)) < \Delta R_n^{ESr}(y, \xi(y))$. We must show that $\sum_{t=1}^{n-1} \Delta R_t^{ES}(y, \xi(y)) + \Delta R_n^{ES}(y, \xi(y)) \geq \sum_{t=1}^{n-1} \Delta R_t^{ESr}(y, \xi(y)) + \Delta R_n^{ESr}(y, \xi(y))$. To do so, we will demonstrate that $\Delta I_n^{ESr}(y, \xi(y)) - \Delta I_n^{ES}(y, \xi(y)) \leq \sum_{t=1}^{n-1} \Delta R_t^{ES}(y, \xi(y)) - \sum_{t=1}^{n-1} \Delta R_t^{ESr}(y, \xi(y))$. Each unit difference comprising $\Delta I_n^{ESr}(y, \xi(y)) - \Delta I_n^{ES}(y, \xi(y))$ can only occur when $\sum_{t=1}^{n-1} \Delta R_{t,m}^{ES}(y, \xi(y)) > \sum_{t=1}^{n-1} \Delta R_{t,m}^{ESr}(y, \xi(y))$ for some rental unit m . Therefore, $\sum_{t=1}^{n-1} \Delta R_t^{ES}(y, \xi(y)) + \Delta R_n^{ES}(y, \xi(y)) \geq \sum_{t=1}^{n-1} \Delta R_t^{ESr}(y, \xi(y)) + \Delta R_n^{ESr}(y, \xi(y))$.

Thus, the inductive hypothesis holds to show that the value of an additional rental unit cannot decrease with the conversion of the restricted allocation to an unrestricted allocation while maintaining the even spread policy; i.e., $\sum_{n=1}^N \Delta R_n^{ES}(y, \xi(y)) \geq \sum_{n=1}^N \Delta R_n^{ESr}(y, \xi(y))$. Since it was also shown above that $\mathbb{E}[\sum_{n=1}^N \Delta R_n^{ESr}(y)] \geq \mathbb{E}[\sum_{n=1}^N \Delta R_n^{ES}(y+1)]$, we see that $\mathbb{E}[\mathcal{R}^{ES}(y, \xi)]$ is concave and non-decreasing in y for the even spread recirculation rule. \square

Proof of Proposition 4. We start with the definition of $Q_i(y, \xi(y)) := \sum_{m=1}^y \mathbf{1}\{\sum_{n=1}^N R_{n,m}^{ES}(y, \xi(y)) = i\}$ as the number of rental units that are rented exactly i times during the rental horizon. Because rental unit lifetimes are independent and identically distributed, we can also express the expected number of returns $\mathbb{E}[\sum_{n=1}^{N+A_{max}} W_n^{ES}(y, \xi)]$ as $\sum_{i=1}^N (1 - \ell_i) \mathbb{E}[Q_i(y, \xi)]$. The goal is to demonstrate the concavity of $\mathbb{E}[\sum_{n=1}^{N+A_{max}} W_n^{ES}(y, \xi)]$ in y . We do this by investigating whether the forward difference of the expected number of returns is non-increasing in y ; i.e., $\Delta \mathbb{E}[\sum_{n=1}^{N+A_{max}} W_n^{ES}(y, \xi)] \geq \Delta \mathbb{E}[\sum_{n=1}^{N+A_{max}} W_n^{ES}(y+1, \xi)]$, which can be equivalently written as follows for $y \geq 0$:

$$\sum_{i=1}^N (1 - \ell_i) \Delta \mathbb{E}[Q_i(y, \xi)] \geq \sum_{i=1}^N (1 - \ell_i) \Delta \mathbb{E}[Q_i(y+1, \xi)].$$

To prove this condition, we next consider a sample path in which both the rental unit lifetimes and the duration of each demand that they satisfy have already been revealed for rental units $1, 2, \dots, y$. However,

the lifetimes for rental units $y + 1$ and $y + 2$, which are used in the analysis of forward differences, and the duration of all demands satisfied by those rental units are random. Since the lifetimes of rental units $y + 1$ and $y + 2$ are stochastically equivalent, it holds that $\Delta\mathbb{E}[Q_i(y, \xi)] \geq \Delta\mathbb{E}[Q_i(y + 1, \xi)]$ for $i = 1, 2, \dots, j$. This condition, however, does not necessarily imply that $\Delta\mathbb{E}[Q_i(y, \xi)] \geq \Delta\mathbb{E}[Q_i(y + 1, \xi)]$ for $i > j$ as the addition of a rental unit might cause other rental units to be allocated fewer times to satisfy the demand (in expectation over random rental durations) and $\Delta\mathbb{E}[Q_i(y, \xi)]$ takes a negative value for some i .

To account for this scenario in which $\sum_{i>j} \Delta\mathbb{E}[Q_i(y + 1, \xi)] > \sum_{i>j} \Delta\mathbb{E}[Q_i(y, \xi)]$ holds for some j , we build on Proposition 3 and the property of the increasing loss probability. Specifically, it follows from Proposition 3 that $\sum_{i=1}^j \Delta\mathbb{E}[Q_i(y, \xi)] + \sum_{i>j} \Delta\mathbb{E}[Q_i(y, \xi)] \geq \sum_{i=1}^j \Delta\mathbb{E}[Q_i(y + 1, \xi)] + \sum_{i>j} \Delta\mathbb{E}[Q_i(y + 1, \xi)]$; i.e., $\sum_{i \geq 1}^j \Delta\mathbb{E}[Q_i(y, \xi)] - \sum_{i \geq 1}^j \Delta\mathbb{E}[Q_i(y + 1, \xi)] \geq \sum_{i>j} \Delta\mathbb{E}[Q_i(y + 1, \xi)] - \sum_{i>j} \Delta\mathbb{E}[Q_i(y, \xi)]$. It also follows from the increasing loss probability property that ℓ_i is non-decreasing in i . Consequently, we obtain $\sum_{i=1}^j (1 - \ell_i) \Delta\mathbb{E}[Q_i(y, \xi)] \geq \sum_{i=1}^j (1 - \ell_i) \Delta\mathbb{E}[Q_i(y + 1, \xi)]$. Therefore, for the rental unit lifetime distribution with an increasing failure rate, the expected number of returns is concave in y , resulting in the concavity of the expected profit function in y under the assumption that $\sum_{a=A_{min}}^{A_{max}} r_a h(a) + c \geq (s^b - s^g) \ell_i$. \square

Proof of Proposition 5. We first focus on the proof of Step 5 listed in the text leading up to the proposition, and begin by showing that $\mathcal{R}^S(\xi_{(1)}, \xi_{(2)}) + \mathcal{R}^S(\xi_{(2)}, \xi_{(1)}) \geq \mathcal{R}^V(\xi_{(1)}, \xi_{(2)}) + \mathcal{R}^V(\xi_{(2)}, \xi_{(1)})$. First, by the increasing failure rate property, we note that for a given value of $\eta_{(1)}$ or $\eta_{(2)}$, rental unit i will have a remaining lifetime that is at least as long as that of rental unit j .

Next, we show that $\mathcal{R}^S(\xi_{(1)}, \xi_{(2)}) \geq \mathcal{R}^V(\xi_{(2)}, \xi_{(1)})$ and $\mathcal{R}^S(\xi_{(2)}, \xi_{(1)}) \geq \mathcal{R}^V(\xi_{(1)}, \xi_{(2)})$. We consider each inequality separately:

- $\mathcal{R}^S(\xi_{(1)}, \xi_{(2)}) \geq \mathcal{R}^V(\xi_{(2)}, \xi_{(1)})$: In this case, we compare the loss period of rental unit i defined by $\xi_{(1)}$ in the switched allocation to rental unit j defined by $\xi_{(1)}$ in the violating allocation. Due to sample path coupling, unit j has a remaining lifetime in the violating allocation that is less than or equal to that of unit i in the switched allocation. By induction, we can show that $\sum_{t=1}^n R_t^V(\xi_{(1)}, \xi_{(2)}) \leq \sum_{t=1}^n R_t^S(\xi_{(1)}, \xi_{(2)})$ for $n = 1, 2, \dots, N$. We omit this induction argument for its similarity to that of Proposition 3.

- $\mathcal{R}^S(\xi_{(2)}, \xi_{(1)}) \geq \mathcal{R}^V(\xi_{(1)}, \xi_{(2)})$: In this case, we compare the loss period of rental unit i defined by $\xi_{(2)}$ in the switched allocation to rental unit j defined by $\xi_{(2)}$ in the violating allocation. Again by sample path coupling, unit j has a remaining lifetime in the violating allocation that is less than or equal to that of unit i

in the switched allocation. As in the previous case, we can show that $\sum_{t=1}^n R_t^V(\xi_{(2)}, \xi_{(1)}) \leq \sum_{t=1}^n R_t^S(\xi_{(2)}, \xi_{(1)})$ for $n = 1, 2, \dots, N$.

Because rental unit lifetimes and durations are independent and identically distributed, the inequality $\mathcal{R}^S(\xi_{(1)}, \xi_{(2)}) + \mathcal{R}^S(\xi_{(2)}, \xi_{(1)}) \geq \mathcal{R}^V(\xi_{(1)}, \xi_{(2)}) + \mathcal{R}^V(\xi_{(2)}, \xi_{(1)})$ implies that $\mathbb{E}[\mathcal{R}^S(y)] \geq \mathbb{E}[\mathcal{R}^V(y)]$. Thus, under the assumption that $\sum_{a=A_{min}}^{A_{max}} r_a h(a) + c \geq (s^b - s^g)\ell_i$, the even spread allocation policy maximizes the expected profit. \square

Proof of Proposition 6. The proof proceeds analogously to Propositions 3 and 4 by analyzing the marginal effect of an additional rental unit in a sample path coupling framework for which \tilde{s}_i is the state of the additional unit after it has been rented i times. We analyze the best-first policy denoted by superscript BF but note that the same logic applies for the worst-first policy. We rely on the reasoning of Proposition 3 but must confirm the inductive argument for the condition-based model. Specifically, we define a restricted allocation BFr analogous to that of Proposition 3 and show that the change in the number of rentals is non-decreasing as the allocation restrictions are removed. Based on the definition of s_{mi} , the period in which the i th allocation of a rental unit m occurs is non-decreasing as the restrictions are relaxed. As each demand restriction is relaxed, $\sum_{n=1}^N R_n(y)$ either remains the same or increases by one. Thus, the expected number of rentals can only increase with the conversion of the restricted allocation to an unrestricted allocation while maintaining the best-first policy, and $\sum_{n=1}^N \Delta R_n^{BF}(y) \geq \sum_{n=1}^N \Delta R_n^{BFr}(y)$. Since it also holds that $\mathbb{E}[\sum_{n=1}^N \Delta R_n^{BFr}(y)] \geq \mathbb{E}[\sum_{n=1}^N \Delta R_n^{BF}(y+1)]$, $\mathbb{E}[\mathcal{R}^{BF}(y, \xi)]$ is concave in y for the best-first (and worst-first) policies. Using the assumption provided above about the transition probability matrix, the argument presented in Proposition 2 holds and the concavity of the expected number of rentals implies the concavity of the expected profit function. \square

Proof of Proposition 7. The proof proceeds as in Proposition 5 with the need to only modify the inverse probability mass function values $\eta_{(1)}$ and $\eta_{(2)}$ for the conditional lifetime distributions of the two rental units. For P that is totally positive of order 2 and given either $\eta_{(1)}$ or $\eta_{(2)}$, rental unit i will always have a longer remaining lifetime (i.e., the number of possible allocations after period $n-1$) than rental unit j . The remainder of the proof follows as in Proposition 5. Thus, with the assumption that $\sum_{a=A_{min}}^{A_{max}} r_a h(a) + c \geq (s^b - s^g)P(i, S)$ for $i = 1, \dots, S-1$, the best-first recirculation policy maximizes the expected profit. \square