

Now Playing: DVD Purchasing for a Multilocation Rental Firm

Opher Baron, Iman Hajizadeh, Joseph Milner

Rotman School of Management, University of Toronto, Toronto, Ontario M5S 3E6, Canada
{opher.baron@rotman.utoronto.ca, i.hajizadeh@utoronto.ca, milner@rotman.utoronto.ca}

This paper studies the problem of purchasing and allocating copies of movies to multiple stores of a movie rental chain. A unique characteristic of this problem is the return process of rented movies. We formulate this problem for new movies as a newsvendor-like problem with multiple rental opportunities for each copy. We provide demand and return forecasts at the store-day level based on comparable movies. We estimate the parameters of various demand and return models using an iterative maximum-likelihood estimation and Bayesian estimation via Markov chain Monte Carlo simulation. Test results on data from a large movie rental firm reveal systematic underbuying of movies purchased through revenue-sharing contracts and overbuying of movies purchased through standard (nonrevenue-sharing) ones. For the movies considered, our model estimates an increase in the average profit per title for new movies by 15.5% and 2.5% for revenue sharing and standard titles, respectively. We discuss the implications of revenue sharing on the profitability of the rental firm.

Key words: service operations; supply chain management; inventory theory and control

History: Received: April 29, 2008; accepted: July 16, 2010. Published online in *Articles in Advance* December 30, 2010.

1. Introduction

The \$24 billion home entertainment industry in 2007 consisted of two major parts, movie sales (\$16 billion) and movie rentals (\$8 billion). Consumers spent, on average, about three times as much money buying and renting movies than in purchasing tickets at theater box offices (EMA 2008). Although movie sales have increased steadily at an average annual rate of 11% since 1990, the movie rental industry has remained almost the same size. However, its constant size does not imply that the industry is in steady state. In fact, the movie sales and rental industry has undergone dramatic technological changes affecting all aspects of the industry during the last 15 years.

Introduced in 1997, DVDs have by far surpassed traditional video cassettes in both sales and rentals. In 2007, DVDs accounted for 99% of rentals and movies sold (EMA 2008). This technology may soon be supplanted by high-definition DVDs. Also, emerging technologies such as Internet movie downloading, video on demand, and self-destructing discs, as well as innovative business models such as rental through the mail (e.g., Netflix) threaten traditional business models. Because of these changes, movie rental firms are under increasing pressure to reduce costs and increase efficiency.

We use data from a multistore movie rental firm to determine the number of copies of a newly released movie to place in each of its stores. A number of

factors affect this decision, including estimates of the uncertain demand; the return process, i.e., the process by which copies are returned to the firm; revenues received and costs incurred to purchase copies; and restrictions on the number of copies the firm can purchase. The latter two points are directly related to the contract by which the firm purchases its movies. Depending on the movie and studio (the supplier), movies may either be purchased outright (a standard contract) or obtained at a significant discount in exchange for a share of rental revenue (a revenue-sharing contract).

Historically, the firm purchased movies under standard and revenue-sharing contracts. Further, the studios fluctuated between both types of agreements several times over the last few years. Because of the difference in the terms, the number of times a copy has to be rented to cover its purchase cost, referred to as the break-even rentals per copy, differs between these purchase contracts. This break-even point drives all purchasing decisions. Managers at the rental firm said that the firm's break-even rentals per copy are three and one for standard and revenue-sharing contracts, respectively. Given a customer rental price of \$5 and a typical 40% revenue sharing with the studio, this would imply a \$15 purchase price net of any salvage value under standard contracts and a \$3 purchase price under revenue sharing. These values are in line with publicly available contract terms

(e.g., Rentrak 2008). We test robustness of these terms in §5. Further, the firm is restricted in the number of copies it purchases under a revenue-sharing contract. Managers at the rental firm confirmed that these constraints were binding for their purchases.

This paper has three main contributions. First, we formulate and solve the stochastic optimization problem faced by the firm to purchase inventory for its multiple stores that rent units over multiple periods. We note that the problem can be solved using a Lagrangian approach and, except for a constraint on the number purchased, is separable by store. We show that under a reasonable assumption on the return process, the problem may be solved through a greedy approach.

Second, we propose and empirically test several demand and return estimation models on data provided by the movie rental chain. Our data set consists of the number of copies allocated and the rental transactions (rentals and returns) for 52 movies at 450 stores for the first 27 rental days. The 52 movies in the data set are 20 new releases (10 revenue-sharing titles and 10 standard titles) and for each title, one or two comparable titles that are used to estimate demand and returns for the new movies. These movies were chosen by the rental firm from among numerous titles. In total there are over 9.5 million rental transactions in our database. As detailed below, for each movie we estimate the demand and return process for each store and day. As such, our data are aggregated by day, so that each movie consists of 24,300 data points. The data provided were relatively clean, especially for the higher-demand movies. However, data cleaning was necessary to adjust for rare cases of missing data, negative rentals, and sales of copies within the first 27 days. Some data have been disguised for reasons of confidentiality.

The main challenge in estimating the demand is that the observed demand, i.e., rents, is often censored (when there are no movies remaining in the store at the end of the day). Further, demand is autocorrelated and nonidentically distributed over the days in the month, and correlated across stores. Thus, our demand estimation models extend similar censoring models used in the OM literature to include variations in demand by store and day. A problem in estimating the returns process is that the data for this process only provides the number of copies returned on a given day and not the duration of the rental period. Thus we estimate the return process by accounting for inventory flows into and out of each store. Using these estimates and expert forecasting opinions, we use data from all of the stores simultaneously to forecast the inventory availability and the demand at each store on each day of the

planning horizon. We emphasize that we do not forecast individual movie demand based on attributes of the movies, e.g., a movie's director or actors. Rather, we transform experts' forecasts using inventory data from comparable movies across all the chain's stores to improve the purchase and allocation of movies to these stores.

Our third contribution is an examination of how standard and revenue-sharing contracts are used in practice in the movie rental supply chain. Previous research indicates that revenue-sharing agreements benefit supply chains (Dana and Spier 2001). For the standard contract titles, we show that the firm generally purchases too many copies of each movie. By purchasing the optimal number of copies for each store, the firm can increase its profits modestly, by approximately 2.5%. By reallocating the number of copies they purchase, they can achieve a profit improvement of 1.1%. This indicates that the profit function is flat near the optimal solution, and that by combining expert opinion with previous rental data, we can improve results across the chain. In contrast, we show that for the revenue-sharing titles, the firm would want to purchase additional copies, increasing average profit per title by 15.5%. However, the constraint on the purchase quantity for the revenue-sharing contracts restricts the firm from doing so. Our optimization achieves a 2.1% improvement under this constraint on average. The difference of these values measures the loss to the firm resulting from the lack of coordination in the supply chain. We discuss this point in our conclusions.

The remainder of the paper is organized as follows. A brief review of the related literature is presented in §2. We model the purchase and allocation decisions for the rental firm in §3. We propose and test several demand and return models in §4. In §5 we compare our model's results to the current practice of the movie rental firm. We also consider how alternate demand and returns estimates perform and conduct sensitivity analysis on the break-even returns per copy. Finally, in §6 we make some observations on the implications of our results and propose future research directions.

2. Literature Review

Analysis of the movie rental industry has recently become a subject of interest in the operations management literature. Lehmann and Weinberg (2000) study the industry from the studio's point of view. They focus on the optimal release times through sequential distribution channels with sales cannibalization (e.g., theaters and rental companies). Pasternack and Drezner (1999) focus on the purchasing problem from the rental firm's point of view. Based on the demand

pattern, they divide the lifetime of a movie into three phases (the first 30 days, the next t periods, and the remainder of time). Tang and Deo (2008) investigate the impact of rental duration on the stocking level, rental price, and retailer's profit. Our work differs from these papers in that they assume some aggregate demand pattern for a rental store, whereas we investigate several demand patterns empirically for a rental chain at a store-day level. Then, given a forecast based on data, we consider the allocation to stores alongside the purchase decision. Moreover, we test our purchase and allocation decisions on real data for a rental chain.

Much of the research in the movie rental industry focuses on designing optimal contracts; see, for example, Cachon and Lariviere (2005). For example, using evidence from this industry, Dana and Spier (2001) prove that revenue sharing successfully integrates a supply chain with intrabrand competition among downstream firms. Gerchak et al. (2006) provide evidence that, in addition to quantity, any contract between studios and rental chains should focus on the shelf-retention time of movies. They propose the addition of a license fee or subsidy to the contract to coordinate the chain when considering shelf retention. Mortimer (2008) provides an extensive empirical analysis of the movie rental industry in the United States. Her regression analysis shows that revenue-sharing contracts have a small positive effect on retailer's profit for popular titles and a small negative effect for less popular titles. In our numerical analysis we consider both standard and revenue-sharing contracts, taking the contract type as exogenous, and comment on the effectiveness of revenue-sharing contracts.

Other papers study a movie rental firm focusing mainly on asymptotic analysis of subscription-based rentals, e.g., the Netflix model. Bassamboo and Randhawa (2007) study the dynamic allocation of new releases to customers that are divided into two segments based on their rental time distribution (slow, fast). Bassamboo et al. (2009) extend the analysis to multiple customer segments focusing on the asymptotic behavior of the usage process. Randhawa and Kumar (2008) show that under some demand functions, subscription-based rental services provide superior profit for the rental firm compared to pay-per-use ones, whereas no option is dominant in service quality, consumer surplus, and social welfare. The context of these papers differs greatly from ours.

Previous related work considers statistical estimation of demand from sales data in the presence of stockouts. The importance of sales as censored demand data for the newsvendor problem was highlighted by Conrad (1976). Wecker (1978) shows that

using sales data instead of demand causes a negative forecasting bias that increases with stockout frequency. Bell (1978, 1981) presents a newsvendor-type analysis to optimize the purchasing and distribution decisions for a magazine and newspaper wholesaler or distributor. Hill (1992) assumes demand to depend on the number of customers as well as customer order sizes, and estimates demand by inflating sales using historical data to adjust for stockouts. Lau and Lau (1996) extend the work of Conrad (1976) to allow for general demand distributions and random censoring levels. The estimation methods in these papers assume that demand among different stores and over different periods is independent and identically distributed (i.i.d.). In our study, demand is autocorrelated and nonidentically distributed; thus, we cannot use either of the methods presented in the above papers. We use two methods to estimate the demand based on sales data. The first is a Bayesian analysis via Markov chain Monte Carlo simulation (see Best et al. 1996). We use the BUGS software discussed in detail by Lunn et al. (2000). The second method is an iterative maximum-likelihood estimation algorithm, similar in nature to the EM algorithm in Dempster et al. (1977).

In our approach to determining the appropriate quantity to purchase for each store, we first estimate the demand and subsequently optimize. There has been some recent related work on joint estimation and optimization of models. Examples include Liyanage and Shanthikumar (2005), Besbes et al. (2010), and Cooper et al. (2006). Broadly speaking, these papers emphasize using operational objectives when estimating or fitting a model as opposed to more traditional measures such as least squares or maximum likelihoods. These papers apply this concept in relatively simple cases, e.g., Liyanage and Shanthikumar (2005) apply their approach to a newsvendor with a single unknown demand parameter to estimate based on i.i.d. demand data. Besbes et al. (2010) considers a statistical test that incorporates decision performance into a measure of statistical validity in the context of fitting a demand curve. Even in these cases the machinery of deriving a best test or optimum decision is significant. Although there may be benefits from considering operational performance in our problem, the size of the estimation problem we investigate limits the applicability of these approaches at this time.

3. A Model for Purchase and Allocation Decisions

In this section we present the model for determining the purchase quantity for movies and their allocation to stores. We consider first a deterministic formulation that allows us to introduce the problem and its

solution algorithm. We then generalize the model to the stochastic case.

3.1. Deterministic Problem

We first present a mathematical programming formulation of the deterministic problem. Let \mathcal{S} be the set of stores and $T = 27$ be the number of days within the release month. Because approximately 90% of a movie's rentals occur in the first month after its release, we consider how many copies of a movie should be purchased for rent during this period. Let c_i be the number of copies purchased for store i , $i \in \mathcal{S}$. Let d_{ij} be the demand at store i on day j , $j = 1, \dots, T$ and let $s_i = \sum_j d_{ij}$ be the total demand at store i . For each store i , let r_{ij} be the number of rentals on day j and l_{ij} be the number of copies left on the shelf at the end of day j . Let $\mathbf{r}_i^j = \{r_{i1}, r_{i2}, \dots, r_{i,j-1}\}$ be the history of rentals through day $j - 1$. Observe $r_{ij} = d_{ij}$ if copies of the movie are left on the shelf at the end of the day, i.e., $l_{ij} > 0$. Otherwise, $r_{ij} \leq d_{ij}$, i.e., demand is censored, and the observed rentals is a lower bound on demand. If the purchase quantity is restricted, let c be the maximum number of copies of a movie that the rental firm can purchase.

Let $u_{ij}(\mathbf{r}_i^j)$ be the number of copies returned to store i on day j expressly written to depend on the rental history. We assume that these copies are returned at the beginning of day j and placed on the shelf immediately (alternate treatments can be easily accommodated). In the simplistic deterministic problem, d_{ij} is known and $u_{ij}(\mathbf{r}_i^j)$ is a deterministic function of \mathbf{r}_i^j . Let π be the number of rentals per copy of a movie required for the firm to break even. This is an exogenous factor determined by the rental firm. Note that π is typically larger for copies purchased under standard contracts compared to those

purchased under revenue-sharing ones. Table 1 provides a summary of our notation.

We use the following integer programming formulation to define the firm's problem of determining the allocation of copies of a movie to the stores (c_i, l_{ij}, r_{ij} are decision variables):

$$\max \sum_{i \in \mathcal{S}} \sum_{j=1}^T r_{ij} - \pi c_i \quad (1a)$$

$$\text{s.t.} \sum_{i \in \mathcal{S}} c_i \leq c \quad (1b)$$

$$r_{ij} = \min\{d_{ij}, l_{ij-1}\} \quad \text{for all } i \in \mathcal{S}, j = 1, \dots, T \quad (1c)$$

$$l_{ij} = l_{ij-1} - r_{ij} + u_{ij}(\mathbf{r}_i^j) \quad \text{for all } i \in \mathcal{S}, j = 1, \dots, T \quad (1d)$$

$$l_{i0} = c_i \in \text{integer} \quad \text{for all } i \in \mathcal{S} \quad (1e)$$

$$c_i, l_{ij}, r_{ij} \geq 0 \quad (1f)$$

$$\text{for all } i \in \mathcal{S}, j = 1, \dots, T. \quad (1g)$$

Assuming, without loss of generality, that the rental price is \$1 and cost per unit is π , the objective (1a) maximizes the profit within the release month. Constraint (1b) enforces the purchase quantity restriction. Without this restriction, e.g., for titles purchased under standard contracts, the problem is separable by store and is not difficult to solve. Constraint (1c) ensures that the rentals for each store-day are less than the demand. Constraint (1d) presents the inventory balance equations that define the interaction between the rental process and the return process. The initial allocation of copies to stores, c_i , is the only real decision we make, and all other variables are calculated based on estimated demand and return and the dynamics of the problem. Therefore, we only impose integrality on the initial allocations in (1e). Integrality itself is not important in our context, and so Problem 1 could be solved as an LP with rounding to achieve a near-optimal integer solution. However, the greedy approach we outline next solves the integral problem and will be applied to the stochastic problem as well.

We solve Problem (1a)–(1g) directly by making several observations. First, because the rental price is constant over the time period, there is no reason not to rent an available copy, given demand. Second, under reasonable assumptions, we can show that each copy allocated to a store will have a nonincreasing number of rentals compared to the previous copy allocated. Therefore, one can iteratively allocate copies to the stores based on which store will provide the greatest number of rentals until c copies are distributed or until the marginal cost of purchasing an additional copy at any store exceeds the marginal revenue.

Table 1 Notation

\mathcal{S} :	Set of stores
T :	Number of days within the release month
π :	Break-even rentals per copy
c :	Maximum number of copies that the rental firm can purchase
c_i :	Number of copies assigned to store i
s_i :	Store demand, total demand at store i within the release month
d_{ij} :	Demand at store i on day j
r_{ij} :	Number of rentals at store i on day j
l_{ij} :	Number of copies left on shelf at store i at the end of day j ($l_{i0} = c_i$)
h_{ij} :	Number of copies not at store i during day j
$p_i(c_i)$:	Total number of rentals at store i in the release month if c_i copies are allocated to it
\mathbf{r}_i^j :	The history of rentals up to day $j - 1$ at store i , i.e., $\{r_{i1}, r_{i2}, \dots, r_{i,j-1}\}$
$u_{ij}(\mathbf{r}_i^j)$:	Number of copies returned to store i on day j
ρ_j :	Daily multiplier, percentage of total demand that occurs in day j
α_{ijt} :	Fraction of rentals made at store i on day j returned in exactly t days
A_{ij} :	A unit-mean random variable distributed as demand normalized by its mean, $E[s_i]$

That is, a greedy approach can be used. We detail this approach below, using what we refer to as the rental frontier.

Let $\rho_i(c_i) = \sum_{j=1}^T r_{ij}$ be the total number of rentals from store i as a function of the number of copies allocated to it, i.e., the rental frontier. Given demand, d_{ij} , and return, $u_{ij}(\mathbf{r}_i^j)$, we can determine $\rho_i(c_i)$ as follows. Let h_{ij} be the number of copies not at store i during day j (i.e., rented before day j and not returned on or before it, so $h_{ij} = c_i - l_{ij} - r_{ij}$). The number of rentals on each day is the minimum of demand and availability, that is,

$$r_{ij} = \min\{d_{ij}, c_i - h_{ij}\}$$

for all $i \in \mathcal{S}$ and $j = 1, \dots, T$. (2)

Let $\rho_{ij}(c_i)$ be the total number of rentals at store i through day j given c_i copies. Then,

$$\rho_{ij}(c_i) = \rho_{i, j-1}(c_i) + \min\{d_{ij}, c_i - h_{ij}\}$$

for all $i \in \mathcal{S}$, $j = 1, \dots, T$, (3)

with $\rho_{i,0}(c_i) = 0$. Thus, $\rho_i(c_i) = \rho_{i,T}(c_i)$.

The rental frontier depends greatly on the return process. Let u_{ijk} be the number of copies returned to store i on day k from rentals made on day j , $k > j$. Then $u_{ik}(\mathbf{r}_i^k) = \sum_{j=1}^{k-1} u_{ijk}$. Let α_{ijt} be the fraction of rentals made on day j returned in exactly t days. Then $u_{ijk} = \alpha_{ij, k-j} r_{ij}$. We define the return process $u_{ik}(\mathbf{r}_i^k)$ to be *monotone* if the fraction of rentals made on day j returned by day k is at least as large as the fraction returned from any subsequent day $j+1, j+2, \dots$. Mathematically, the return process is monotone if

$$\sum_{t=1}^{k-j} \alpha_{ijt} \geq \sum_{t=1}^{k-j-1} \alpha_{i, j+1, t}$$

for all $j = 1, \dots, T$, $k = 2, \dots, T$, $i \in \mathcal{S}$. (4)

PROPOSITION 1. *If the return process in a store is monotone, the rental frontier of that store, $\rho_i(c_i)$, is a concave nondecreasing function of the number of copies allocated c_i .*

All proofs appear in Appendix A.

Using $\rho_i(c_i)$, the problem (1a)–(1g) is reformulated as

$$\max_{\sum_{i \in \mathcal{S}} c_i \leq c} \sum_{i \in \mathcal{S}} \rho_i(c_i) - \pi c_i \quad (5)$$

with c_i integer where $\rho_i(c_i)$ is defined by (3). The slope of the rental frontier, $d\rho_i(c_i)/dc_i$ defines the marginal number of rentals obtained for an additional copy allocated to store i . Because of the linking constraint, $\sum_{i \in \mathcal{S}} c_i \leq c$, we propose the following greedy algorithm to iteratively allocate copies to the stores, stopping when c copies are allocated or the slope of all stores is less than π .

Algorithm 1 (Greedy approach—Deterministic demand)

1. Initialization: Let $\rho_i(0) = 0$, $c_i = 0$, find $\rho_i(1)$ using (3), and let $Slope_i = \rho_i(1)$ for all $i \in \mathcal{S}$.
2. Find maximum slope: Find

$$i^* = \min \left[\arg \max_i \{Slope_i\} \right].$$

3. Check stopping rule: If $Slope_{i^*} < \pi$ or $\sum_{i \in \mathcal{S}} c_i = c$, STOP.

4. Allocate copy: Let $c_{i^*} = c_{i^*} + 1$.

5. Update slope: Find $\rho_{i^*}(c_{i^*} + 1)$ using (3). Let $Slope_{i^*} = \rho_{i^*}(c_{i^*} + 1) - \rho_{i^*}(c_{i^*})$. Go to step 2.

Algorithm 1 finds the store i^* with the maximum slope of rental frontier in step 2 and allocates a copy to that store in step 4 if this slope is higher than the break-even point and the purchase quantity restriction is not violated. The slope is then updated in step 5, and the procedure repeats. Finding $\rho_i(c_i)$ using (3) for each store $i \in \mathcal{S}$ has a complexity $O(T)$. Therefore, the initialization step of Algorithm 1 has a complexity $O(T|\mathcal{S}|)$. Each iteration of the algorithm involves finding $\rho_i(c_i)$ for one store and finding the store with maximum slope that has a complexity $O(T + |\mathcal{S}|)$. If D represents the total demand in all store-days, Algorithm 1 goes through $O(D)$ iterations, one for each copy allocated to a store. Because, typically, D is much larger than $|\mathcal{S}|$, the overall complexity of Algorithm 1 is $O((T + |\mathcal{S}|)D)$. We show,

PROPOSITION 2. *If the return process is monotone, the greedy allocation in Algorithm 1 obtains an optimal solution to problem (1a)–(1f).*

3.2. Stochastic Problem

We now present the more general problem where demand is viewed as uncertain. For store i and day j , let random variables D_{ij} be the demand and R_{ij} be the number of rentals. We can write the stochastic problem as

$$\max \sum_{i \in \mathcal{S}} \sum_{j=1}^T E[R_{ij}] - \pi c_i \quad (6a)$$

$$\text{s.t.} \sum_{i \in \mathcal{S}} c_i \leq c \quad (6b)$$

$$R_{i1} = \min\{D_{i1}, c_i\} \quad \text{for all } i \in \mathcal{S} \quad (6c)$$

$$R_{ij} = \min \left\{ D_{ij}, c_i - \sum_{t=1}^{j-1} \left(1 - \sum_{k=1}^{j-t} \alpha_{tk} \right) R_{it} \right\}$$

for all $i \in \mathcal{S}$, $j = 2, \dots, T$ (6d)

$$c_i \in \text{integer}, \quad c_i \geq 0 \quad \text{for all } i \in \mathcal{S}. \quad (6e)$$

In this formulation we use the α_{tk} notation to express the fraction of day t rentals returned after

k days as opposed to the equivalent $u_{ik}(\mathbf{r}_i^k)$ notation. Based on our estimation procedures discussed in §4, we find this to be a convenient notation. Further, we show there that we can best estimate the demand process, D_{ij} , as the product of two independent random variables, S_i , the total demand for store i , and P_j , a multiplier representing the fraction of demand realized on day j . Let S_i have mean s_i and variance $\sigma_{s_i}^2$; and P_j have mean p_j and variance $\sigma_{p_j}^2$. Then,

$$\Pr(D_{ij} \leq d_{ij}) = \Pr(S_i P_j \leq d_{ij}). \quad (7)$$

Observe that problem (6a)–(6e) is a generalization of the newsvendor problem. There is one purchase that is used to satisfy demand over a finite horizon. If the release month consisted of only one day, then it would be a newsvendor problem where for each store i the problem is reduced to maximizing $E[\min\{D_{i1}, c_i\}] - \pi c_i$. However, given the return process and multiple days, the problem is more involved.

We approach the problem by first letting A_{ij} be a unit-mean random variable distributed as D_{ij} normalized by its mean, i.e.,

$$\Pr(D_{ij} \leq d_{ij}) = \Pr\left(A_{ij} \leq \frac{d_{ij}}{s_i p_j}\right).$$

The rental frontier is now given by the random variable $R_i(c_i) = \sum_{j=1}^T R_{ij}$, where we define R_{ij} in (6c) and (6d). We can then substitute $D_{ij} = s_i p_j \times A_{ij}$. For any realization $\mathbf{a}_i = \{a_{ij}\}_{j \in [1, T]}$ of $\{A_{ij}\}_{j \in [1, T]}$, the rental frontier, $\rho_i(c_i, \mathbf{a}_i)$ can be calculated as

$$r_{i1} = \min\{s_i p_1 a_{i1}, c_i\} \quad (8a)$$

$$r_{ij} = \min\left\{s_i p_j a_{ij}, c_i - \sum_{t=1}^{j-1} \left(1 - \sum_{k=1}^{j-t} \alpha_{tk}\right) r_{it}\right\}$$

for all $j = 2, \dots, T$ (8b)

$$\rho_i(c_i, \mathbf{a}_i) = \sum_{j=1}^T r_{ij}. \quad (8c)$$

Given distributions for the A_{ij} s, the expected rental frontier for each store can be found. We have

PROPOSITION 3. *If the return process is monotone, the expected slope of the rental frontier for store i is non-increasing in the number of copies allocated to it, c_i .*

We estimate the slope of the rental frontier in Algorithm 2 using a sample path approach. To do so, we require distributions for the A_{ij} s and estimates of the mean and variance of S_i and P_j . In §4.4 we study how to estimate these means and variances. To simplify the optimization we make several assumptions regarding the A_{ij} 's. First, we propose to aggregate A_{ij} , $j = 1, \dots, T$ into A_i , to reduce the number of samples required. That is, we assume that the A_{ij} are

identically distributed across days, i.e., $A_{ij} \equiv A_i$ for all j for some unit mean A_i . The variance of A_{ij} is given by

$$\begin{aligned} \sigma_{A_{ij}}^2 &= \frac{\sigma_{D_{ij}}^2}{(s_i p_j)^2} = \frac{\sigma_{p_j}^2 s_i^2 + \sigma_{s_i}^2 (p_j^2 + \sigma_{p_j}^2)}{(s_i p_j)^2} \\ &= \frac{\sigma_{s_i}^2}{s_i^2} \left[1 + \frac{\sigma_{p_j}^2}{p_j^2} \left(1 + \frac{s_i^2}{\sigma_{s_i}^2}\right)\right]. \end{aligned}$$

We propose to let the variance of A_i be given by

$$\sigma_{A_i}^2 = \frac{\sigma_{s_i}^2}{s_i^2} \left[1 + v \left(1 + \frac{s_i^2}{\sigma_{s_i}^2}\right)\right],$$

where $v = 1/T \sum_{j=1}^T \sigma_{p_j}^2 / p_j^2$ is the mean squared coefficient of variation of the P_j s. We performed paired t -tests per store per movie to find whether the transformation preserves the average standard deviation of A_{ij} , i.e., $H_0: \sigma_{A_i} = 1/T \sum_{j=1}^T \sigma_{A_{ij}}$. The average p -value of all tests was 0.15. Therefore, with a significance level of 5%, we do not reject this null hypothesis. To avoid the inherent assumption of normality in the paired t -test, we also performed Wilcoxon signed-rank tests and obtained the same conclusion.

Second, we assume A_i has a Gamma distribution. Overall demand at each store may be viewed as an observation of a nonnegative random variable. The relatively high coefficient of variation observed in the data (the average C.V. = 0.58) and nonnegativity of demand, suggests many distributions including the Gamma distribution. We note that alternate demand distributions (truncated normal, log normal) were tested and gave similar results to those presented below.

Based on Proposition 3 and the above discussion, we propose the following algorithm to find the solution to problem (6a)–(6e). This provides the optimal number of movies for each store $i \in \mathcal{S}$.

Algorithm 2 (Greedy approach—stochastic demand)

1. Initialization: Let n be the number of intervals used to discretize A_i , $prob = 1/n$, $U = \{\mathcal{S}\}$, $c_i = 0$, $\rho_i(0, a_i) = 0$ for any a_i , and $E[slope_i] = 0$ for all $i \in \mathcal{S}$.
2. Discretize A_i : For each $i \in U$, let $a_i = F_{A_i}^{-1}(prob)$.
3. Find the slope: For each $i \in U$, find $\rho_i(c_i + 1, a_i)$ using recursion (8a)–(8c). Let $\rho'_i(c_i, a_i) = \rho_i(c_i + 1, a_i) - \rho_i(c_i, a_i)$.
4. Find expected slope: For each $i \in U$, let $E[slope_i] = E[slope_i] + \rho'_i(c_i, a_i)/(n - 1)$.
5. $prob = prob + 1/n$. If $prob < 1$, then go to step 2.
6. Find the maximum slope: Find

$$i^* = \min_i \left[\arg \max \{E[slope_i]\} \right].$$

Let $U = \{i^*\}$.

7. Check stopping rule: If $E[\text{slope}_{i^*}] < \pi$ or $\sum_{i \in \mathcal{S}} c_i = c$, then STOP; otherwise, $c_{i^*} = c_{i^*} + 1$, $E[\text{slope}_{i^*}] = 0$, $\text{prob} = 1/n$, and go to step 2.

Algorithm 2 discretizes A_i into n equiprobable intervals and includes two loops: an expectation loop, steps 2 to 5, that finds the expected slope of the rental frontier for a given store and allocation, and an allocation loop, steps 2 to 7, that, while the purchase quantity restriction is not violated, adds a copy to the store with maximum expected slope if the latter exceeds π and updates that store's expected slope. In the first iteration, Algorithm 2 finds the expected slope for all stores, whereas in later iterations the expected slope is updated only for the store with maximum expected slope. Therefore, the complexity of steps 2 and 3 are $O(T|\mathcal{S}|)$ in the first iteration and $O(T)$ in later iterations. Each iteration consists of n repetitions of steps 2 and 3. If D represents the total demand for all store-days, Algorithm 2 goes through $O(D)$ iterations, one for each copy allocated to a store. Thus, the overall complexity of Algorithm 2 is $O(nTD)$.

4. Estimating Demand and Return from Rental Data

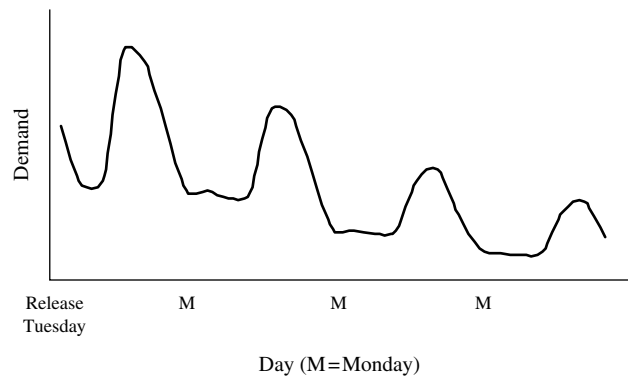
The main operations of a rental store consist of two processes: the rental (demand) process and the return process. We require models of these two processes to optimize the rental store's performance. In this section we propose four models for the returns process and four models for the demand process. Using the data set over all stores and days for each of the comparables, we estimate the parameters of these models using a Markov Chain Monte Carlo Simulation. For each new movie, we then average the estimates for its comparables to provide a method of forecasting its demand and returns. Then, using the data for each of the new movies, we forecast the number of returns for each of the four returns models. We compare the forecasted returns to the observed returns using the root mean squared error (RMSE) and the bias of the forecast. We proceed similarly for the demand models, forecasting the demand for each store and day for the four proposed models. We compare the forecasted demand to the observed rentals, both for the uncensored days (where actual demand is observed) and censored days (where demand is censored by lack of inventory). For the censored days, using simulation, we measure the error in the number of censored days, the likelihood of observing the number of censored days, and the likelihood of observing each day as being censored. In doing so we determine the best demand and return model of those we propose. Prior to doing so, we describe some common practices in demand forecasting in the movie rental industry and explain how we build on these practices in our models.

4.1. Forecasting in Practice

The problem faced by the rental firm is to forecast the demand at each store and day for a new release. The demand for a movie at each store-day depends on many factors, including demand in previous days, the day of the week, the store's location, and the number of active users. Thus, demands at different stores or on different days are nonidentically distributed. Operational data from rental stores demonstrate that observed demand, i.e., sales, follows a combined decay and cyclical pattern over time (see Figure 1). Demand is high when the movie is released and decreases over time until it reaches a low level after one month. In addition, rental stores observe a weekly pattern of higher demand during the weekends. In the example, the movie is released on a Tuesday (as is typical in the United States) and achieves its highest demand on the following Saturday.

It is common in practice to forecast demand and return of new releases based on the demand and return processes of comparable movies. This forecasting process includes three steps performed by an expert movie forecaster. However, in some cases, the output of this process is not an explicit demand and return forecast, but rather the advised allocation of movies to stores based on an implicit forecast. In the first step, comparable movies, i.e., ones with demand and return patterns that are believed to be similar to that of the new release, are chosen. Typically, one to three such comparable movies are selected for a new release. For example, when forecasting the demand and return for *Shrek 3*, *Shrek* and *Shrek 2* are reasonable comparables. Second, the expert inflates or deflates the store allocations for each comparable subjectively to account for how the title fared during its initial month in the stores. Third, the expert chooses weights to apply in averaging the comparables when forecasting demand for the new release. The weights are chosen subjectively to account for perceived differences in size between the new release and its comparables. For example, U.S. box office sales for *Shrek*,

Figure 1 Typical Observed Demand Pattern for a Movie



Shrek 2, and *Shrek 3* were \$268 million, \$436 million, and \$320 million, respectively. Therefore, a possible set of weights would be $320/268 = 1.09$ for *Shrek* and $320/436 = 0.73$ for *Shrek 2*. In practice, the weights may be based on any number of factors, such as box office sales or marketing initiatives undertaken by the firm or the studio.

In the next subsection, we describe several statistical methods that can be used to estimate the return processes for the comparable movies based on observed rents and returns. We explain how these can be combined to develop an estimate of the return process for the new release. We then compare the different return models and choose the best among them. We subsequently do the same for the demand process.

4.2. Estimating Return for a New Release

The return process is what distinguishes a rental store from a sales-oriented store. Although customers can return products in sales-oriented stores, returns are, normally, a small fraction of the sales and are neglected in most operational analyses. In contrast, returns account for approximately half of the daily activity at a rental store, and are thus a vital part of this business. For each comparable movie, we estimate the return process based on its r_{ij} , h_{ij} , and u_{ij} . For a given movie, let \hat{U}_{ij} be the forecasted returns for store i and day j . We then combine the estimates for the comparables to forecast the return process for the new release. We consider four possible models of the return process:

Return Model 1: $\hat{U}_{ij} = \alpha_{ij} \times h_{ij}$ (a multiplicative model with store-specific parameters). This model assumes that a fixed fraction α_{ij} of copies off the shelf are returned to store i on day j . In this model, the α_{ij} s are determined by the observations of u_{ij} and h_{ij} . Thus, although it requires the estimation of $|\mathcal{S}|(T-1)$ parameters, this estimation is straightforward.

Return Model 2: $\hat{U}_{ij} = \alpha_j \times h_{ij}$ (a multiplicative model with parameters common across stores). This model is similar to Model 1, but assumes that the return pattern is identical for all stores. Model 2 is a parsimonious variation of Model 1 that requires estimating $T-1$ parameters.

Return Model 3: $\hat{U}_{ij} = \sum_{l=1}^{j-1} \alpha_{j-l} \times r_{il}$ (a time-dependent return rate). Here α_k represents the fraction of rentals that are returned in exactly k days. This model assumes that all rentals follow the same time-dependent return process. Model 3 requires the estimation of $T-1$ parameters.

Return Model 4: $\hat{U}_{ij} = \sum_{k=\max\{j-14,1\}}^{j-1} \alpha_{k,j-k} \times r_{ik}$ (a time- and day-dependent return rate). Here α_{jk} represents the fraction of rentals on day j that are returned in exactly k days. This model is similar to Model 3 but allows the return pattern to change over time. Although the rental duration is flexible,

we observed that almost all rentals are returned within 14 days. Therefore, to reduce the number of parameters estimated, we limit the return pattern to 14 days. Model 4 requires the estimation of fewer than $14(T-1)$ parameters.

We use the observed values of u_{ij} , h_{ij} , and r_{ij} to estimate the parameters of the return models. The parameters for Model 1 are calculated directly. Parameter estimates for Models 2–4 are made through a Bayesian procedure using Markov Chain Monte Carlo simulation (MCMCS). The BUGS statistical software integrates Gibbs sampling of a MCMCS and Bayesian updating to form estimates of the full joint distribution of the parameters. The approach starts from the specified initial distribution for the parameters and successively samples from the conditional distribution of each parameter given all the others in the model. We use *Normal* initial distributions for unbounded parameters, *Gamma* for nonnegative parameters, and *Beta* for parameters defined between 0 and 1. We note that BUGS does not require explicit specification of error terms in its model (see Best et al. 1996, p. 329). We use the WinBUGS software implementation of the approach for Microsoft Windows (Lunn et al. 2000). The initial distribution for the α s in BUGS is *Beta*(2, 2). The results are not sensitive to the initial distributions, converge in fewer than 10,000 iterations, and take, on average, 0.5, 16, and 10 minutes for Models 2–4, respectively.

We forecast the return parameters for the new releases by averaging the estimated parameters for each comparable movie, m . That is,

$$\text{Model 1: } \alpha_{ij}^{\text{New}} = \frac{1}{|M|} \sum_{m \in M} \alpha_{ij}^m,$$

$$\text{Model 2: } \alpha_j^{\text{New}} = \frac{1}{|M|} \sum_{m \in M} \alpha_j^m,$$

$$\text{Model 3: } \alpha_k^{\text{New}} = \frac{1}{|M|} \sum_{m \in M} \alpha_k^m,$$

$$\text{Model 4: } \alpha_{jk}^{\text{New}} = \frac{1}{|M|} \sum_{m \in M} \alpha_{jk}^m.$$

4.3. Comparison of Return Models

To compare the performances of our return models, we estimate the return parameters for all comparables using each return model and combine the estimates as discussed in §4.2 to obtain a forecast of the return process. We use the observed rental data for the new release to calculate two measures of fit. The measures are:

1. The RMSE of the forecasted returns versus the observed returns.

$$\text{RMSE} = \sqrt{\frac{\sum_{i \in \mathcal{S}} \sum_{j \in \mathcal{D}} (\hat{u}_{ij} - u_{ij})^2}{N}}$$

Table 2 Comparison of the Return Models Using the RMSE and BIAS Measures

Measure	Model 1	Model 2	Model 3	Model 4
RMSE	7.8	7.4	7.2	6.2
BIAS	1.14	1.30	0.56	0.10

2. The mean bias of forecasted returns.

$$BIAS = \frac{\sum_{i \in \mathcal{S}} \sum_{j \in \mathcal{D}} \hat{u}_{ij} - u_{ij}}{N}$$

Table 2 presents the weighted average of the RMSE and BIAS of the proposed return models for the 20 movies weighted by their size. (We note that the ranking of these measures for the different models does not depend on the contract type used to purchase the movies.) From the table it seems that Model 4, which predicts the fraction returned as dependent on the day the movie is borrowed and the length of time it is held, performs the best. Indeed, a pairwise *t*-test across all 20 movies confirmed that for both measures Model 4 performs better than the other models at a 95% confidence level.

4.4. Estimating Demand for a New Release

In order to estimate the demand for a new release, we proceed in much the same way as the movie rental firm, albeit more objectively. For each comparable we estimate the true demand based on the censored demand provided by rents. Using weights based on the firm’s estimates of the overall demand for the new release (the expert opinion), we then forecast the demand for the new title on each day at each store. We note that any day with three or fewer copies left on the shelf is assumed to be a censored day (i.e., zero left on the shelf). This adjustment expresses the reality that some units returned to the store may not be available for rent on the same day. The validity of this assumption was affirmed by the rental firm.

We note that standard demand estimation models with censored data typically assume that the demand process is independent and identically distributed across different periods. However, because of weekly cycles in demand and differences in store size, demand is dependent and nonidentically distributed over time and stores.

We propose several estimators for the demand for the comparables. These include a multiplicative model, and three models that express the weekly and daily patterns observed in Figure 1.

Demand Model 1: $\hat{D}_{ij} = s_i \times p_j$ (a multiplicative model). As in §3.2, s_i s denote store sizes and p_j s are multipliers that represent the daily pattern of demand. This model assumes that demand at all stores follows the same daily pattern scaled up by the store size. To estimate the demand for all store-days,

we estimate s_i for all stores and p_j for all days within the release month for a total of $|S| + 26$ parameter estimates—we consider the first 27 days of demand and let $p_1 = 1$ to normalize the values. The initial distributions we used in BUGS are $s_i \sim \text{Gamma}(5, 10)$, $p_j \sim \text{Gamma}(1, 1)$.

Demand Model 2: $\hat{D}_{ij} = \sum_{k=1}^7 \beta_k y_{kj} + \sum_{k'=1}^4 \beta'_{k'} y'_{k'j} + \rho d_{ij-1}$ (a cyclic, autoregressive model). Let y_{kj} , $k = 1, \dots, 7$, be binary dummy variables representing days of the week, i.e., $y_{kj} = 1$ if day j is the k th day of the week, and $y'_{k'j}$, $k' = 1, \dots, 4$, be binary dummy variables representing weeks of the month, i.e., $y'_{k'j} = 1$ if day j is in week k' . This model assumes that demand follows a combined weekly cyclic and time-diminishing process. To estimate the demand for all store-days, we estimate six β s, three β' s, ρ , and d_{i0} for each store (we let $\beta_7 = \beta_4 = 0$ to normalize the values) for a total of $|S| + 10$ parameter estimates. The initial distributions we used in BUGS are $d_{i0} \sim \text{Gamma}(5, 10)$, β_k , and $\beta'_{k'} \sim \text{Normal}(0, 10,000)$, $\rho \sim \text{Beta}(9, 1)$. Tang and Deo (2008) used a similar autoregressive demand model but without the cyclic term.

Demand Model 3: $\hat{D}_{ij} = \sum_{k=1}^7 \beta_k y_{kj} + \sum_{k'=1}^4 \beta'_{k'} y'_{k'j} + \rho^{j-1} d_{i0}$. (a cyclic, decreasing model). Here, in comparison to Model 2, the two effects of a weekly cyclic pattern and a time-diminishing demand are separated. We estimate the parameters as in Model 2. Several studies use decreasing, although not necessarily exponentially, demand models without the cyclic term (see Pasternack and Drezner 1999, Lehmann and Weinberg 2000, Gerchak et al. 2006).

Demand Model 4: $\hat{D}_{ij} = s_i (\sum_{k=1}^7 \beta_k y_{kj} + \sum_{k'=1}^4 \beta'_{k'} y'_{k'j})$ (a multiplicative, cyclic model). This model is a combination of the previous models in that it captures the multiplicative nature of Model 1 but enforces the weekly cyclic pattern of Models 2 and 3. We estimate $|S| + 9$ parameters. The initial distributions used in BUGS are $s_i \sim \text{Gamma}(5, 10)$, β_k and $\beta'_{k'} \sim \text{Normal}(0, 1)$.

We use the observed values of r_{ij} and l_{ij} to estimate the parameters of the demand models. Mean values of the initial distributions are based on subjective knowledge of the data. We use initial distributions with large variances to allow proper Bayesian updating. Note that, after convergence, the variances of posterior distributions are small. We examined several alternative initial distributions (with large variances) and obtained the same results, but with more iterations required for convergence. In implementation, the Bayesian estimates generated by the MCMCS converge in less than 10,000 iterations and take, on average, 6 minutes, 16 hours, 22 hours, and 5 minutes for demand Models 1, 2, 3, and 4, respectively. In addition, we use a joint maximum-likelihood estimator (MLE)/method of moments procedure to estimate

the values for Model 1. In this procedure we iteratively find an MLE estimate for the p_j s and a method of moments estimate for the s_i s until convergence is achieved. See Appendix B for details. The MLE estimator for Model 1 requires, on average, three minutes.

As was done for the returns models, we combine the parameter estimates for the comparable titles, m , into a forecast for the new release. In this case, the process is more involved. In particular, to combine the estimates of s_i for Models 1 and 4, or alternatively d_{i0} for Models 2 and 3, we propose weights that are intended to inflate or deflate the demand for each comparable such that it becomes equal in size—i.e., demand—to the new release. Because the demand of the new release is not known, we rely on the best estimate we have, namely, the true purchase quantity of the new release made by the firm. Let c^{New} be this quantity.

For each comparable m , we use the deterministic demand model (1a)–(1f) given in §3.1 on its ex post demand to find the optimal purchase quantity, c^m . That is, for each comparable, having observed its rents, we estimate the demand process, d_{ij} , for the comparables to use in (1c). We then use return Model 4 to provide $u_{ij}(r_i^j)$ in (1d). We let c in (1b) be a decision variable. Finally, we let $c^m = c^*$, the optimal number of copies to purchase (n.b., in solving program (1a)–(1f), we use the value of π corresponding to the new release to ensure that the quantity determined is scaled appropriately). The weight for comparable m is then $w^m = c^{\text{New}}/c^m$. The weights for each new release express the relative demand for the new movie versus its comparables. They do not convey any information on the demand pattern (i.e., the cyclic or declining nature of the demand), nor the relative importance of one comparable versus another in determining the correct demand pattern. Therefore, these weights relate only to the s_i s and d_{i0} s. For the other parameters in the models, we take a simple average. That is, our demand process estimate for the new release is given by

$$\begin{aligned} s_i^{\text{New}} &= \frac{1}{|M|} \sum_{m \in M} w^m s_i^m, & d_{i0}^{\text{New}} &= \frac{1}{|M|} \sum_{m \in M} w^m d_{i0}^m, \\ p_j^{\text{New}} &= \frac{1}{|M|} \sum_{m \in M} p_j^m, & \beta_k^{\text{New}} &= \frac{1}{|M|} \sum_{m \in M} \beta_k^m, \\ \beta_{k'}^{\text{New}} &= \frac{1}{|M|} \sum_{m \in M} \beta_{k'}^m. \end{aligned} \tag{9}$$

4.5. Comparison of Demand Models

We consider several methods to measure the accuracy of different demand models for the new releases. Note that our demand forecasts for the new releases, d_{ij}^{New} , must be compared to observed rentals for the new release that may be censored because of stockouts. For days without a stockout, the observed rentals may be

directly compared. However, for days with a stockout, it is not valid to compare the rentals to the forecast. Rather, one might compare the predicted likelihood of a stockout for these days. Let $G(\cdot)$ and $g(\cdot)$ be the predicted cumulative distribution function and probability density function. Assuming N independent observations, y_1, \dots, y_N , such that the first $N_c \leq N$ of them are censored, i.e., $d_i \geq y_i$ for $i = 1, \dots, N_c$, the likelihood measure

$$p(y | Model) = \prod_{i=1}^{N_c} (1 - G(y_i)) \prod_{i=N_c+1}^N g(y_i), \tag{10}$$

multiplies the likelihood of observing a stockout on the days with stockouts by the likelihood of the observed demand on the days without stockouts. Lockwood and Schervish (2005) present a version of (10) that allows for dependency among observations. Recall that for each movie we have approximately 12,150 demand observations (450 stores times 27 days) that are mutually dependent. As Lockwood and Schervish (2005) discuss, evaluating the version of (10) that allows correlation between these observations is impractical.

For our purposes, let \mathcal{D} be the set of days in the release month, $\mathcal{D}_i^c \subset \mathcal{D}$ be the set of censored days for store i , $N = |\mathcal{S}| \times T$ be the total number of store-days, and $N_c = \sum_{i \in \mathcal{S}} |\mathcal{D}_i^c|$ be the total number of censored store-days. Dropping the superscript, “New,” for store i and day j , let \hat{D}_{ij} be the forecasted demand for the new release with PDF and CDF of f_{ij} and F_{ij} (assumed Gamma distributed), respectively, and r_{ij} be the observed rentals for the new release. Our Bayesian estimation of the demand parameters provides a mean and variance for each parameter. We follow (9) to calculate the mean and variance of the demand parameters for the new release given those of the comparables. We then calculate the mean and variance of \hat{D}_{ij} using the appropriate demand model. For example, Model 1 defines the demand as store size, S_i , multiplied by daily parameter, P_j . If S_i has a mean s_i and variance $\sigma_{s_i}^2$ and P_j has a mean p_j and variance $\sigma_{p_j}^2$, then assuming S_i and P_j are independent, \hat{D}_{ij} has a mean $\hat{d}_{ij} = s_i p_j$ and variance $\sigma_{p_j}^2 s_i^2 + \sigma_{s_i}^2 (p_j^2 + \sigma_{p_j}^2)$. Our MLE algorithm for Model 1 provides a mean and variance for P_j but only a mean for S_i . In this case, we find the average coefficient of variation for store i , cv_i , over all comparable movies and estimate $\sigma_{s_i}^2 = (cv_i \times s_i)^2$.

The first two measures we use for the accuracy of different demand models are as follows:

1. The RMSE of the forecasted demand versus observed rentals over the noncensored store-days.

$$RMSE\text{-UNCENSORED} = \sqrt{\frac{\sum_{i \in \mathcal{S}} \sum_{j \in \mathcal{D} - \mathcal{D}_i^c} (\hat{d}_{ij} - r_{ij})^2}{N - N_c}}$$

This is a standard measure in the absence of censoring. It captures fidelity to the uncensored results and should be argued as trying to fit known data well.

2. The average log-likelihood (LL) of the observed rentals using the forecasted demand, as in (10).

$$\text{LL-RENTALS} = - \frac{\sum_{i \in \mathcal{S}} (\sum_{j \in \mathcal{D} - \mathcal{D}_i^c} \ln(f_{ij}(r_{ij})) + \sum_{j \in \mathcal{D}_i^c} \ln(1 - F_{ij}(r_{ij})))}{N}$$

This measure captures the overall likelihood of the observed data assuming that demand among different store-days is independent. As discussed following (10), this measure is not perfect in our setting. Still, we believe it provides some indication for the goodness of fit of different models.

In addition to the theoretical measures above, we performed a simulation analysis of the demand models. Using a sample path approach, we generate a random realization of demand for all store-days using the joint parameter distributions. This considers the correlation among demands over different store-days. We then identify whether each store-day is censored by comparing the simulated demand to observed availability. That is, letting \hat{d}_{ij} be the simulated demand, a day is considered censored if $\hat{d}_{ij} \geq l_{ij} + r_{ij}$. Let \hat{e}_i be the forecasted number of censored days for store i for a sample path and let \bar{e}_i be its average over all iterations. Let $\text{freq}(x)$ be the relative frequency of observing event x , i.e., frequency of event x in the simulation divided by the total number of sample paths. We propose the following measures of accuracy:

3. RMSE of the forecasted versus observed number of censored days over all stores.

$$\text{RMSE-CENSORED} = \sqrt{\frac{\sum_{i \in \mathcal{S}} (\bar{e}_i - |\mathcal{D}_i^c|)^2}{|\mathcal{S}|}}$$

This measure captures the fidelity to the censored results, i.e., what is the difference between the number of observed censored days and the one forecasted by the model.

4. The average LL of the observed number of censored days using the forecasted demand.

$$\text{LL-CENSORED} = - \frac{\sum_{i \in \mathcal{S}} \ln(\text{freq}(\hat{e}_i = |\mathcal{D}_i^c|))}{N}$$

This measure also captures the fidelity to the censored results, i.e., how well the model forecasts the number of censored days.

5. The average LL of the observed censorship pattern using the forecasted demand.

$$\text{LL-PATTERN} = - \frac{\sum_{i \in \mathcal{S}} (\sum_{j \in \mathcal{D} - \mathcal{D}_i^c} \ln(\text{freq}(\hat{d}_{ij} < l_{ij} + r_{ij})) + \sum_{j \in \mathcal{D}_i^c} \ln(\text{freq}(\hat{d}_{ij} \geq l_{ij} + r_{ij})))}{N}$$

Table 3 Comparison of Demand Models Using the Five Proposed Measures

Measure	Model 1	Model 1	Model 2	Model 3	Model 4
	MLE	Bayesian	Bayesian	Bayesian	Bayesian
1. RMSE-UNCENSORED	11.5	16.3	24.1	21.5	16.8
2. LL-RENTALS	3.3	3.8	4.9	4.9	4.7
3. RMSE-CENSORED	6.8	7.0	7.1	6.8	8.5
4. LL-CENSORED	3.5	6.5	6.5	8.6	7.8
5. LL-PATTERN	0.5	1.0	1.3	2.0	1.9

This measures the likelihood that in the simulation a stockout occurs on the observed censored days plus the likelihood that no stockout occurs on the observed uncensored days. As such, this is an overall measure of the pattern of censoring in the data.

We present the results of measures 1 to 5 for the four proposed demand models in Table 3 (again, we note that the ranking of these measures for the different models does not depend on the contract type used to purchase the movies). We observe that Model 1-MLE dominates the others. Using a pairwise t -test across the 20 movies, we find that measures 1, 2, 4, and 5 are statistically significant at a 95% confidence level. Measure 3 is inconclusive.

An intuitive explanation for the better fit of the multiplicative demand model is that the daily variations in demand observed in the data seem to depend on the store size, i.e., larger stores have larger daily variations and vice versa. Model 1 captures this effect by multiplying the daily multipliers by the store size. However, Models 2 and 3 assume that daily variations are equal among all stores regardless of their sizes. Although Model 4 (the multiplicative, cyclical pattern model) is more parsimonious, in actuality we need to estimate 459 parameters for it in contrast to 477 parameters for Model 1 (for 450 stores). The additional parameters allow us to capture effects not available in the cyclic model. For example, demand in the first week may not follow the cyclicity of the following weeks (typically, there is high demand on the midweek release date) and this may be significant because, typically, a high percentage of demand occurs in this week.

5. Numerical Results

In this section we compare the results of our purchase and allocation decisions to those of the firm and to those given by several other heuristics. We proceed as follows. For each new release we forecast the demand and return processes using the data from its comparable movies, demand Model 1-MLE, and return Model 4 as chosen in §4. These models provide inputs to Algorithm 2, which is used to determine the number of copies to purchase for each store. In solving Algorithm 2, we consider two cases. In the fixed-copies case, the total quantity purchased is set equal

to the actual quantity purchased by the firm, c , i.e., we require $\sum_{i \in S} c_i = c$. Thus, the algorithm allocates the purchased copies to the stores. In the optimized-copies case, we include no constraints on purchase quantity. Then, as noted, the problem disaggregates by store. Also note that Proposition 1 indicates that a monotone return process, as defined by (4), is a sufficient condition for the concavity of the rental frontier. Based on our data set, 94% of all estimated return parameters for Model 4 satisfy monotonicity. We plotted the rental frontier for the remaining 6%, and they are concave as well.

After the observed rentals and returns become available for the new release, we estimate (again using demand Model 1-MLE and return Model 4) the demand and return processes for the new release for each store and day. This estimation is the best assessment of the demand and return processes for the new release. Using this assessment, we estimate the profit that would be generated from any purchase and allocation decision. That is, for each initial inventory, c_i , we estimate the number of rentals that would occur on each day, and using the returns model, we estimate the number of returns and available inventory for subsequent days. Thus, we are able to compute for any given inventory allocation the total number of rentals and, given a value of π , the associated profit. We do so using the solutions given by the fixed-copies and optimized-copies cases, as well as the actual purchase quantity of the firm.

Table 4 summarizes the results of our analysis for the two sets of movies, standard and revenue sharing, with 10 titles each. For each title we give the firm's actual purchase quantity, the number of rentals generated and the implied profit based on the observed rentals and demands. The fixed-copies column presents the percentage change in rentals and profit. The optimized-copies column presents the percentage change in the number of copies purchased, rentals, and profit. We also report the average changes for each contract type. Because we optimize in the fixed- and optimized-copies cases based on the forecasted process, but evaluate the decisions based on the observed demand and return processes, there is no guarantee that our results will outperform those of the firm. This is shown by the negative profit changes in Table 4. Similarly, the fixed-copies decision could outperform the optimized-copies decision based on the observed demand. The optimized solutions only ensure that we make the best decision ex ante demand (i.e., using the forecasted demand and return processes).

For the standard contract we observe that under the optimized-copies case the firm could achieve a 2.51% higher profit by reducing the number of copies by approximately 9.98% and the number of rentals by 5.76%. A profit gain of 1.14% is achieved by the fixed-copies case. That is, by optimizing the allocation of copies to stores, without changing their quantity, we are able to obtain 45% of the profit improvement.

Table 4 Comparison of the Firm's Actual Purchase Quantity and Resulting Rentals and Profits to Those Given by the Optimization Fixing the Total Number of Copies Purchased, and the Optimization Allowing the Number of Copies Purchased to be Optimized

Contract	Movie	Firm's decision			Fixed copies (%)		Optimized copies (%)		
		Copies	Rentals	Profit	Rentals	Profit	Copies	Rentals	Profit
Standard	S1	67,254	351,948	150,186	2.04	4.79	-9.80	-5.14	1.12
	S2	35,349	213,798	107,751	-0.96	-1.91	-2.11	-2.13	-2.14
	S3	49,999	204,081	54,084	-0.36	-1.35	-6.14	-3.62	3.34
	S4	24,730	107,677	33,487	1.86	5.97	-2.57	-0.29	4.75
	S5	28,491	122,458	36,985	-1.81	-6.00	-14.56	-12.44	-7.56
	S6	39,446	137,753	19,415	0.52	3.70	-22.17	-12.59	45.80
	S7	28,545	110,215	24,580	0.97	4.34	-13.74	-7.38	14.77
	S8	5,550	23,167	6,517	1.06	3.76	-16.29	-11.12	2.09
	S9	12,082	57,097	20,851	-2.63	-7.21	-5.64	-6.74	-8.67
	S10	6,495	22,339	2,854	2.14	16.78	-4.59	-1.28	21.32
Average		29,794	135,053	45,671	0.38	1.14	-9.98	-5.76	2.51
Revenue sharing	RS1	59,923	311,256	251,333	0.98	1.21	37.50	22.30	18.68
	RS2	58,949	262,718	203,769	2.94	3.79	23.59	19.35	18.12
	RS3	45,782	256,087	210,305	-0.38	-0.47	36.56	19.75	16.09
	RS4	68,863	248,552	179,689	2.71	3.75	33.08	16.15	9.67
	RS5	29,955	137,532	107,577	2.63	3.36	34.86	27.97	26.05
	RS6	43,664	202,629	158,965	1.72	2.19	27.66	15.18	11.75
	RS7	40,111	190,298	150,187	-0.01	-0.01	28.42	14.89	11.28
	RS8	41,922	151,590	109,668	2.68	3.71	44.59	22.11	13.51
	RS9	24,681	101,399	76,718	2.35	3.11	24.16	18.79	17.06
	RS10	22,905	94,261	71,356	2.05	2.71	30.17	17.14	12.96
Average		43,676	195,632	151,957	1.63	2.10	32.37	19.28	15.52

For the revenue-sharing contracts, we observe that the fixed-copies cases have a modest 2.10% improvement achieved through better allocation. However, without the constraint, a 32.37% increase in copies results in 19.28% more rentals and a 15.52% increase in profit. We observe that the firm overbuys standard titles and underbuys revenue-sharing titles compared with the optimized-copies decision. The difference between the 15.52% and 2.10% values provides a measure for the loss resulting from the purchase restrictions imposed by the studios. We discuss the impact of these restrictions on the benefits of revenue sharing in the next section. (The average profit increases are all significant at a 95% confidence level.)

Next, we compare our results to those given by simpler estimation approaches. In this regard we test several alternatives, comparing the average results of our estimation approach (from Table 4) to those for the alternative approaches in Table 5.

Test 1—Weighted Demand. A naive approach for estimating demand might be given by a simple weighted average of observed demand of each comparable movie for each store and day using the weights w^m as above. In Test 1, we use this demand model instead of Model 1-MLE to estimate the ex ante demand. We continue to use return Model 4.

Test 2—Common p_j 's. In Demand Model 1, we estimate the demand process for each movie separately. It may be argued that the daily demand pattern is common across movies. To this end, we test a model where a common set of p_j 's, given by their average, is used rather than the specific one for each movie. All else remains the same.

Test 3—Returns Model 1. We use the simpler returns Model 1 instead of Model 4 to forecast the returns process of the new release. We continue to use demand Model 1-MLE. This studies the value of capturing the more-detailed return process.

Test 4—Common Returns Model. We have estimated the returns process for each of the movies separately. It may be argued that the returns process is not movie specific and can be estimated using a single model across all of the comparable movies at the same time. We do so in this test.

We observe that the results for the fixed-copies case for both the standard and revenue-sharing contracts are similar to those of our policy; i.e., the differences are not statistically significant at a 95% confidence level. Because the fixed-copies case amounts to reallocating supply among the stores, it is not surprising that given a reasonable estimate of the demand and returns, there is only so much improvement that can be made. Thus, the relative improvement in profits of approximately 1.1% and 2.1% for the standard and revenue-sharing cases, respectively, result primarily from the optimization made through Algorithm 2.

In contrast, for the optimized-copies cases, we observe, using a 95% confidence level, that the naive demand and return models fail to achieve the same level of profit improvement as our policy. Test 2-Common p_j 's and Test 4-Common Returns models do achieve, on average, approximately two-thirds of the profit of our policy. This suggests that although there are similarities between the demand and return patterns, a 50% improvement in profit can be achieved by estimating these separately for each movie. In addition, in the event of sparse movie rental data, it is still possible to capture some of the realized gains using a pooled demand and return model. Further, the simple demand estimation procedure (averaging the observed demands for the comparables) and the simple returns Model 1 do not perform particularly well.

A heuristic approach to the fixed-copies case (with the quantity purchase restriction) is given by scaling the optimized-copies solution so that the constraint

Table 5 Comparison of Alternate Demand and Returns Process Estimation Procedures to Our Policy Given by Demand Model 1-MLE and Returns Model 4

Contract	Test	Fixed copies (%)		Optimized copies (%)		
		Rentals	Profit	Copies	Rentals	Profit
Standard	Our policy	0.38	1.14	−9.98	−5.76	2.51
	Test 1—Weighted demand	−0.02	−0.04	−21.80	−14.12	0.92
	Test 2—Common p_j 's	0.39	1.16	−12.18	−7.65	1.24
	Test 3—Returns model 1	0.34	1.01	−20.72	−13.46	0.76
	Test 4—Common returns	0.35	1.04	−4.28	−2.25	1.72
	Test 5—Apportionment	−0.07	−0.22			
Revenue sharing	Our policy	1.63	2.10	32.37	19.28	15.52
	Test 1—Weighted demand	1.55	2.00	11.73	8.30	7.32
	Test 2—Common p_j 's	1.66	2.13	23.81	15.29	12.84
	Test 3—Returns model 1	1.60	2.06	10.27	7.32	6.48
	Test 4—Common returns	1.64	2.12	18.68	12.12	10.24
	Test 5—Apportionment	1.30	1.67			

Note. For each we present the change in rentals and profit versus those of the firm for the fixed copies and optimized copies cases.

holds. That is, one can solve the optimized-copies problem independently for each store and then proportionately scale the number of copies apportioned to each store so that the c_i s sum to c . In Test 5-Apportionment in Table 5 we compare this apportionment heuristic to the optimal fixed-copies solution (both differences are statistically significant at a 95% confidence level). We observe for the standard contract case that doing so results in a loss compared with the firm's solution. For the revenue-sharing case, the heuristic results in a solution that may be improved upon by 25% using our optimization approach. This demonstrates the need to jointly optimize for all locations simultaneously.

We next test the robustness of the results of our policy with respect to π . The previous analysis uses the values of $\pi = 1$ and 3 as given by the firm's manager. Note that π depends on the salvage value obtained by the marginal unit, which is not known until after demand occurs; and on the purchase price, which may vary based on supplier and quantity purchased (see Cachon and Kok 2007 for a discussion on salvage value estimation). That is, $\pi = (P - S')/(\phi F)$, where P is the unit purchase price, S' is the firm's salvage value on the marginal unit purchased, and F is the rental fee. Let subscript "std." signify standard contracts; "r.s.," revenue sharing. The publicly available rental fee is \$5. For reasons of confidentiality the actual purchase price and marginal salvage value were not revealed to us. However, publicly available data indicate $P_{\text{std.}} = \$20$ and $P_{\text{r.s.}} = \$3$ may be valid numbers. Then $\phi = 0.6$, $S'_{\text{std.}} = \$5$, and $S'_{\text{r.s.}} = \$0$ would imply $\pi_{\text{std.}} = 3$ and $\pi_{\text{r.s.}} = 1$. Note that these values may be appropriate because the greater purchase quantity resulting from $\pi_{\text{r.s.}} = 1$ would lead to additional units to salvage, lowering their salvage value to the point where some units are returned to the studio (at 0 salvage value) rather than sold, generally after about six months. Because the purchase price and marginal salvage value may vary, we present in Table 6 representative summary results for the optimized-copies case in the neighborhood of the given values of π . These correspond to varying $S'_{\text{std.}}$ from \$0 to \$7.50 holding $P_{\text{std.}} = \$20$, and varying $P_{\text{r.s.}}$ from \$1.50 to \$4.50 holding $S'_{\text{r.s.}} = \$0$. These represent reasonable limits on these values given publicly available data and salvage prices observed at retail locations of the firm. As would be expected from a newsvendor analysis, the optimal purchase quantity is very sensitive to the value of π while the profit, as shown by its percentage change, is less so. Our conclusion that the firm overbuys standard titles is valid unless π lies below approximately 2.7. The firm underbuys revenue sharing titles at all reasonable values of π , greatly if the price per unit is low.

Table 6 Optimized Quantity Results While Varying π

Contract	π	Copies (%)	Rentals (%)	Profit (%)
Standard	2.5	12.02	7.87	2.76
	3	-9.98	-5.76	2.51
	3.5	-28.15	-19.52	9.71
	4	-43.32	-33.16	43.10
Revenue sharing	0.5	69.17	32.72	28.14
	1	32.37	19.28	15.52
	1.5	3.84	3.94	4.05

Note. The fixed-copies case is not sensitive to π by definition.

6. Discussion and Future Work

Recently, movie rental firms have been subject to intense competition from alternative business models, e.g., online rentals, movies-on-demand, and download-to-rent. For example, Blockbuster, a leading global video rental firm, reported negative net income in four out of five years (2004–2008), and its share price as of December 2009 has decreased by about 98% since its peak in 2002. Similarly, Movie Gallery, the second-largest video rental firm in the United States, filed for bankruptcy in October 2007. Our study considers how rental firms may use detailed transaction history of comparable movies to better predict demand and supply to improve their performance.

We model the demand and return processes for DVD rental to better forecast the supply requirements of the firm. We use the firm's entire data set—i.e., past demand and returns to all its stores—to forecast demand at the store and day level, rather than relying on the data for each store individually. Our procedure is limited in several ways. First, to prove concavity and optimality of our greedy algorithm, we require monotonicity in the return rate of movies. Although we find concavity holds throughout in our data set, we do not incorporate this sufficient constraint in our estimation procedure. Second, we develop the approach to the stochastic optimization problem using notation based on demand Model 1 and returns Model 4 (see (6d) and (7)). Although not a limitation, one would need to change the notation and Algorithm 2 to accommodate alternate demand or returns estimation models. Our approach, however, is limited in that we reduce the numerical computations by assuming demand is only dependent on store size (reducing A_{ij} to A_i). Further, we assume that demand is Gamma distributed.

We find that by purchasing the correct quantity we are able to modestly improve (2.5% on average) the profitability of a rental firm for movies purchased under a standard contract. In particular, this is achieved by reducing the number of copies purchased. In contrast, it seems appropriate to consider the fixed-copies case for the revenue-sharing titles.

There we observe that a 2.1% increase in the rental firm’s gross profit may be possible, in this case by better allocation to stores. We find that by increasing the purchase quantity, the firm could significantly increase its profits (15.5%). However, the firm cannot take advantage of this potential gain because of purchase quantity restrictions imposed by studios. We note that even if the differences are statistically significant, the forecasted gains might be nullified by estimation errors and in out-of-sample testing. This is especially important if the firm uses reallocation with fixed copies because the margin of gain there is relatively small.

We have several comments with respect to these results. First, to place this in perspective, a 2.1% increase in gross profits attributable to rentals (as opposed to DVD or other sales) at Blockbuster, Inc., based on their 2008 annual report, results in an average increase in net income by \$37.5 million over the last three years. Blockbuster reported an average annual operating loss of \$60.2 million during this period. So a 2.1% improvement would reduce the average operating loss by 62%. We emphasize that we present these numbers only to provide a context to understanding what 2.1% may mean in this industry. Second, our results represent only a small sample of the movies purchased and rented by the firm, and generalization to firmwide measures ignores various aspects of actual operations. In particular, our analysis is based on treating each movie individually and ignores firm policies that might link them, such as budgetary constraints, requirements on assortments and availability of substitutes, and requirements on merchandising, such as shelf-space presentation. With respect to the standard titles, it is possible that because of marketing or budgetary reasons, the overpurchase we observe is a response to the restrictions placed on the revenue-sharing titles.

Finally, we have observed the deleterious effect on profits of the purchase restrictions for the revenue-sharing contracts. That such restrictions could have such an effect is not surprising because the theory

behind revenue-sharing contracts is that without purchase restrictions they can align supply chains to the benefit of all parties. Thus, the question is raised: why do such restrictions exist? Addressing this questions calls for future work. We hypothesize that the quantity restrictions imposed may be related to studio concerns regarding the sale of copies of DVDs by the rental firms after the first month. That is, the studio may be concerned with cannibalization of sales and thus acts to restrict the number of copies.

As a demonstration of this potential reasoning, we present in Table 7 estimates of the average change in the mean profit of the firm and mean revenue of the studio, for the revenue-sharing contract and the quantity-restricted revenue-sharing contract versus the standard contract (this is the average for the 10 revenue-sharing titles in our data set). For example, for a standard price $P_{std.} = 20$, and salvage values $\bar{S}_{std.} = 10$ and $\bar{S}_{r.s.} = 2.5$, we observe that the rental firm’s profits are 55% higher under revenue sharing than they would have been under a standard contract (using $\phi = 0.6$ and $P_{r.s.} = 3$). The results are sensitive to a number of parameters, particularly the standard contract price and average salvage value (details of the estimation procedure are in Appendix C.) The implication of the table is that both the studio and the rental firm may benefit from a revenue-sharing contract or a quantity-restricted revenue sharing when there is no cannibalization. Further, in concert with supply chain coordination theory, the quantity restriction reduces the benefit for both parties. However, we observe that if units purchased by the firm are sold and fully cannibalize sales (on a one-to-one basis) of the studio, revenue sharing can lead to worse performance for the studio than a standard contract, although quantity restrictions can mitigate these losses.

These observations lead to several questions: What is the effect of cannibalization on contract choice? If studios actually do lose money under revenue-sharing contracts compared with standard contracts, why are they used? Why were quantity restrictions

Table 7 The Mean Change in Supply Chain Benefits over the Standard Contract

Contract	Supplier/Buyer	$P_{std.} = 20, \bar{S}_{std.} = 10$				$P_{std.} = 15, \bar{S}_{std.} = 10$			
		$\bar{S}_{r.s.} (%)$				$\bar{S}_{r.s.} (%)$			
		2.5	5	7.5	10	2.5	5	7.5	10
Revenue sharing	Rental firm	55	77	99	122	18	35	52	69
	Studio (no cannibalization)	44	56	68	80	91	107	124	140
	Studio (full cannibalization)	-61	-49	-37	-25	-49	-32	-17	-1
Quantity-restricted revenue sharing	Rental firm	31	47	63	79	0	12	25	37
	Studio (no cannibalization)	14	23	32	40	52	64	75	87
	Studio (full cannibalization)	-42	-33	-24	-15	-22	-10	1	13

put in place as opposed to other contract terms such as a buy-back agreement? We emphasize that it is only a hypothesis that concerns on cannibalization have lead to the quantity restrictions. Future work, which is beyond the scope of the current paper, should address these questions.

Acknowledgments

The authors thank the editor, associate editor, and referees for their many helpful suggestions that greatly improved the manuscript. This paper is based on a Ph.D. thesis by I. Hajizadeh. This work was partially supported by grants from the Natural Sciences and Engineering Research Council of Canada.

Appendix A. Proofs

PROOF OF PROPOSITION 1. $\rho_{i1} = \min\{d_{i1}, c_i\}$ is concave and nondecreasing in c_i . Assume $\rho_{il}(c_i)$ is concave and nondecreasing in c_i for $l = 1, \dots, j - 1$. Observe,

$$\begin{aligned} \rho_{ij}(c_i) &= \rho_{i,j-1}(c_i) + \min\{d_{ij}, c_i - h_{ij}\} \\ &= \rho_{i,j-1}(c_i) + \min\left[d_{ij}, c_i - \sum_{t=1}^{j-1} r_{it} + \sum_{t=2}^j u_{it}(r_{it}^t)\right] \\ &= \min\left[d_{ij} + \rho_{i,j-1}(c_i), c_i + \sum_{t=2}^j u_{it}(r_{it}^t)\right]. \end{aligned}$$

Therefore, by induction, if $\sum_{t=2}^j u_{it}(r_{it}^t)$ is concave and nondecreasing in c_i , we are done. Suppressing the first subscript $i \in S$, we have

$$\begin{aligned} \sum_{t=2}^j u_t(r^t) &= \sum_{t=2}^j \sum_{k=1}^{t-1} u_{kt} = \sum_{t=2}^j \sum_{k=1}^{t-1} \alpha_{k,t-k} r_k \\ &= \alpha_{11} r_1 + (\alpha_{12} r_1 + \alpha_{21} r_2) \\ &\quad + \dots + (\alpha_{1,j-1} r_1 + \alpha_{2,j-2} r_2 + \dots + \alpha_{j-1,1} r_{j-1}) \\ &= \alpha_{j-1,1} r_{j-1} + (\alpha_{j-2,1} + \alpha_{j-2,2}) r_{j-2} + \dots + \left(\sum_{t=1}^{j-1} \alpha_{1t}\right) r_1 \\ &= \alpha_{j-1,1} \sum_{k=1}^{j-1} r_k + (\alpha_{j-2,1} + \alpha_{j-2,2} - \alpha_{j-1,1}) \\ &\quad \cdot \sum_{k=1}^{j-2} r_k + \dots + \left(\sum_{t=1}^{j-1} \alpha_{1t} - \sum_{t=1}^{j-2} \alpha_{2t}\right) r_1 \\ &= \sum_{l=1}^{j-1} \left(\sum_{t=1}^{j-l} \alpha_{1t} - \sum_{t=1}^{j-l-1} \alpha_{l+1,t}\right) \rho_l. \end{aligned}$$

The coefficients of the ρ_l terms are nonnegative by (4). From the induction assumption, $\sum_{t=2}^j u_{it}(r_{it}^t)$ is the sum of concave nondecreasing functions and is, therefore, concave and nondecreasing. \square

PROOF OF PROPOSITION 2. Let $A^* = \{c_i^*\}_{i \in \mathcal{S}}$ denote the vector of optimal allocations provided by the optimization problem given in (1a)–(1g) and $A^G = \{c_i^G\}_{i \in \mathcal{S}}$ denote the vector of greedy allocations obtained from Algorithm 1. Under no purchase quantity restrictions, we have for each $i \in \mathcal{S}$, by Proposition 1.

a. If $c_i^* > c_i^G$, the slope of the rental frontier at $c_i^* - 1$ is no more than π . So, decreasing c_i^* by 1 will not decrease the objective function (1a).

b. If $c_i^* < c_i^G$, the slope of the rental frontier at c_i^* is no less than π . So, increasing c_i^* by 1 will not decrease the objective function (1a).

By repeating this analysis we can transform A^* into A^G without worsening the optimal solution. Therefore, A^G is an optimal solution of problem (1a)–(1g). The proof for the case with purchase quantity restrictions is similar and is omitted. \square

PROOF OF PROPOSITION 3. For a given vector of demand realizations, the rental frontier for store i is concave by Proposition 1. Therefore, the slope of the rental frontier is nonincreasing. The expected slope of the rental frontier is a convex combination of the slope per each demand realization. So, it is nonincreasing in c_i . \square

Appendix B. Maximum-Likelihood Demand Estimation

Let $f_j(\cdot; \bar{x})$ and $F_j(\cdot; \bar{x})$ be the parametric probability density and cumulative distribution functions, respectively, for p_j , the daily multiplier of day j , where \bar{x} is the parameter vector for the distribution, e.g., \bar{x} includes the mean and standard deviation if the distribution is normal. Let $\mu_j(\bar{x})$ be the mean of $f_j(\cdot; \bar{x})$. Let \mathcal{C} be the subset of stores with at least one day of censored demand. We estimate d_{ij} for all $i \in \mathcal{S}$ and $j = 1, \dots, T$ using the following algorithm. The algorithm terminates when demand estimates converge as measured by a two-criteria test: the relative change in total demand at each store must be less than $\epsilon_1 > 0$, and the average relative change in total demand among all stores must be less than $\epsilon_2 > 0$. Upon completion (in step 4), Algorithm 3 provides estimates for store sizes and daily multipliers.

Algorithm 3 (Maximum-likelihood demand estimation)

1. Initialization: Choose ϵ_1 and $\epsilon_2 > 0$. Let $n = 0$ and

$$\hat{p}_j^{(0)} = \frac{1}{|\mathcal{S} - \mathcal{C}|} \sum_{i \in \mathcal{S} - \mathcal{C}} \frac{r_{ij}}{\sum_{j=1}^T r_{ij}} \quad \text{for all } j = 1, \dots, T.$$

2. Estimate store demand: For all $i \in \mathcal{S}$, let

$$\hat{s}_i^{(n)} = \frac{\sum_{j \in \mathcal{D}_i} r_{ij}}{\sum_{j \in \mathcal{D}_i} \hat{p}_j^{(n)}} \quad \text{where } \mathcal{D}_i = \{j \mid l_{ij} > 0, 1 \leq j \leq T\}.$$

3. Estimate daily demand: For all $i \in \mathcal{S}$ and $j = 1, \dots, T$, let

$$\hat{d}_{ij}^{(n)} = \begin{cases} r_{ij} & \text{if } l_{ij} > 0, \\ \max\{r_{ij}, \hat{s}_i^{(n)} \times \hat{p}_j^{(n)}\} & \text{if } l_{ij} = 0. \end{cases}$$

4. Check convergence: For $n \geq 1$, if

$$\left| \frac{\sum_{j=1}^T (\hat{d}_{ij}^{(n)} - \hat{d}_{ij}^{(n-1)})}{\sum_{j=1}^T \hat{d}_{ij}^{(n-1)}} \right| \leq \epsilon_1 \quad \text{for all } i \in \mathcal{S}$$

and

$$\frac{1}{|\mathcal{S}|} \left| \sum_{i \in \mathcal{S}} \frac{\sum_{j=1}^T (\hat{d}_{ij}^{(n)} - \hat{d}_{ij}^{(n-1)})}{\sum_{j=1}^T \hat{d}_{ij}^{(n-1)}} \right| \leq \epsilon_2,$$

set $\hat{d}_{ij} = \hat{d}_{ij}^{(n)}$, $\hat{s}_i = \hat{s}_i^{(n)}$, $\hat{p}_j = \hat{p}_j^{(n)}$, and STOP.

5. Maximum-likelihood estimation of daily multipliers:
 For all $j = 1, \dots, T$, let

$$L_j(\bar{x}) = \prod_{i \in \mathcal{S}} g_j \left(\frac{r_{ij}}{\hat{s}_i^{(n)}}; \bar{x} \right)$$

where

$$g_j(y; \bar{x}) = \begin{cases} f_j(y; \bar{x}) & \text{if } l_{ij} > 0, \\ 1 - F_j(y; \bar{x}) & \text{if } l_{ij} = 0 \end{cases}$$

and $\bar{x}^* = \arg \max_{\bar{x}} \{L_j(\bar{x})\}$. Then, $\hat{p}_j^{(n+1)} = \mu_j(\bar{x}^*)$.

6. Let $n = n + 1$. Go to step 2.

In step 1, we estimate the daily multiplier for each day as the average daily multiplier among stores with no censored days. In step 2 we use the latest daily multipliers to estimate the store demand for each store based only on days in which actual demand was observed. In step 3, we estimate the daily demand for censored days as the maximum of the observed rentals and $\hat{s}_i^{(n)} \times \hat{p}_j^{(n)}$. The estimated demand for uncensored days is equal to the observed rentals. Step 4 tests for convergence of the algorithm. In step 5, we use maximum-likelihood estimation to reestimate \hat{p}_j (see Dempster et al. 1977 for MLE). Because the rentals among different stores are not identically distributed, we perform the maximum-likelihood estimation on the observed rentals normalized by the store demand. Finally, we return to step 2 to estimate new store demands.

In implementation of Algorithm 3, we assume that the daily multipliers are normally distributed, let $\epsilon_1 = 5\%$ and $\epsilon_2 = 1\%$, and use Excel Premium Solver to find the maximum of the likelihood function. In our computational experiments, demand converges in four iterations on average. Our initial estimates of daily multipliers in Algorithm 3 are based on only the stores with no censored days. To assess the robustness of the algorithm, we test different initial estimates of daily multipliers: (i) equating them for all days; (ii) choosing them randomly from a uniform distribution; and (iii) by inverting the original multipliers (normalized to sum to one). The algorithm converges in at most one additional iteration to within 0.3% of the demand estimated using the original multipliers. Therefore, we conclude that Algorithm 3 is robust to initial estimates of daily multipliers.

Appendix C. Estimation Procedure for Table 7

Using the ex post (not forecasted) rents and returns data from the revenue-sharing titles, we estimate the actual demand and return pattern for each revenue-sharing title. Then, using Algorithm 2 we determine the optimal purchase quantity, $c_{\text{std},i}$ assuming a standard contract ($\pi = 3$). We determine the expected resulting number of rentals, $r_{\text{std},i}$, and the rental firm's expected profit given by $Fr_{\text{std},i} - P_{\text{std},i}c_{\text{std},i} + \bar{S}_{\text{std},i}c_{\text{std},i}$, where \bar{S} is the average salvage value per unit; this is generally higher than the marginal value $S'_{\text{std},i}$. We compare these metrics (copies purchased, rentals, and profit) assuming a standard contract, i.e., $\pi = 3$, to

the cases where the optimal purchase quantity assumes a revenue-sharing contract ($\pi = 1$), and to a quantity-restricted revenue-sharing contract ($\pi = 1$), where the total number of copies purchased is the firm's actual purchase quantity. For these cases, the firm's profit is given by $\phi Fr - P_{\text{r.s.}}c + \phi \bar{S}_{\text{r.s.}}c$. In accordance with our previous analysis, we let $P_{\text{std.}} = \$20$, $P_{\text{r.s.}} = \$3$, $F = \$5$, $\phi = 0.6$, $\bar{S}_{\text{std.}} = \10 (here $\bar{S}_{\text{std.}}$ expresses a reasonable average salvage value for a previously viewed copy), and obtain the results for values of $\bar{S}_{\text{r.s.}}$ between \$2.5 and \$10. Revenue for the studio is given by $P_{\text{std.}}c$ for standard contracts and $(1 - \phi)Fr + P_{\text{r.s.}}c + (1 - \phi)\bar{S}_{\text{r.s.}}c$ for revenue-sharing and quantity-restricted contracts. For the studio we present the cases where copies sold to the rental firm either have no effect on or fully cannibalize studio sales (on a one-to-one basis). We assume that the studio loses \$15 per cannibalized sale.

References

- Bassamboo, A., R. S. Randhawa. 2007. Optimal control in a Netflix-like closed rental system. Working paper, McCombs School of Business, University of Texas at Austin, Austin.
- Bassamboo, A., S. Kumar, R. S. Randhawa. 2009. Dynamics of new product introduction in closed rental systems. *Oper. Res.* 57(6) 1347–1359.
- Bell, P. C. 1978. A new procedure for the distribution of periodicals. *J. Oper. Res. Soc.* 29(5) 427–434.
- Bell, P. C. 1981. Adaptive sales forecasting with many stockouts. *J. Oper. Res. Soc.* 32(10) 865–873.
- Besbes, O., R. Phillips, A. Zeevi. 2010. Testing the validity of a demand model: An operations perspective. *Manufacturing Service Oper. Management* 12(1) 162–183.
- Best, N. G., D. J. Spiegelhalter, A. Thomas, C. E. G. Brayne. 1996. Bayesian analysis of realistically complex models. *J. Roy. Statist. Soc. Ser. A (Statist. in Soc.)* 159(2) 323–342.
- Cachon, G. P., A. G. Kok. 2007. Implementation of the newsvendor model with clearance pricing: How to (and how not to) estimate a salvage value. *Manufacturing Service Oper. Management* 9(3) 276–290.
- Cachon, G. P., M. A. Lariviere. 2005. Supply chain coordination with revenue-sharing contracts: Strengths and limitations. *Management Sci.* 51(1) 30–44.
- Conrad, S. A. 1976. Sales data and the estimation of demand. *Oper. Res. Quart.* 27(1) 123–127.
- Cooper, W. L., T. Homem-de-Mello, A. J. Kleywegt. 2006. Models of the spiral-down effect in revenue management. *Oper. Res.* 54(5) 968–987.
- Dana, J. D. Jr., K. E. Spier. 2001. Revenue sharing and vertical control in the video rental industry. *J. Indust. Econom.* 49(3) 223–245.
- Dempster, A. P., N. M. Laird, D. B. Rubin. 1977. Maximum likelihood from incomplete data via the EM algorithm. *J. Roy. Statist. Soc. Ser. B (Methodological)* 39(1) 1–38.
- Entertainment Merchants Association, The (EMA). 2008. 2007 annual report on the home entertainment industry. Accessed October 16, 2008, <http://www.entmerch.org/>.
- Gerchak, Y., R. K. Cho, S. Ray. 2006. Coordination of quantity and shelf-retention timing in the video movie rental industry. *IIE Trans.* 38(7) 525–536.
- Hajizadeh, I. 2010. Problems in supply chain location and inventory under uncertainty. Ph.D. thesis, University of Toronto, Toronto, ON.
- Hill, R. M. 1992. Parameter estimation and performance measurement in lost sales inventory systems. *Internat. J. Production Econom.* 28(2) 211–215.

- Lau, H. S., A. H. Lau. 1996. Estimating the demand distributions of single-period items having frequent stockouts. *Eur. J. Oper. Res.* **92**(2) 254–265.
- Lehmann, D. R., C. B. Weinberg. 2000. Sales through sequential distribution channels: An application to movies and videos. *J. Marketing* **64**(3) 18–33.
- Liyanage, L. H., J. G. Shanthikumar. 2005. A practical inventory control policy using operational statistics. *Oper. Res. Lett.* **33**(4) 341–348.
- Lockwood, J. R., M. J. Schervish. 2005. MCMC strategies for computing Bayesian predictive densities for censored multivariate data. *J. Comput. Graphical Statist.* **14**(2) 395–414.
- Lunn, D. J., A. Thomas, N. Best, D. Spiegelhalter. 2000. WinBUGS—A Bayesian modelling framework: Concepts, structure, and extensibility. *Statist. Comput.* **10**(4) 325–337.
- Mortimer, J. H. 2008. Vertical contracts in the video rental industry. *Rev. Econom. Stud.* **75**(1) 165–199.
- Pasternack, B. A., Z. Drezner. 1999. The videotape rental model. *J. Appl. Math. Decision Sci.* **3**(2) 163–170.
- Randhawa, R. S., S. Kumar. 2008. Usage restriction and subscription services: Operational benefits with rational users. *Manufacturing Service Oper. Management.* **10**(3) 429–447.
- Rentrak. 2008. Pricing and terms. Accessed March 1, 2008, <http://www.rentrakppt.com/>.
- Tang, C. S., S. Deo. 2008. Rental price and rental duration under retail competition. *Euro. J. Oper. Res.* **187**(3) 806–828.
- Wecker, W. E. 1978. Predicting demand from sales data in the presence of stockouts. *Management Sci.* **24**(10) 1043–1054.