

Capacity Management in Rental Businesses with Two Customer Bases

Sergei V. Savin

Graduate School of Business, Columbia University, 404 Uris Hall, New York, New York 10027, sv30@columbia.edu

Morris A. Cohen, Noah Gans, Ziv Katalan

Operations and Information Management Department, The Wharton School, University of Pennsylvania,
3730 Walnut Street, Suite 500, Philadelphia, Pennsylvania 19104

{cohen@wharton.upenn.edu, gans@wharton.upenn.edu, katalan@wharton.upenn.edu}

We consider the allocation of capacity in a system in which rental equipment is accessed by two classes of customers. We formulate the problem as a continuous-time analogue of the one-shot allocation problems found in the more traditional literature on revenue management, and we analyze a queueing control model that approximates its dynamics. Our investigation yields three sets of results.

First, we use dynamic programming to characterize properties of optimal capacity allocation policies. We identify conditions under which “complete sharing”—in which both classes of customers have unlimited access to the rental fleet—is optimal.

Next, we develop a computationally efficient “aggregate threshold” heuristic that is based on a fluid approximation of the original stochastic model. We obtain closed-form expressions for the heuristic’s control parameters and show that the heuristic performs well in numerical experiments. The closed-form expressions also show that, in the context of the fluid approximation, revenues are concave and increasing in the fleet size.

Finally, we consider the effect of the ability to allocate capacity on optimal fleet size. We show that the optimal fleet size under allocation policies may be lower, the same as, or higher than that under complete sharing. As capacity costs increase, allocation policies allow for larger relative fleet sizes. Numerical results show that, even in cases in which dollar profits under complete sharing may be close to those under allocation policies, the capacity reductions enabled by allocation schemes can help to lift profit margins significantly.

Subject classifications: service systems; queueing control; stochastic knapsack; fluid models.

Area of review: Manufacturing, Service, and Supply Chain Operations.

History: Received September 2000; revisions received July 2001, May 2003; accepted March 2004.

1. Introduction

Rental businesses found in many sectors of the economy share some fundamental attributes. The rental company invests in equipment for which there is a potential demand, and a stream of customers patronizes the company, renting its equipment. After each rental, the equipment is returned to the company, and rental durations are typically significantly shorter than the life of the equipment, so that each unit may be used repeatedly.

For those who manage rental businesses, important managerial decisions focus on matching rental demand with the equipment supply. These decisions create a hierarchy of managerial controls at the company’s disposal. Longer-term capital-investment decisions set the company’s overall level of rental capacity and attempt to capture as much demand for rental services as is (marginally) profitable. While they provide for long-term matching between supply and demand, fleet-sizing decisions may not be used to counterbalance short-term supply and demand mismatches. On a tactical time scale, capacity allocation decisions may be needed to determine which customers are served when rental capacity becomes scarce.

In this paper, we consider a simple stationary model of a rental problem in which capacity must be rationed among two classes of arriving customers. We address both the lower-level allocation problem and the higher-level capacity-sizing problems, with an emphasis on the former.

Our approach to the tactical allocation problem follows in the spirit of early formulations of seat-allocation problems in the airline yield-management literature. (For example, see Littlewood 1972, Alstrup et al. 1986, and Belobaba 1989.) When should arriving customers of each of the classes be allowed to rent equipment, and when would they be “closed out?”

Two common assumptions made in traditional revenue management models make them inadequate for our purposes, however: They assume that there exists a finite horizon over which units of capacity can be sold and that each unit of capacity can be used only once. For example, in aviation there are c seats on a flight, and once they are sold or the plane takes off they are not available for sale.

While hotel problems could (and perhaps should) in principle be formulated as rental problems, most academic literature only addresses the problem of allocating the rooms

available on a single night. (For example, see Rothstein 1974, Ladany 1977, Williams 1977, Liberman and Yechiali 1978, Bitran and Gilbert 1996.) An exception is the application of linear-programming (LP)-based “bid price” controls to hotel stays. (See Williamson 1992, Weatherford 1995.) In this case, multiple nights are considered, but the problem is modeled as deterministic.

In rental businesses, however, the problem is most naturally treated as a problem in dynamic and stochastic control. An arriving customer rents a unit, which becomes unavailable for the length of the person’s rental. When the rental period ends, the unit becomes available again. Over any short period of time, the numbers of arriving and departing customers may be uncertain, and managers must develop effective policies for controlling the rental of system capacity.

We view the allocation of rental capacity as a continuous-time infinite-horizon problem in which arrivals of customers and durations of rentals are both uncertain. We formulate this problem as one of admission control to a multiple-server loss system. We assume that, if admitted into service, a customer pays a daily rental fee that depends on the class to which she belongs. If the rental request is rejected, then a class-dependent, lump-sum penalty is incurred. We show that this capacity allocation problem can be reduced to a special case of the *stochastic knapsack problem* introduced in the telecommunications literature (Ross and Tsang 1989), one in which arriving “objects” (demands) are all of size one.

We note that this formulation does not capture the use of prior information on rental duration. In some contexts, such as truck-trailer leasing (the application that originally motivated this paper) and storage-locker rentals, this information may not be available. In others, such as hotel systems, customer-stated projections of expected duration are readily available and can be of great value in improving the effectiveness of capacity allocation decisions. Thus, our approach has important limits.

Nevertheless, the simplicity of our approach allows us to make a number of contributions:

1. We demonstrate that the allocation problem with lump-sum penalties can be reduced to one with no penalties by appropriately adjusting the values of the rental fees. The adjustment factors are proportional to the penalty values and the service rates.

2. We characterize two conditions under which the “complete-sharing” (CS) policy that is often used in practice is optimal: The first is in the “off-season,” when the overall demand for service is low relative to capacity; the second is in the “peak season” of high demand, given that different customer classes are sufficiently similar.

3. We analyze a fluid approximation to the original system, and we derive closed-form expressions that characterize the controls and the performance obtained when allocating capacity using an “aggregate threshold” (AT) policy. These expressions allow us to efficiently calculate

admission thresholds that appear to perform well in the original stochastic model.

4. Closed-form expressions for the fluid model also allow us to demonstrate the concavity of the fluid model’s revenues with respect to the fleet size when the aggregate threshold policy is used. This concavity is the essential property required for the efficient solution of the related, long-term problem of capacity sizing.

5. We show that, in the presence of capacity rationing, the optimal fleet size can be either higher or lower than that obtained when no rationing is employed. The relationship between the two fleet sizes varies systematically with the cost of capacity.

6. We present numerical experiments that highlight the potential benefit of jointly optimizing fleet size and tactical controls. In particular, there appear to be cases in which the suboptimal use of complete sharing results in near-optimal dollar profits. Even in these cases, however, the return on investment in capacity suffers significantly.

More broadly, these numerical results complement our characterization of sufficient conditions for the optimality of CS policies. CS policies maximize physical measures of system utilization. When complete sharing is optimal, this physical measure of system utilization is a good proxy for economic utilization. When complete sharing is not optimal, however, its use can degrade profit margins and, by extension, economic measures of resource efficiency, such as return on assets. In this case, physical and economic measures of efficiency do not coincide.

Thus, within the context of the stationary problem developed in this paper, we are able to characterize how the use of tactical controls affects longer-term decisions regarding fleet size, as well as longer-term and economic efficiency. While a complete analysis of the problem, which should account for seasonal changes in demand patterns, is beyond the scope of this paper, our current results represent a promising first step.

Finally, we note that our analysis and results complement that of two recent papers that have independently considered the stochastic knapsack problem. Our analysis parallels that of Altman et al. (2001), which uses dynamic programming techniques to study optimal capacity allocation rules and develops and solves (numerically) a fluid approximation to the problem. Our special problem structure, however, allows us to more fully characterize properties of optimal and heuristic admission controls. We are able to develop a number of additional useful structural results concerning optimal policies and to develop precise, closed-form characterizations in the context of fluid control. Örmeci et al. (2001) also uses dynamic programming techniques to develop similar characterizations of structural properties of the optimal policy. It does not, however, consider heuristic controls. Neither of these papers considers how the use of tactical controls affects longer-term fleet-sizing decisions.

The remainder of this paper is organized as follows. In the next section, we formulate and analyze the capacity allocation problem and demonstrate how the problem

with lump-sum penalties can be reduced to one without penalties. We also discuss properties of optimal capacity allocation policies and establish conditions for the optimality of the CS policy. In §3, we introduce a heuristic AT policy based on a fluid-model version of our system, and we compare the performance of this heuristic to that of the optimal policy. In §4, we investigate the interaction between capacity sizing and capacity allocation problems and establish how optimal fleet capacity changes in the presence of capacity rationing. We then conclude with a discussion of the results and describe open issues and worthwhile extensions. All proofs may be found in the appendix at <http://or.pubs.informs.org/Pages/collect.html>.

2. The Capacity Allocation Problem

In this section, we analyze the capacity allocation decision. We formulate it as a problem in the control of queues, and we use dynamic programming techniques to investigate properties of the optimal control policies.

2.1. Model Description

Consider a fleet of c identical vehicles or pieces of rental equipment accessed by two customer classes whose arrival processes are independent and Poisson with intensities λ_1 and λ_2 . Let the durations of their rentals be independent, exponentially distributed random variables of mean μ_1^{-1} and μ_2^{-1} . Suppose, further, that each arrival wishes to rent exactly one unit of capacity.

At each arrival epoch a system controller, such as the manager of the rental location, can decide whether or not to admit an arriving customer for service—if one of the c units of capacity is free—or to reject the arrival. Arrivals that are admitted to service are permitted to complete the duration of their (randomly distributed) rental periods uninterrupted. Rejected customers do not queue; they exit the system. Similarly, customers that arrive when all c units of capacity are rented are lost.

Rewards and penalties associated with the system state and action are as follows. Arrivals that are admitted to service pay respective rental fees of $\$a_1$ and $\$a_2$ per unit of time. When a customer’s rental request is denied—either due to the absence of available rental capacity or because of the particular capacity allocation policy used—a lump-sum penalty of $\$\pi_1$ or $\$\pi_2$ is incurred, depending on the customer’s class. (For more on rejection penalties and their relationship to service-level constraints, please see Appendix A.)

The assumption that interarrival and service times are exponentially distributed implies that, at times between these event epochs, the system evolves as a continuous-time Markov chain. At these times, the system state can be completely described by the numbers of class-1 and class-2 customers renting units. Furthermore, system control—in the

form of acceptance or rejection of an arriving customer—is exercised only at arrival epochs, and it is sufficient to consider only the discrete-time process embedded at arrival and departure epochs when determining the form of effective system controls (see Puterman 1994, Chapter 11). That is, the system can be modeled as a discrete-time Markov Decision Process (MDP).

In Appendix B, we formally define discounted and average-cost versions of this MDP. For both cases, we also indicate why there exist stationary, deterministic policies that are optimal. Therefore, we will only consider policies of this class. Furthermore, rather than directly analyze the MDPs’ objective functions, we use well-known results concerning the convergence of the value-iteration procedure to analyze the problems.

2.2. Value-Iteration Formulation

We begin our definition of the value-iteration procedure by “uniformizing” the system. (See Lippman 1975 and Serfozo 1979.) Formally, we let $\Gamma = \lambda_1 + \lambda_2 + c\mu_1 + c\mu_2$ and, for the discounted problem with a continuous-time discount rate of α , we uniformize the system at rate $\alpha + \Gamma$.

Without loss of generality, we can define the time unit so that $\alpha + \Gamma = 1$. Thus, $\lambda_i \equiv \lambda_i/(\alpha + \Gamma)$ and $\mu_i \equiv \mu_i/(\alpha + \Gamma)$ become, respectively, the probability that the next uniformized transition is a type- i arrival or service completion. Similarly, $a_i \equiv a_i/(\alpha + \Gamma)$ is the expected discounted revenue per type- i rental until the time of the next uniformized transition.

Note that the uniformization rate includes the discount factor, α . In fact, it is well known that discounting at rate α is equivalent to including a constant intensity at which the process terminates, after which no more profits will be earned. Thus, one may think of α as the per-period probability that the next transition is a terminating one. (For example, see Puterman 1994, §5.3.)

The rate also includes rental completions of “phantom” customers. For example, if the current system state is (k_1, k_2) , then the probability that one of $(c - k_1)$ phantom type-1 customers or $(c - k_2)$ phantom type-2 customers completes a rental is $(c - k_1)\mu_1 + (c - k_2)\mu_2$. At the end of such a phantom rental, the observed state remains the same, (k_1, k_2) .

Given these uniformized system parameters, we define the value-iteration operator T as

$$\begin{aligned}
 Tf(k_1, k_2) &= a_1k_1 + a_2k_2 + \lambda_1H_1[f(k_1, k_2)] + \lambda_2H_2[f(k_1, k_2)] \\
 &\quad + \mu_1k_1f(k_1 - 1, k_2) + \mu_2k_2f(k_1, k_2 - 1) \\
 &\quad + ((\mu_1 + \mu_2)c - \mu_1k_1 - \mu_2k_2)f(k_1, k_2). \tag{1}
 \end{aligned}$$

The heart of the procedure is carried out via the maximizations

$$H_1[f(k_1, k_2)] = \begin{cases} \max[f(k_1, k_2) - \pi_1, f(k_1 + 1, k_2)] & \text{when } k_1 + k_2 < c, \\ f(k_1, k_2) - \pi_1 & \text{when } k_1 + k_2 = c, \end{cases} \tag{2}$$

and

$$H_2[f(k_1, k_2)] = \begin{cases} \max[f(k_1, k_2) - \pi_2, f(k_1, k_2 + 1)] & \text{when } k_1 + k_2 < c, \\ f(k_1, k_2) - \pi_2 & \text{when } k_1 + k_2 = c, \end{cases} \quad (3)$$

which are specified for any function f defined on the state space $S = \{(k_1, k_2) \in Z^2 \mid k_1 \geq 0, k_2 \geq 0, k_1 + k_2 \leq c\}$.

Let $v_0(k_1, k_2) \equiv 0$ represent an initial estimate of the optimal expected discounted profit, and v_n represent the estimated value after n iterations of the value-iteration algorithm:

$$\begin{aligned} v_n(k_1, k_2) &= a_1 k_1 + a_2 k_2 + \lambda_1 H_1[v_{n-1}(k_1, k_2)] + \lambda_2 H_2[v_{n-1}(k_1, k_2)] \\ &\quad + \mu_1 k_1 v_{n-1}(k_1 - 1, k_2) + \mu_2 k_2 v_{n-1}(k_1, k_2 - 1) \\ &\quad + ((\mu_1 + \mu_2)c - \mu_1 k_1 - \mu_2 k_2) v_{n-1}(k_1, k_2). \end{aligned} \quad (4)$$

Then, the fact that

$$\lambda_1 + \lambda_2 + (\mu_1 + \mu_2)c < 1 \quad (5)$$

for $\alpha > 0$ ensures that T is a contraction operator and that $\{v_n\}$ converges to the optimal “value function”

$$\begin{aligned} v(k_1, k_2) &= a_1 k_1 + a_2 k_2 + \lambda_1 H_1[v(k_1, k_2)] + \lambda_2 H_2[v(k_1, k_2)] \\ &\quad + \mu_1 k_1 v(k_1 - 1, k_2) + \mu_2 k_2 v(k_1, k_2 - 1) \\ &\quad + ((\mu_1 + \mu_2)c - \mu_1 k_1 - \mu_2 k_2) v(k_1, k_2), \end{aligned} \quad (6)$$

whose value equals that of the MDP’s optimal objective function (see Porteus 1982).

The first two terms on the right-hand side of (6) represent the expected discounted revenue earned until the next uniformized transition. The following four represent the probabilities and associated profits-to-go associated with system arrivals and service completions. The last term represents the probability and profit-to-go of a “phantom” rental completion. (Without loss of generality, we omit the probability, α , and value, 0, associated with a terminating transition.)

If no rejection penalties are used ($\pi_1 = \pi_2 = 0$), then (6) directly reduces to the *stochastic knapsack problem*, well known from the telecommunications literature (Ross and Tsang 1989). Furthermore, for any given rental fees and penalty values (a_1, a_2, π_1, π_2) , there exists an equivalent stochastic knapsack formulation with adjusted rental fees: $(\hat{a}_1, \hat{a}_2, 0, 0)$.

THEOREM 1. *For any problem with rewards and penalties (a_1, a_2, π_1, π_2) , and optimal value function $v(k_1, k_2)$, there exists an alternative formulation with rewards*

$$\hat{a}_i = a_i + \pi_i(\mu_i + \alpha), \quad i = 1, 2, \quad (7)$$

zero penalties, and optimal value function $\hat{v}(k_1, k_2)$, for which

$$\begin{aligned} \hat{v}(k_1, k_2) &= v(k_1, k_2) + \left(\lambda_1 \frac{\pi_1}{\alpha} + \lambda_2 \frac{\pi_2}{\alpha} \right) \\ &\quad + \pi_1 k_1 + \pi_2 k_2. \end{aligned} \quad (8)$$

Furthermore, a policy is optimal for the original problem if and only if it is optimal for the transformed problem with adjusted revenues and zero penalties.

Therefore, in the analysis that follows we will consider only the transformed problem $\hat{v}_n(k_1, k_2)$ with adjusted fees (\hat{a}_1, \hat{a}_2) . Observe that the adjustment factors are linear in the penalty values and the service rates.

We note that the paper’s numerical results are performed using an average-cost MDP formulation. (Because they do not depend on the starting state, “average-cost” results are easier than discounted results to interpret.) In this case, a similar result holds, with

$$\hat{a}_i = a_i + \mu_i \pi_i, \quad i = 1, 2. \quad (9)$$

For a formal development of the value-iteration procedure and the analogue of Theorem 1 for the average-cost problem, please see Appendix C at <http://or.pubs.informs.org/Pages/collect.html>.

2.3. Optimality of Switching-Curve Policies

To establish structural properties of the optimal control policy, it is sufficient to show that certain properties of the functions defined on S are preserved under the action of the value-iteration operator, T (see Porteus 1982). In particular, we are interested in submodularity. We say that $f(k_1, k_2)$ is submodular in k_1 and k_2 if

$$\begin{aligned} f(k_1 + 1, k_2 + 1) - f(k_1, k_2 + 1) \\ \leq f(k_1 + 1, k_2) - f(k_1, k_2), \quad k_1 + k_2 + 2 \leq c. \end{aligned} \quad (10)$$

Let F be the set of all f defined on S that are submodular in k_1 and k_2 .

The following theorem states that F is closed under T , so that the value-iteration operator preserves submodularity of the value function. This, in turn, implies that the optimal capacity allocation policy is of a special form; it is a “switching curve” policy.

THEOREM 2 (ALTMAN ET AL. 2001, ÖRMECI ET AL. 2001, SAVIN 2001).

- (a) $f \in F \Rightarrow Tf \in F$, and therefore $\hat{v}(k_1, k_2) \in F$.
- (b) In turn, for each k_1 it is optimal to admit customers of class 1 when in state (k_1, k_2) if and only if $k_2 < k_2^{\min}(k_1)$, where

$$\begin{aligned} k_2^{\min}(k_1) &= \begin{cases} c - k_1 & \text{if } \hat{v}(k_1 + 1, c - k_1 - 1) > \hat{v}(k_1, c - k_1 - 1), \\ \min\{k_2 : 0 \leq k_2 \leq c - k_1 - 1, \hat{v}(k_1 + 1, k_2) \leq \hat{v}(k_1, k_2)\} & \text{otherwise.} \end{cases} \end{aligned}$$

Similarly, for each k_2 it is optimal to admit customers of class 2 when in state (k_1, k_2) if and only if $k_1 < k_1^{\min}(k_2)$, where

$$k_1^{\min}(k_2) = \begin{cases} c - k_2 & \text{if } \hat{v}(c - k_2 - 1, k_2 + 1) > \hat{v}(c - k_2 - 1, k_2), \\ \min(k_1 : 0 \leq k_1 \leq c - k_2 - 1, \hat{v}(k_1, k_2 + 1) \leq \hat{v}(k_1, k_2)) & \\ \text{otherwise.} \end{cases}$$

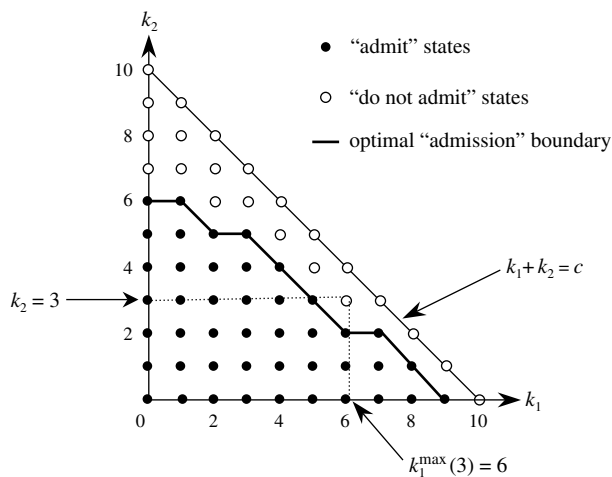
Part (b) of the theorem can be interpreted as follows: When a given number of customers of a particular class is already renting equipment, the “next” customer of the same class is admitted if and only if the number of customers of the other class present in the system does not exceed some critical value. This is a switching-curve policy, characterized by c critical indices for each of the customer classes.

For the average-cost case we can develop analogous results. At every pass of the value-iteration procedure, the operator preserves the submodularity of the estimate of the value function. This ensures that the results of Theorem 2 apply to the optimal control policy for this case as well.

As an illustration of the optimal capacity allocation policies, we consider an example with $\hat{a}_1 = 10$, $\hat{a}_2 = 5$, $\lambda_1 = 25$, $\lambda_2 = 10$, $\mu_1 = 5$, $\mu_2 = 1$, and $c = 10$ for the case when average revenue per period is maximized. Figure 1 describes the capacity allocation decisions for class-2 customers and illustrates the notion of the “switching curve.”

One feature of this example worth noting is the following: class-1 customers are always allowed to rent equipment, i.e., $k_2^{\min}(k_1) = c - k_1$ for all feasible k_1 (and so we did not include the graph of optimal allocation for class 1). In this case, we say that class-1 customers are a preferred class. While in every numerical example we tested there existed a preferred class, we have not been

Figure 1. The optimal capacity allocation policy for class-2 customers when the average adjusted revenue per period is maximized ($\hat{a}_1 = 10$, $\hat{a}_2 = 5$, $\lambda_1 = 25$, $\lambda_2 = 10$, $\mu_1 = 5$, $\mu_2 = 1$, and $c = 10$).



able to prove that such a class exists universally. Nevertheless, we have been able to characterize a great deal about preferred-customer classes.

2.4. Preferred Classes and the Optimality of the Complete-Sharing Policy

In this section, we investigate the conditions that make a particular customer class a preferred one. Closely connected to the question about the nature of preferred classes is the issue of the optimality of the CS policy: CS is optimal when both customer classes are preferred. The following theorem provides sufficient conditions under which one—or both—classes may be preferred.

THEOREM 3. (a) Define $\lambda = \lambda_1 + \lambda_2$, $\bar{\mu} = \min(\mu_1, \mu_2)$, $\bar{a} = \max(\hat{a}_1, \hat{a}_2)$, and

$$c_i^* = 2 + \frac{\lambda}{\bar{\mu}} \left(\frac{\bar{a}}{\hat{a}_i} \left(\frac{\mu_i + \alpha}{\mu_i} \left(6 + 4 \left(\frac{\lambda + 2\bar{\mu}}{\mu_i} \right) \right) + \frac{\mu_i}{\bar{\mu} + \alpha} \left(2 + \frac{\lambda + 2\bar{\mu}}{\mu_i + \alpha} \right) \right) - 1 \right), \quad i = 1, 2. \quad (11)$$

Then, for systems with capacity $c > c_i^*$, it is always optimal to admit class- i customers, $i = 1, 2$.

(b) In turn, for $c \geq \max(c_1^*, c_2^*)$, the policy of CS of the service fleet is optimal.

Theorem 3 provides a lower bound on the amount of capacity sufficient to ensure that a particular customer class (or both classes) has unrestricted access to the available equipment. Of course, for profit-maximizing firms, capacity costs may prevent c from becoming large enough to optimally operate in the complete-sharing regime. In §4, we investigate the interaction among capacity cost, fleet size, and tactical control in more detail.

We note that for each customer class this lower bound is, as expected, a nonincreasing function of the penalty-adjusted fee paid by customers of this class. We observe that in the simple case of $\mu_1 = \mu_2 \gg \alpha$, (11) implies that $c_1^*, c_2^* \gg \lambda/\bar{\mu}$. Thus, in the presence of seasonal demand patterns, these results describe the “off-peak” season when the demand for rentals may be significantly lower than the available capacity.

Note that Theorem 3 is stronger than a limiting statement. In general, it is not hard to imagine that as $c \rightarrow +\infty$, a CS policy will be asymptotically optimal. Theorem 3, however, says that there is a fixed, finite c above which CS is optimal. This is because, as more and more pieces of equipment are rented, the probability that the next event is a service completion, rather than an arrival, grows. Thus, the busier the system, the stronger its drift toward emptying out. For large enough c , the expected loss of revenue due to blocking becomes small when compared to the immediate gain of taking the next customer, no matter which class she belongs to.

Theorem 3 states that for sufficiently high service capacity, the CS policy is optimal. It is also possible to show

that CS is optimal even in the “peak season,” when capacity is tight, provided that the customer classes are similar in terms of their penalty-adjusted rental fees:

THEOREM 4. For either class $i \in \{1, 2\}$, and $j \neq i$, if

$$\frac{\hat{a}_i}{\max[\mu_i, \mu_j]} \geq \frac{\lambda_j}{\lambda_j + \mu_i} \frac{\hat{a}_j}{\mu_j}, \quad (12)$$

then it is always optimal to admit type- i customers.

The statement of Theorem 4 is intuitively appealing: All other parameters of the problem being fixed, there exists a minimum value of the adjusted rental fee \hat{a}_i that ensures that customers of this class should be freely admitted into the system. CS of service fleet is optimal when (12) is satisfied for both classes, i.e., when \hat{a}_1 and \hat{a}_2 are “close.”

Furthermore, recall that $\hat{a}_i = a_i + \pi_i(\mu_i + \alpha)$ depends on both the revenue earned when accepting a class- i customer and the penalty paid when rejecting class- i demand. That is, a preferred customer may be profitable to serve, unprofitable not to serve, or some combination of the two. For example, a high-volume customer, such as a national account, may receive a favorable rental rate in return for a large stream of rentals. At the same time, contractual service-level requirements or the customer’s market power may imply a large rejection penalty so that class- i arrivals become VIP. (For more on the relationship between service-level constraints and rejection penalties, see Appendix A at <http://or.pubs.informs.org/Pages/collect.html>.)

The sufficient conditions of Theorem 4 are direct analogues to expressions for protection levels in airline seat-allocation models. (For example, see Belobaba 1989.) Both sets of inequalities can be interpreted in terms of simple marginal analysis. For instance, for $i = 1$, the right-hand side of (12) describes (a bound on) the expected cost of admitting an arriving class-1 customer. It is the expected revenue lost from a blocked class-2 customer that might have been served. Here $\lambda_2/(\lambda_2 + \mu_1)$ is the probability that a class-2 arrives before the admitted class-1 finishes service, and \hat{a}_2/μ_2 is the expected revenue lost, given the blocking occurs.

In fact, Örmeci et al. (2001) develops a characterization of preferred classes that mirrors this “marginal analysis” result. The left-hand side of (12) is more complex—and more stringent—than simply \hat{a}_1/μ_1 , however. This difference better reflects the more complex dynamics of our system.

Observe that there exists a broad range of circumstances under which a class of customers may be preferred. First, note that if $\hat{a}_i > \hat{a}_j$ and $\hat{a}_i/\mu_i \geq \hat{a}_j/\mu_j$, then type- i customers have higher penalty-adjusted rental rates and higher expected rental durations—and they are preferred. Second, even though $\hat{a}_j/\mu_j \leq \hat{a}_i/\mu_i$, type- j customers may also be preferred, as long as \hat{a}_j is not too far below \hat{a}_i .

Conversely, it is possible to construct examples in which neither of the sufficient conditions of Theorem 4 is satisfied. This occurs when $\hat{a}_i > \hat{a}_j$, $\mu_i > \mu_j$, and $\hat{a}_i/\mu_i < \hat{a}_j/\mu_j$.

Of course, failure to satisfy the sufficient conditions does not demonstrate that there exists no preferred class.

Finally, we note that the conditions of Theorem 4 are broadly applicable in that they do not depend on the service capacity, c , or on the intensity of arrivals of the customer class being considered for admission. The required parameters are simple to estimate from observable data, and the results are simple to interpret.

3. Heuristic Capacity Allocation Policies

In general, it is optimal to base the control of admissions into the service on the numbers of customers of both classes 1 and 2 that are in the system at the time each control decision is made. In practice, however, these “vector” policies may be difficult to implement, especially for rental systems with large capacities.

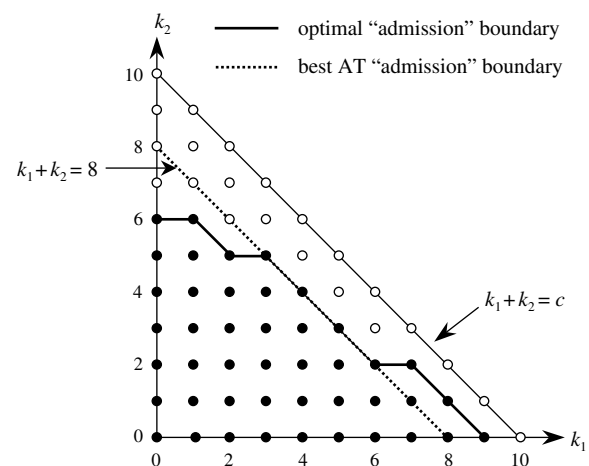
Admission control decisions that are based on the value of a particular scalar metric derived from this vector state, rather than the detailed state of the system, may also provide effective (if suboptimal) controls. One of the most widely used heuristics is the aggregate threshold (trunk reservation) policy.

The AT policy assumes that there exists a preferred-customer class, and it is the class that offers higher revenue per unit of time. The AT policy admits second-class customers as long as the total number of customers already in the system does not exceed some critical threshold value.

Besides being intuitively appealing, AT policies have been proven to be optimal whenever $\mu_1 = \mu_2$ (see Miller 1969). More generally, we expect them to perform well in cases when the expected service times for different customer classes are similar.

Figure 2 illustrates the best AT policy, as well as the optimal control policy, for the same example shown in Figure 1. While the control exercised by the AT policy differs

Figure 2. The optimal and the best AT policies for class-2 customers ($\hat{a}_1 = 10$, $\hat{a}_2 = 5$, $\lambda_1 = 25$, $\lambda_2 = 10$, $\mu_1 = 5$, $\mu_2 = 1$, and $c = 10$).



from that of the optimal policy, the revenues it generates are nearly optimal, falling below optimality by about 0.15%.

AT policies, however, do not yield closed-form expressions for system performance measures. In general, the task of computing the value of the best AT level can be comparable in its complexity to the task of computing the optimal control policy.

Ideally, we would like to have a policy that combines ease of calculation with the robust performance of AT controls. In the following section, we develop such a heuristic. It uses a fluid-model approximation of the stochastic model to derive closed-form expressions for the AT values.

3.1. Fluid Models and Scaling

In many practical situations, both the size of rental fleet c and the offered rental intensities $\rho_1 = \lambda_1/\mu_1$ and $\rho_2 = \lambda_2/\mu_2$ are large. Under these conditions, a deterministic fluid model may offer a good approximation to the original control problem. Indeed, Altman et al. (2001) offer a heuristic derivation of such a fluid model as the limit of a linearly scaled sequence of MDPs, and they numerically evaluate the resulting Hamilton-Bellman-Jacobi equations.

We follow the approach of Altman et al. (2001), but given the underlying structure of our problem, in which there are two classes of customers, we can directly analyze the trajectory of the fluid system. This allows us to develop an AT heuristic whose performance is robust and whose closed-form expressions allow for immediate calculation of policy parameters. Furthermore, our analysis also allows us to demonstrate the concavity of discounted revenues (of a “ μ -scaled” version of our model), with respect to the fleet size, c , a property that becomes important in the capacity-sizing analysis of §4.

We start by defining the state space and dynamics for fluid approximations (in general). Time t is continuous, and the state parameters $k_1(t)$ and $k_2(t)$ of the original model become continuous-state variables, restricted to set $S_f = \{k_1(t) \geq 0, k_2(t) \geq 0, k_1(t) + k_2(t) \leq c\}$. Poisson customer arrivals are replaced by the deterministic continuous “flow” arrivals with intensities λ_1 and λ_2 . The departure process becomes deterministic as well: For the state $(k_1(t), k_2(t))$ it is represented by an outflow at rate $\mu_1 k_1(t) + \mu_2 k_2(t)$.

Arrivals are controlled as follows: At time t , a control policy $(u_1(t), u_2(t))$ results in the total customer inflow of $u_1(t)\lambda_1 + u_2(t)\lambda_2$. Thus, for the control trajectories $(u_1(t), u_2(t))$ ($0 \leq u_i(t) \leq 1, i = 1, 2$), the Kolmogorov evolution equations for the original system are replaced by

$$\begin{aligned} \frac{dk_1(t)}{dt} &= u_1(t)\lambda_1 - \mu_1 k_1(t) \quad \text{and} \\ \frac{dk_2(t)}{dt} &= u_2(t)\lambda_2 - \mu_2 k_2(t), \end{aligned} \quad (13)$$

with a constraint that reflects the finite size of the service fleet

$$\begin{aligned} \lambda_1 u_1(t) + \lambda_2 u_2(t) &\leq \mu_1 k_1(t) + \mu_2 k_2(t), \\ i = 1, 2, \quad \text{whenever } k_1(t) + k_2(t) &= c. \end{aligned} \quad (14)$$

The total discounted revenue is then the objective to be maximized. If at $t = 0$ the system is in the state (k_1, k_2) , then—for a feasible (under (14)) control policy Δ that uses $(u_1(t), u_2(t))$ —the total discounted revenue is

$$\begin{aligned} \hat{R}_\alpha(k_1, k_2, \Delta) &= \int_0^\infty (\hat{a}_1 k_1(t) + \hat{a}_2 k_2(t)) e^{-\alpha t} dt \\ &= \frac{\hat{a}_1 k_1}{\mu_1 + \alpha} + \frac{\hat{a}_2 k_2}{\mu_2 + \alpha} + R_\alpha(k_1, k_2, \Delta), \end{aligned} \quad (15)$$

where

$$R_\alpha(k_1, k_2, \Delta) = \int_0^\infty \left(\frac{\hat{a}_1 \lambda_1 u_1(t)}{\mu_1 + \alpha} + \frac{\hat{a}_2 \lambda_2 u_2(t)}{\mu_2 + \alpha} \right) e^{-\alpha t} dt \quad (16)$$

is the part of the revenue that actually depends on the control policy chosen. In what follows, the term “revenue” is used to designate $R_\alpha(k_1, k_2, \Delta)$.

Our AT heuristic is based on a “scaled” version of the fluid model:

DEFINITION 1. A μ -scaled version of the fluid model with parameters $\lambda_1, \lambda_2, \mu_1$, and μ_2 is the problem with parameters $\lambda_1^s = \lambda_1 \mu / \mu_1$, $\lambda_2^s = \lambda_2 \mu / \mu_2$, and $\mu_1^s = \mu_2^s = \mu$ for $\mu \in [\mu_1, \mu_2]$.

Note that in every μ -scaled version of the fluid model, the departure rates of both customer classes are equal and $\lambda_1^s / \mu_1^s = \lambda_1 / \mu_1$, $\lambda_2^s / \mu_2^s = \lambda_2 / \mu_2$. Because the departure rates of both classes are the same, one can use arguments similar to those in Miller (1969) to show that the optimal admission control decisions only depend on the total number of customers $k(t) = k_1(t) + k_2(t)$ in the system. Thus, given $\hat{a}_1 > \hat{a}_2$, a control policy that admits as many class-1 customers as possible and limits the admissions of class-2 customers is optimal for any μ -scaled problem.

3.2. Fluid Aggregate Threshold Heuristic

In the μ -scaled model, system dynamics simplify to

$$\frac{dk(t)}{dt} = u_1(t)\lambda_1^s + u_2(t)\lambda_2^s - \mu k(t). \quad (17)$$

In turn, a fluid analogue of the original, stochastic system’s AT policy admits class 2 customers if and only if the total system occupancy, $k(t)$, does not exceed a “fluid aggregate threshold” (FAT), k_{FAT} . When $\rho_1 \geq c$ or $\rho_1 + \rho_2 \leq c$, such a FAT policy is a direct analogue of AT policies in the original stochastic system. For $\rho_1 < c < \rho_1 + \rho_2$, however, there does not exist a neat correspondence. Therefore, in the following sections we define and analyze the FAT policy within each subset of the relevant parameter range.

3.2.1. The FAT Policy When $\rho_1 \geq c$. For systems with $\rho_1 \geq c$, class-1 traffic alone is sufficient to ensure complete utilization of the rental fleet, and a threshold policy can be defined and analyzed in a straightforward fashion. In this case, the control $(u_1(t), u_2(t))$ is defined as follows:

$$(u_1(t), u_2(t)) = \begin{cases} (1, 1) & \text{for } k(t) < k_{\text{FAT}}, \\ (1, 0) & \text{for } k_{\text{FAT}} \leq k(t) < c, \\ (c\mu/\lambda_1^s, 0) & \text{for } k(t) = c. \end{cases} \quad (18)$$

Note that, once the system hits the boundary and $k(t) = c$, customers continue to be admitted at the maximum feasible rate, and the system state remains at the boundary thereafter.

Control (18) then implies that, at time t , the revenue-generation rate

$$r(t) = \frac{\hat{a}_1 \lambda_1^s u_1(t) + \hat{a}_2 \lambda_2^s u_2(t)}{\mu + \alpha} \quad \text{for } \rho_1 \geq c$$

is given by

$$r(t | \rho_1 \geq c) = \begin{cases} \frac{\hat{a}_1 \lambda_1^s + \hat{a}_2 \lambda_2^s}{\mu + \alpha} & \text{for } k(t) < k_{\text{FAT}}, \\ \frac{\hat{a}_1 \lambda_1^s}{\mu + \alpha} & \text{for } k_{\text{FAT}} \leq k(t) < c, \\ \frac{\hat{a}_1 \mu c}{\mu + \alpha} & \text{for } k(t) = c. \end{cases} \quad (19)$$

To compute the total discount revenues for a given k_{FAT} , we must also account for the starting state $k \equiv k(0)$.

When $k < k_{\text{FAT}} \leq c$, there are three elements to the discounted revenues: those earned as $k(t)$ approaches k_{FAT} ; those earned when $k_{\text{FAT}} \leq k(t) \leq c$; and those earned after the boundary has been hit. We calculate each in turn. Let

$$t_{\text{FAT}} = \frac{1}{\mu} \ln \left(\frac{\rho_1 + \rho_2 - k}{\rho_1 + \rho_2 - k_{\text{FAT}}} \right)$$

be the time that the system state hits k_{FAT} , so that $k(t_{\text{FAT}}) = k_{\text{FAT}}$. Then, from (19) we have

$$\begin{aligned} & \int_0^{t_{\text{FAT}}} \left(\frac{\hat{a}_1 \lambda_1^s u_1(t) + \hat{a}_2 \lambda_2^s u_2(t)}{\mu + \alpha} \right) e^{-\alpha t} dt \\ &= \left(\frac{\hat{a}_1 \lambda_1^s + \hat{a}_2 \lambda_2^s}{\mu + \alpha} \right) \frac{1 - \exp(-\alpha t_{\text{FAT}})}{\alpha}. \end{aligned} \quad (20)$$

Similarly, let

$$t_c = t_{\text{FAT}} + \frac{1}{\mu} \ln \left(\frac{\rho_1 + \rho_2 - k_{\text{FAT}}}{\rho_1 + \rho_2 - c} \right)$$

be the time at which the system state hits c , so that $k(t_c) = c$. Then, using (19), we have

$$\begin{aligned} & \int_{t_{\text{FAT}}}^{t_c} \left(\frac{\hat{a}_1 \lambda_1^s u_1(t) + \hat{a}_2 \lambda_2^s u_2(t)}{\mu + \alpha} \right) e^{-\alpha t} dt \\ &= \left(\frac{\hat{a}_1 \lambda_1^s}{\mu + \alpha} \right) \frac{\exp(-\alpha t_{\text{FAT}}) - \exp(-\alpha t_c)}{\alpha}. \end{aligned} \quad (21)$$

Finally, from (19) the revenues earned after reaching the boundary are given by

$$\begin{aligned} & \int_{t_c}^{+\infty} \left(\frac{\hat{a}_1 \lambda_1^s u_1(t) + \hat{a}_2 \lambda_2^s u_2(t)}{\mu + \alpha} \right) e^{-\alpha t} dt \\ &= \left(\frac{\hat{a}_1 \mu c}{\mu + \alpha} \right) \frac{\exp(-\alpha t_c)}{\alpha}. \end{aligned} \quad (22)$$

Collecting the revenue terms (20)–(22), substituting for t_{FAT} and t_c , and simplifying, we then obtain the discounted revenues for the FAT policy when $\rho_1 \geq c$ and $k \leq k_{\text{FAT}} < c$:

$$\begin{aligned} R_{\alpha}^{\text{FAT}}(k, k_{\text{FAT}} | \rho_1 \geq c, k \leq k_{\text{FAT}} < c) &= \frac{\mu}{\alpha(\alpha + \mu)} \left(\hat{a}_1 \rho_1 + \hat{a}_2 \rho_2 - \left(\frac{\rho_1 + \rho_2 - k_{\text{FAT}}}{\rho_1 + \rho_2 - k} \right)^{\alpha/\mu} \right. \\ &\quad \left. \cdot \left(\hat{a}_2 \rho_2 + \hat{a}_1 \frac{(\rho_1 - c)^{(\alpha + \mu)/\mu}}{(\rho_1 - k_{\text{FAT}})^{\alpha/\mu}} \right) \right). \end{aligned} \quad (23)$$

When $k_{\text{FAT}} \leq k < c$, only type-1 customers are admitted to the system. In this case, in the above analysis we replace t_{FAT} by 0 and k_{FAT} by k . Then, analogous calculations yield

$$\begin{aligned} R_{\alpha}^{\text{FAT}}(k, k_{\text{FAT}} | \rho_1 \geq c, k_{\text{FAT}} \leq k < c) &= \frac{\mu \hat{a}_1}{\alpha(\alpha + \mu)} \left(\rho_1 - \frac{(\rho_1 - c)^{(\alpha + \mu)/\mu}}{(\rho_1 - k)^{\alpha/\mu}} \right). \end{aligned} \quad (24)$$

3.2.2. FAT Policy When $\rho_1 + \rho_2 < c$. When $\rho_1 + \rho_2 < c$, a threshold policy with $k_{\text{FAT}} < c$ leads to incomplete utilization of the rental fleet and may be trivially improved by setting $k_{\text{FAT}} = c$ so that all customers are admitted for service, no matter what the initial state of the system, $k(0)$. Here, the policy is, again, a direct analogue of AT policies in the original stochastic system. Specifically, the optimal fluid threshold of c corresponds to complete sharing, an AT policy with a threshold of c .

Because $\rho_1 + \rho_2 < c$, even with no control the boundary $k(t) = c$ is never hit (for $t > 0$). In this case, the optimal control is

$$(u_1(t), u_2(t)) = (1, 1)$$

for any system state, $k(t)$, and the rate at which revenue is earned is

$$r(t | \rho_1 + \rho_2 < c) = \frac{\hat{a}_1 \lambda_1^s + \hat{a}_2 \lambda_2^s}{\mu + \alpha}.$$

In turn, the revenue calculation is

$$\begin{aligned} R_{\alpha}^{\text{FAT}}(k, c | \rho_1 + \rho_2 < c) &= \int_0^{\infty} \left(\frac{\hat{a}_1 \lambda_1^s u_1(t) + \hat{a}_2 \lambda_2^s u_2(t)}{\mu + \alpha} \right) e^{-\alpha t} dt \\ &= \frac{\mu}{\alpha(\mu + \alpha)} (\hat{a}_1 \rho_1 + \hat{a}_2 \rho_2). \end{aligned} \quad (25)$$

3.2.3. FAT Policy When $\rho_1 < c \leq \rho_1 + \rho_2$. Finally, when $\rho_1 < c \leq \rho_1 + \rho_2$, there does not appear to exist a fluid analogue of a threshold policy that is both effective and straightforward to implement. On the one hand, a threshold of $k_{\text{FAT}} < c$ results in incomplete utilization of the rental fleet and can be improved upon by admitting *some* class-2

customers. On the other hand, setting $k_{\text{FAT}} = c$ and admitting all class-2 customers is infeasible because the maximum rate at which the system can be cleared is strictly less than the rate at which customers are arriving: $c\mu < \lambda_1^s + \lambda_2^s$.

In this case, a natural interpretation of the threshold rule defines a “soft” threshold when $k(t) = c$, one that limits, but does not eliminate, the flow of class-2 customers into the system:

$$(u_1(t), u_2(t)) = \begin{cases} (1, 1) & \text{for } k(t) < c, \\ \left(1, \frac{\mu c - \lambda_1^s}{\lambda_2^s}\right) & \text{for } k(t) = c, \end{cases} \quad (26)$$

so that

$$r(t \mid \rho_1 < c \leq \rho_1 + \rho_2) = \begin{cases} \frac{\hat{a}_1 \lambda_1^s + \hat{a}_2 \lambda_2^s}{\mu + \alpha} & \text{for } k(t) < c, \\ \frac{\hat{a}_1 \lambda_1^s + \hat{a}_2 (\mu c - \lambda_1^s)}{\mu + \alpha} & \text{for } k(t) = c. \end{cases} \quad (27)$$

Thus, for $\rho_1 < c < \rho_1 + \rho + 2$, the control generates system behavior and revenue that differ from those when $k_{\text{FAT}} < c$ or $k_{\text{FAT}} = c$, and we denote this soft threshold as $k_{\text{FAT}} = c^-$.

Given $k_{\text{FAT}} = c^-$ and any $k \equiv k(0) \in [0, c]$, the system’s revenues can be split into two components: those earned before reaching c , and those earned after. If

$$t_c = \frac{1}{\mu} \ln \left(\frac{\rho_1 + \rho_2 - k}{\rho_1 + \rho_2 - c} \right)$$

is the time required for the system to reach the boundary, then the first revenue component in (27) gives us

$$\int_0^{t_c} \left(\frac{\hat{a}_1 \lambda_1^s u_1(t) + \hat{a}_2 \lambda_2^s u_2(t)}{\mu + \alpha} \right) e^{-\alpha t} dt = \left(\frac{\hat{a}_1 \lambda_1^s + \hat{a}_2 \lambda_2^s}{\mu + \alpha} \right) \frac{1 - \exp(-\alpha t_c)}{\alpha}. \quad (28)$$

After the full capacity is reached, we use the bottom revenue-generation rate within (27) to obtain

$$\int_{t_c}^{+\infty} \left(\frac{\hat{a}_1 \lambda_1^s u_1(t) + \hat{a}_2 \lambda_2^s u_2(t)}{\mu + \alpha} \right) e^{-\alpha t} dt = \left(\frac{\hat{a}_1 \lambda_1^s + \hat{a}_2 (\mu c - \lambda_1^s)}{\mu + \alpha} \right) \frac{\exp(-\alpha t_c)}{\alpha}. \quad (29)$$

Adding (28) and (29), and using the expression for t_c , we then have

$$R_\alpha^{\text{FAT}}(k \mid \rho_1 < c \leq \rho_1 + \rho_2) = \frac{\mu}{\alpha(\alpha + \mu)} \left((\hat{a}_1 \rho_1 + \hat{a}_2 \rho_2) - \hat{a}_2 \frac{(\rho_1 + \rho_2 - c)^{(\alpha/\mu) + 1}}{(\rho_1 + \rho_2 - k)^{\alpha/\mu}} \right). \quad (30)$$

3.2.4. Optimal Thresholds and Revenues for the FAT Policy. We can use the expressions we have derived for discounted revenues to determine both optimal thresholds and optimal discounted revenues. In both cases, we obtain simple, closed-form expressions.

First, we address the optimal threshold, k_{FAT}^* . For $\rho_1 \geq c$, its determination follows from differentiation of (23) with respect to k_{FAT} :

THEOREM 5. *The optimal value of the AT, k_{FAT}^* , is independent of the starting state, k , and is given by*

$$k_{\text{FAT}}^*(c) = \begin{cases} 0 & \text{for } c < \rho_1 (1 - (\hat{a}_2/\hat{a}_1)^{\mu/(\mu+\alpha)}), \\ c - (\rho_1 - c) \left((\hat{a}_1/\hat{a}_2)^{\mu/(\mu+\alpha)} - 1 \right) & \text{for } \rho_1 (1 - (\hat{a}_2/\hat{a}_1)^{\mu/(\mu+\alpha)}) \leq c \leq \rho_1, \\ c^- & \text{for } \rho_1 < c \leq \rho_1 + \rho_2, \\ c & \text{for } \rho_1 + \rho_1 < c. \end{cases} \quad (31)$$

We observe that, all other problem parameters being fixed, the optimal AT value described by (31) is a nondecreasing function of the fleet size c . In particular, if the available rental capacity falls below the critical value

$$c_{\min} = \rho_1 \left(1 - \left(\frac{\hat{a}_2}{\hat{a}_1} \right)^{\mu/(\mu+\alpha)} \right),$$

then it is optimal not to admit any class-2 customers into service. Conversely, if the rental capacity is sufficiently large, exceeding the offered load from class 1, then the control on admissions of class-2 customers should be postponed until the entire rental fleet is utilized. For the rental fleet values in between these two critical quantities, some form of admission control on class-2 customers is optimal, even in states in which some rental capacity is available. We observe that the critical index c_{\min} is a decreasing function of the ratio of penalty-adjusted rental fees \hat{a}_2/\hat{a}_1 .

When the time-discounting factor α is much smaller than μ , the optimal AT level, described in Theorem 5, is not particularly sensitive to the choice of μ . Even for rental durations of several months, the service rates (inverse of the expected service time) are about $\mu \simeq 10^{-3}$ per day and are at least an order of magnitude higher than any realistic values for α (for example, 30%–40% annual discounting rate results in $\alpha \simeq 10^{-4}$ per day). The same argument suggests that k_{FAT}^* is not sensitive to the choice of α . Thus, it is straightforward to use k_{FAT}^* as a threshold for both discounted and “average-cost” versions of the problem.

Using an expression for the optimal AT (31), we obtain:

THEOREM 6. *Given fixed λ_1^s , λ_2^s , μ , \hat{a}_1 , \hat{a}_2 , and α , define*

$$c_{\min} = \rho_1 \left(1 - \left(\frac{\hat{a}_2}{\hat{a}_1} \right)^{\mu/(\mu+\alpha)} \right).$$

(a) *If the rental system starts in state k , then the optimal total discounted revenue is*

$$R_\alpha^{\text{FAT}}(k, k_{\text{FAT}}^*(c))$$

$$= \begin{cases} \frac{\mu}{\alpha(\alpha + \mu)} \left(\hat{a}_1 \rho_1 - \hat{a}_1 \frac{(\rho_1 - c)^{(\alpha + \mu)/\mu}}{(\rho_1 - k)^{\alpha/\mu}} \right) & \text{for } c \leq c_{\min}, \\ \frac{\mu}{\alpha(\alpha + \mu)} \left(\hat{a}_1 \rho_1 + \hat{a}_2 \rho_2 - \hat{a}_2 \frac{(\rho_2 + (\rho_1 - c)(\hat{a}_1/\hat{a}_2)^{\mu/(\mu + \alpha)})^{(\alpha + \mu)/\mu}}{(\rho_1 + \rho_2 - k)^{\alpha/\mu}} \right) & \text{for } c_{\min} \leq c < \rho_1, k < k_{\text{FAT}}^*(c), \\ \frac{\mu}{\alpha(\alpha + \mu)} \left(\hat{a}_1 \rho_1 - \hat{a}_1 \frac{(\rho_1 - c)^{(\alpha + \mu)/\mu}}{(\rho_1 - k)^{\alpha/\mu}} \right) & \text{for } c_{\min} \leq c < \rho_1, k \geq k_{\text{FAT}}^*(c), \\ \frac{\mu}{\alpha(\alpha + \mu)} \left(\hat{a}_1 \rho_1 + \hat{a}_2 \rho_2 - \hat{a}_2 \frac{(\rho_1 + \rho_2 - c)^{(\alpha + \mu)/\mu}}{(\rho_1 + \rho_2 - k)^{\alpha/\mu}} \right) & \text{for } \rho_1 \leq c \leq \rho_1 + \rho_2, \\ \frac{\mu}{\alpha(\alpha + \mu)} (\hat{a}_1 \rho_1 + \hat{a}_2 \rho_2) & \text{for } \rho_1 + \rho_2 < c. \end{cases} \quad (32)$$

(b) For fixed values of rental fees, and demand and service parameters, $R_{\alpha}^{\text{FAT}}(k, k_{\text{FAT}}^*(c))$ is a nondecreasing concave function of the rental fleet size c for every $k \leq c$.

Inspection of (32) shows that $R_{\alpha}^{\text{FAT}}(k, k_{\text{FAT}}^*(c))$, like $k_{\text{FAT}}^*(c)$, is insensitive to the choice of μ for $\alpha \ll \mu$. (Of course, this insensitivity follows from the μ -scaled problem, not necessarily from the two-class problem in which $\mu_1 \neq \mu_2$.) Part (b) of Theorem 6 also states that, for any starting state, FAT revenues are concave in c . Thus, although the concavity of revenue with respect to fleet size is difficult to demonstrate in the context of the original MDP, it emerges naturally from the μ -scaled fluid approximation. This concavity property becomes important in the context of fleet-sizing decisions, which we discuss in §4.

3.3. Numerical Study of the Performance of the FAT Heuristic

Our motivation for developing the FAT heuristic was that it should perform well and be easy to implement. Therefore, to test the policy's performance we have undertaken a series of numerical studies that compare its average revenues to those obtained using the optimal control and the CS policy.

In two of the three cases analyzed above, translation of the FAT policy (31) to the context of a discrete, stochastic system is straightforward. For $\rho_1 \geq c$, we assume that $\mu/(\mu + \alpha) \approx 1$, when necessary, and then round the resulting k_{FAT}^* down to the nearest integer. For $\rho_1 + \rho_2 < c$, we set the aggregate system threshold equal to c , effectively implementing a CS policy.

When $\rho_1 < c < \rho_1 + \rho_2$, however, $k_{\text{FAT}}^* = c^-$, and the inflow of class-2 rentals is partially controlled. In this case, there is not a clear correspondence in a discrete system: Setting the AT to c implements CS, which does not control class-2 customers at all; conversely, setting the threshold to $c - 1$ completely stops the flow of class-2 customers at the boundary.

Because both alternatives of the FAT policy are trivial to compute, we include them both in our numerical tests. In

total, in each numerical experiment we test four policies: the optimal policy; FAT with c^- set to c (" $c^- = c$ "); FAT with c^- set to $c - 1$ (" $c^- = c - 1$ "); and CS. For each set of system parameters, we evaluate the Markov chains induced by the four policies (in the case of the optimal policy, via value iteration) to calculate long-run average revenues.

In our numerical tests, we fix the expected rental duration of class-1 rentals at $1/\mu_1 = 1$, and we run sets of tests in which we systematically vary the offered load, $\theta = (\rho_1 + \rho_2)/c$, as well as the relative processing rate of class-2 customers, μ_2 . Within each test set, θ and μ_2 also remain fixed, and we run $(10 \times \theta + 1)$ experiments in which we systematically vary λ_1 and λ_2 .

More specifically, in each test set we begin with $(10 \times \theta + 1)$ equally spaced λ_1 s—from $\lambda_1 = 0$ to $\lambda_1 = \mu_1 c$ —and then choose λ_2 in each case so that $\lambda_1/\mu_1 c + \lambda_2/\mu_2 c = \theta$. We then modify the endpoints—where either λ_1 or λ_2 equals zero—so that the arrival rate that would be zero actually equals 0.01. For example, for the set in which $\theta = 1$ and $\mu_2/\mu_1 = 1$, there are $10 \times \theta + 1 = 11$ test points, and their (λ_1, λ_2) values are $\{(0.01, 9.99), (1, 9), (2, 8), (3, 7), (4, 6), (5, 5), (6, 4), (7, 3), (8, 2), (9, 1), (9.99, 0.01)\}$.

Table 1 shows results for the 21 sets of experiments. In each experiment within a set we record average penalty-adjusted revenue per period using the optimal policy (R^*), as well as that obtained from the FAT and CS policies (R^{FAT} and R^{CS}). For each experiment, we then calculate the percentage revenue lost when using the heuristic controls ($(1 - R^{\text{FAT}}/R^*) \times 100\%$ and $(1 - R^{\text{CS}}/R^*) \times 100\%$). Finally, within each cell of Table 1 we report two statistics that summarize the results across all $10 \times \theta + 1$ experiments: the average of the percentage shortfalls, as well as the maximum shortfall recorded over all cases (in parentheses).

Table 1's results show that all three policies perform well at low offered loads. For $\theta \leq 1$, none of the three policies controls the inflow of class-2 requests, and all three perform consistently close to optimality. It is also worth noting that, in these examples, the CS policy is consistently optimal at $\theta = 0.5$. While the sufficient c_i^* s of Theorem 3 can be very large—in the thousands in many of these examples—the offered loads at which the CS policy is actually optimal appear to be much less extreme.

As θ climbs above 1, the three policies diverge, and the FAT heuristics outperform CS. At $\theta = 2$ —when the offered load is twice that of the system's capacity—the FAT with $c^- = c - 1$ still performs quite well, with a worst optimality gap of less than 6% and an average gap in each table cell that is consistently below 1%. Here, the performance of FAT with $c^- = c$ is noticeably worse, with the maximum gap of 16.1% and an average gap ranging from 2.7% to 4.7%. The CS policy's worst-case performance is also 16.1% below optimal, and its average performance in each cell trails that of FAT with $c^- = c$, falling 7.0% to 9.6% below optimality.

At very high θ s, the FAT heuristic with $c^- = c - 1$ consistently outperforms the other heuristics. For example,

Table 1. Numerical results. Average and maximum (in parentheses) percent shortfall from optimal revenue for $10\theta + 1$ test cases.

θ	CS policy			FAT with $c^- = c - 1$			FAT with $c^- = c$		
	$\mu_2/\mu_1 = 0.1$	$\mu_2/\mu_1 = 1$	$\mu_2/\mu_1 = 10$	$\mu_2/\mu_1 = 0.1$	$\mu_2/\mu_1 = 1$	$\mu_2/\mu_1 = 10$	$\mu_2/\mu_1 = 0.1$	$\mu_2/\mu_1 = 1$	$\mu_2/\mu_1 = 10$
0.5	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)
0.8	0.0 (0.0)	0.2 (0.4)	1.5 (2.8)	0.0 (0.0)	0.2 (0.4)	1.5 (2.8)	0.0 (0.0)	0.2 (0.4)	1.5 (2.8)
1.0	0.3 (1.8)	0.9 (1.9)	1.0 (1.9)	0.3 (1.8)	0.9 (1.9)	1.0 (1.9)	0.3 (1.8)	0.9 (1.9)	1.0 (1.9)
1.5	3.1 (6.2)	4.1 (6.5)	6.7 (11.7)	1.2 (4.0)	0.7 (3.2)	0.2 (2.8)	1.9 (5.8)	2.8 (7.3)	4.3 (11.6)
2.0	7.0 (13.1)	6.3 (12.7)	9.6 (16.1)	1.5 (5.9)	0.6 (2.7)	0.3 (2.3)	2.7 (10.9)	3.3 (12.5)	4.7 (16.1)
2.5	9.2 (16.9)	9.6 (16.5)	11.6 (19.6)	1.6 (7.9)	0.7 (4.6)	0.2 (1.9)	3.1 (14.6)	3.5 (16.3)	4.6 (19.6)
3.0	10.9 (20.5)	11.3 (20.1)	12.3 (22.3)	1.6 (9.4)	0.7 (5.3)	0.2 (1.2)	3.5 (17.4)	3.6 (19.1)	4.2 (22.3)

Notes. Penalty-adjusted service fees are $\hat{a}_1 = 10$ and $\hat{a}_2 = 5$ for class-1 and class-2 customers, respectively. Service rate for class-1 customers is $\mu_1 = 1$. Rental fleet size is $c = 10$.

when $\theta = 3$, the average revenue generated by the FAT policy with $c^- = c - 1$ ranged from 0.2% to 1.9% below optimal, and the worst-case examples of each of the $10 \times \theta + 1$ test sets ranged from 1.2% to 9.4% below optimal. In contrast, average and worst-case performance of the FAT with $c^- = c$ and CS policies were three to four times worse.

Thus, as the offered load increases, the performance of all three heuristics deteriorates with respect to optimality. In general, the heuristics are exercising insufficient control of class-2 customers. The relatively strong performance of the FAT heuristic with “ c^- ” set to $c - 1$ reflects the benefit of reserving the last unit of rental capacity for “preferred” class-1 customers when the traffic intensity is high.

Figure 3 provides additional detail on how the setting of c^- affects the performance of the FAT heuristic. In the figure, rental capacity is $c = 10$, service rates are $\mu_1 = 1.0$ and $\mu_2 = 0.1$, and penalty-adjusted revenues are $\hat{a}_1 = 10$ and $\hat{a}_2 = 5$. The aggregate offered load is fixed at $\theta = (\rho_1 + \rho_2)/c = 2$, and the x -axis of the figure’s parametric analysis tracks the fraction of the offered load due to

class-1 customers as it is systematically increased from 0% to 100% of the total: from $\rho_1/c = 0$ to $\rho_1/c = 2$. The y -axis reports the two FAT policies’ resulting percentage shortfall from long-run average optimal revenue.

As Figure 3 indicates, whenever $\rho_1/c \geq 1.0$, the two policies are identical—with the same threshold, $k_{\text{FAT}} \leq c - 1$, and the same long-run average revenues. When $\rho_1/c < 1.0$, however, the two heuristics’ recommendations differ— $k_{\text{FAT}} = c - 1$ versus $k_{\text{FAT}} = c$ —and average revenues differ as well. For moderate ρ_1 s, the “ $c^- = c - 1$ ” policy outperforms the “ $c^- = c$ ” one, and for $\rho_1 \ll c$, the reverse is true.

It is worth noting that numerical experiments using other θ s yield plots whose gross features are directly analogous to those of Figure 3. Larger values of θ lead to more extreme performance differences between the $c^- = c - 1$ and $c^- = c$ variants of the FAT at moderate to very low values of ρ_1 .

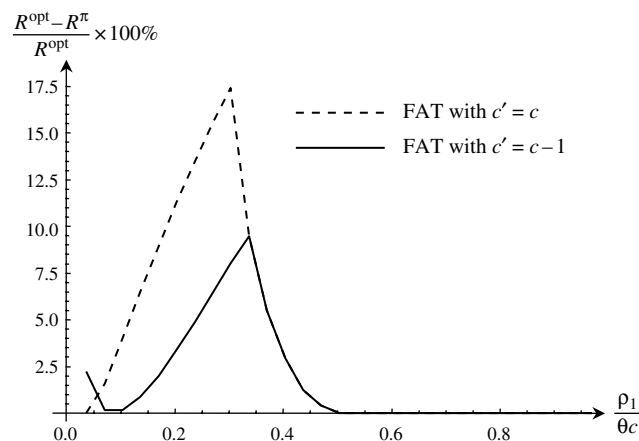
4. The Effect of Capacity Allocation on Optimal Fleet Size

The allocation policies investigated in §§2 and 3 are tactical controls intended to address instances in which the number of rental units available falls short of the anticipated near-term demand. The total fleet size c clearly affects the nature of the control. In particular, Theorem 3 shows that, given ample capacity, the optimal control is to give free access to all customers.

It is also natural to ask the converse question. How does the use of tactical control affect the fleet size the rental company should use? When is the optimal fleet size large enough so that, as in Theorem 3, CS is (nearly) optimal? More generally, given the ability to change fleet size, what is the economic value to a firm of exercising tactical controls? In this section, we address both of these questions.

In fact, the effect of capacity allocation on optimal fleet size is not immediately clear. One might argue that, given any fixed fleet size, optimal rationing increases revenue per unit of time. This revenue increase, in turn, allows the firm to more profitably sustain higher overall capacity levels. Alternatively, one might argue that rationing reduces the

Figure 3. Performance of alternative FAT heuristics with c^- interpreted as c (dashed line) and as $c - 1$ (solid line) (system has $c = 10$, $\theta = (\rho_1 + \rho_2)/c = 2$, $\mu_1 = 1$, $\mu_2 = 0.1$, $\hat{a}_1 = 10$, and $\hat{a}_2 = 5$).



aggregate arrival rate to the rental fleet and that, in turn, fewer units of capacity are required to process the arrivals that are actually served.

We can provide some insight into these trade-offs by directly comparing the optimal fleet size under active allocation policies to that under CS, which passively allows all customers access to rental capacity whenever it is available. We formulate the problem of finding the optimal fleet size as

$$\Pi(\Delta) = \max_c (R(c, \Delta(c)) - hc), \tag{33}$$

where $R(c, \Delta(c))$ is the average revenue per period when operating c units under allocation policy $\Delta(c)$, and the capacity cost of $\$h$ per unit per period is fixed for all c . Note that, given a fixed offered load, $\rho_1 + \rho_2$, the allocation policy, $\Delta(c)$, may vary with c .

We then compare the maximizer of (33) under two regimes. In one we use $\Delta(c) = \text{CS}(c)$, the CS policy, for all c . In the other $\Delta(c) = \Delta^*(c)$, which we define as any family of allocation policies for which the following conditions hold:

1. For any fixed c , $R(c, \Delta^*(c)) \geq R(c, \text{CS}(c))$.
2. There exists a $\tilde{c} < \infty$ such that for all $c \geq \tilde{c}$, $R(c, \Delta^*(c)) = R(c, \text{CS}(c))$.
3. $R(c, \Delta^*(c)) - R(c - 1, \Delta^*(c - 1)) \leq R(c - 1, \Delta^*(c - 1)) - R(c - 2, \Delta^*(c - 2))$.

Condition 1 states that, for any c , $\Delta^*(c)$ performs at least as well as CS.

Condition 2 states that there exists a finite fleet size above which CS performs as well as $\Delta^*(c)$. Note that Theorem 3 demonstrates that such a \tilde{c} exists in the context of the discounted problem.

Condition 3 requires that average revenues per period under Δ^* are concave in c . Theorem 6 proves that this type of concavity exists for the FAT policy in the context of the discounted fluid model, and the result also suggests that the condition (roughly) holds for AT policies more generally. Similarly, although we have not been able to prove that the condition holds for the optimal policy, it has consistently been present in the numerical tests we have run.

Without loss of generality, we assume that $\hat{a}_1 \geq \hat{a}_2$, and we define

$$\begin{aligned} h_{\min}^* &= R(\tilde{c}, \Delta^*(\tilde{c})) - R(\tilde{c} - 1, \Delta^*(\tilde{c} - 1)), \\ h_{\min}^{\text{CS}} &= R(\tilde{c}, \text{CS}(\tilde{c})) - R(\tilde{c} - 1, \text{CS}(\tilde{c} - 1)), \\ h_{\max}^* &= R(1, \Delta^*(1)) - R(0, \Delta^*(0)) \\ &\geq \max\left(\frac{\hat{a}_1 \rho_1}{1 + \rho_1}, \frac{\hat{a}_1 \rho_1 + \hat{a}_2 \rho_2}{1 + \rho_1 + \rho_2}\right), \quad \text{and} \\ h_{\max}^{\text{CS}} &= R(1, \text{CS}(1)) - R(0, \text{CS}(0)) = \frac{\hat{a}_1 \rho_1 + \hat{a}_2 \rho_2}{1 + \rho_1 + \rho_2}. \end{aligned} \tag{34}$$

Observe that h_{\min}^* and h_{\min}^{CS} are the marginal values of adding the last piece of equipment, as it becomes optimal to take all arrivals, first-come first-served. Similarly, h_{\max}^* and h_{\max}^{CS} are the marginal values of the first piece of equip-

ment under the two schemes. It is not difficult to see that $h_{\min}^* \leq h_{\min}^{\text{CS}} \leq h_{\max}^{\text{CS}} \leq h_{\max}^*$.

The following result uses these relationships to parameterize how the fleet size under capacity allocation policies differs from that under CS:

THEOREM 7. *Let $c^*(h)$ and $c^{\text{CS}}(h)$ be the maximizers of (33) under Δ^* and CS.*

- (a) *If $h < h_{\min}^*$, then $c^*(h) = c^{\text{CS}}(h) \geq \tilde{c}$.*
- (b) *If $h \in [h_{\min}^*, h_{\min}^{\text{CS}}]$, then $c^*(h) \leq c^{\text{CS}}(h)$.*
- (c) *If $h \in [h_{\min}^{\text{CS}}, h_{\max}^{\text{CS}}]$, then $c^*(h)$ may be smaller, equal to, or larger than $c^{\text{CS}}(h)$.*
- (d) *If $h \in [h_{\max}^{\text{CS}}, h_{\max}^*]$, then $c^*(h) \geq c^{\text{CS}}(h)$.*
- (e) *If $h > h_{\max}^*$, then $c^*(h) = c^{\text{CS}}(h) = 0$.*

We note that the results of Theorem 7 can be extended to the multiclass case, as well as to multiperiod capacity-sizing models with more complex cost structures. We briefly discuss the latter in the discussion at the end of the paper.

Thus, the optimal fleet size using capacity rationing may be either higher or lower than that under the CS policy. The theorem shows that the relationship between the two depends fundamentally on the unit cost of capacity.

Parts (a) and (e) of Theorem 7 show that optimal capacity levels for CS and Δ^* coincide for very high and very low values of holding costs. If holding costs are extremely high, the expected revenues cannot justify the acquisition of even a single unit of capacity, even under rationing. On the other hand, if the holding costs are extremely low, then Theorem 3 implies that the optimal rationing policy is CS, and in this case the profit-maximizing capacity levels of the two policies again coincide.

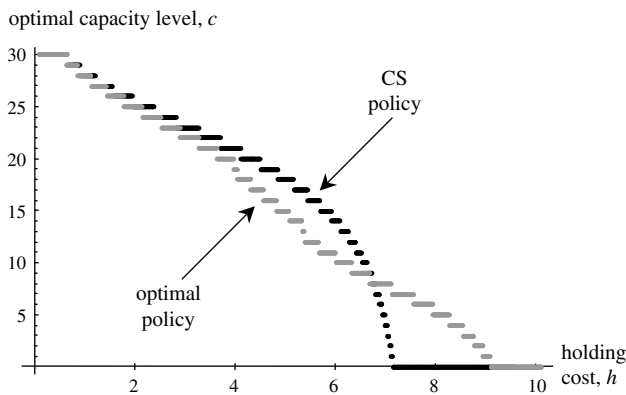
Parts (b) and (d) show ranges for which the $c^*(h)$ unambiguously dominates and is dominated by $c^{\text{CS}}(h)$. Part (b) shows that, for low values of h , the lower marginal value to the rationing policy of adding the “last” unit of capacity (before CS becomes optimal) drives $c^*(h)$ below $c^{\text{CS}}(h)$. Part (d) shows that, for high values of h , the benefit of being able to reject lower-revenue customers allows $c^*(h)$ to climb above $c^{\text{CS}}(h)$.

Finally, part (c) defines a set of intermediate values of h for which $c^*(h)$ can be higher, the same as, or lower than $c^{\text{CS}}(h)$. The ordering of the relationships reflects the proximity of h to the boundaries, h_{\min}^{CS} and h_{\max}^{CS} .

While the relationships described in the theorem are not strict inequalities, it is not difficult to develop examples in which $c^*(h)$ differs from $c^{\text{CS}}(h)$. Figure 4 illustrates an example in which the $\Delta^*(c)$ used for each c is the optimal policy for that c . Given the problem parameters $\tilde{c} = 30$, $h_{\max}^* = 9.09$, $h_{\max}^{\text{CS}} = 7.14$, $h_{\min}^* = 0.65$, and $h_{\min}^{\text{CS}} = 0.63$. Between h_{\max}^{CS} and h_{\min}^* , optimal fleet sizes for the two policies are equal at $h = 6.5$.

We now turn to the economic benefit of capacity rationing. Table 2 presents a set of 25 numerical experiments that compare the performance of the optimal (OPT) and complete sharing (CS) policy. For each example, the table

Figure 4. Optimal capacity size as a function of the holding cost under the optimal and CS policies (fixed problem parameters: $\hat{a}_1 = 10$, $\hat{a}_2 = 5$, $\lambda_1 = \lambda_2 = 10$, and $\mu_1 = \mu_2 = 1$).



Note. Capacity cost per unit of time, h , is systematically varied.

reports the optimal fleet size, profit per period, and (percent) profit margin for both policies. In all of the experiments, the aggregate arrival rate ($\lambda_1 + \lambda_2$), service rates (μ_1 and μ_2), and penalty-adjusted revenues (\hat{a}_1 and \hat{a}_2) remain fixed. Then, the fraction of the offered load due to class-1 customers ($\lambda_1/(\lambda_1 + \lambda_2)$) and the holding cost per unit of time per unit of capacity (h) are systematically varied. The table's results reflect three phenomena that are worth noting.

The first is the effect of increased holding costs on fleet sizes, displayed in Figure 4. At lower relative holding costs, capacity rationing reduces the optimal fleet size relative to that for CS in a company. As one looks down each pair of columns, however, one sees that rationing allows a company to maintain larger capacity than would be optimal

under CS. Interestingly, in looking across each row, one sees that for very small and very large fractions of class-1 customers, optimal capacities from the two policies are the same. In the former case, this is due to the optimality of CS; the policies themselves coincide.

In the latter case, however, the optimal policy reserves capacity for class-1 customers. While complete sharing is suboptimal, the blocking of class-1 customers due to class-2 admissions is a rare enough event that it does not significantly affect optimal fleet size (or, for that matter, profits).

The second is the fact that, when capacity costs are relatively low, CS appears to be fairly robust with respect to average profit per unit of time. In particular, when $h/\hat{a}_2 = 0.5$, the profit advantage derived from capacity rationing is minimal, less than 1%, and when $h/\hat{a}_2 = 0.7$, the advantage is no more than 3%. Here, increased costs, due to additional capacity, are made up for by increased revenues, due to additional class-2 traffic. Rather, it is when capacity costs are high—as high as or higher than class-2 penalty-adjusted revenues—that the profit increase due to restrictions on class-2 access becomes significant.

The last effect is that, conversely, capacity rationing provides for a more consistently significant increase in profit margins over CS. For example, even when capacity costs are half that of penalty-adjusted class-2 revenues, margins may increase by as much as 9%. As capacity costs approach and exceed class-2 fees, the benefit that follows the ability to limit class-2 customers increases far more sharply.

Of dollar profit and profit margin, which is more indicative of the value of rationing to the rental company? For capital-constrained companies, we would argue it is the latter. Indeed, in many rental businesses, capacity is a signif-

Table 2. Numerical results. Optimal fleet sizes, profit per unit time, and profit margins for the optimal (OPT) and complete-sharing (CS) policies.

	h/\hat{a}_2	$\lambda_1/(\lambda_1 + \lambda_2) = 0.1$		$\lambda_1/(\lambda_1 + \lambda_2) = 0.3$		$\lambda_1/(\lambda_1 + \lambda_2) = 0.5$		$\lambda_1/(\lambda_1 + \lambda_2) = 0.7$		$\lambda_1/(\lambda_1 + \lambda_2) = 0.9$	
		OPT	CS	OPT	CS	OPT	CS	OPT	CS	OPT	CS
Fleet sizes	0.5	11	11	12	13	12	13	13	13	14	14
	0.7	9	9	9	10	10	11	11	12	12	12
	0.9	5	5	6	8	8	9	9	10	11	11
	1.1	0	0	3	4	6	7	8	9	9	10
	1.3	0	0	2	0	4	3	6	6	8	8
Profits	0.5	18.52	18.52	27.22	27.02	36.40	36.17	45.66	45.33	54.71	54.60
	0.7	8.47	8.47	16.27	16.05	24.98	24.26	33.46	32.82	41.94	41.62
	0.9	1.48	1.48	8.31	7.00	15.59	14.01	23.01	21.76	30.50	29.99
	1.1	—	—	3.46	0.97	8.71	5.82	14.50	12.28	20.50	19.61
	1.3	—	—	1.12	—	4.15	0.60	7.90	4.82	12.03	10.86
Profit margin	0.5	67.3%	67.3%	90.7%	83.1%	121.3%	111.3%	140.5%	139.5%	156.3%	156.0%
	0.7	26.9%	26.9%	51.7%	45.9%	71.4%	63.0%	86.9%	78.1%	99.9%	99.1%
	0.9	6.6%	6.6%	30.8%	19.4%	43.3%	34.6%	56.8%	48.4%	61.6%	60.6%
	1.1	—	—	21.0%	4.4%	26.4%	15.1%	33.0%	24.8%	41.4%	35.7%
	1.3	—	—	8.6%	—	16.0%	3.1%	20.3%	12.4%	23.1%	20.9%

Notes. In all test cases, the following parameters are fixed: $\hat{a}_1 = 10$, $\hat{a}_2 = 5$, $\mu_1 = \mu_2 = 1$, and $\lambda_1 + \lambda_2 = 10$. The relative value of the holding cost, h/\hat{a}_2 , and the fraction of demand due to the preferred class, $\lambda_1/(\lambda_1 + \lambda_2)$, are systematically varied.

icant capital investment, and measures, such as *return on assets*, that track the (dollar) efficiency of asset utilization become critical measures of performance for managers and for investors. In our numerical experiments, h —the cost per unit of capacity per unit of time—reflects interest expenses of capacity investment (as well as maintenance expenses). In dividing profit by $h \cdot c$ (or, equivalently for us, by c), profit margin accounts for this investment in capacity.

Thus, absolute profits lost due to lack of control may not be large. Nevertheless, the ability to ration capacity and, in turn, to adjust the size of a rental fleet can, at the same time, lead to a significant improvement in the economic utilization of the assets employed. The CS policy maximizes physical—rather than economic—utilization of assets. When CS is optimal, the two notions naturally coincide. When it is not optimal, however, it leads to lower economic productivity.

5. Discussion

Our formulation of the rental capacity allocation problem captures some essential features that the more traditional yield-management literature does not address. In it, we explicitly represent the fact that customers arrive at random, use pieces of equipment for rental periods of uncertain duration, and then return the equipment to be used again.

Using dynamic programming techniques, we are able to characterize “switching-curve” policies as being optimal. We also demonstrate that there are two sets of conditions under which a customer class should be labeled a “VIP” and have unrestricted access to the available service capacity: one in which there is ample excess capacity and another in which the penalty-adjusted revenue and service-rate parameters are favorable. In particular, we find that customers may be assigned the VIP tag even when their rental fees are lower than those of the other class.

When applied to both customer classes, the sufficient conditions for VIP status become conditions in which “complete-sharing” (CS) policies are optimal. These policies are of interest because service companies often use equipment utilization as a criterion for measuring system performance and may be reluctant to turn away customers. Theorem 4 implies that the goals of maximizing utilization and of maximizing revenues are properly aligned, even in the peak season, if the penalty-adjusted rental fees and service rates of the different customer classes are similar.

We also analyze a “fluid aggregate threshold” (FAT) policy that is based on a fluid approximation of the original policy. Our numerical tests show that the performance of the FAT heuristic is close to that of the optimal admission policies over a broad range of operating regimes. In addition to providing a simple and effective capacity allocation policy, the fluid model results in revenue that is a concave function of the rental fleet size. This concavity

is essential for the analysis of related capacity-sizing decisions and for an understanding of how capacity allocation schemes affect them.

We then demonstrate that, given this concavity property, the optimal fleet size using capacity rationing may be either higher or lower than that under the CS policy. As capacity costs grow, the optimal fleet size under rationing grows relative to that under CS.

Finally, we show that appropriate adjustment of the fleet size under CS may produce nearly optimal profits. Even in this case, however, the economic productivity of assets can suffer significantly. When CS is not optimal, its maximization of physical utilization leads to economic underutilization of resources.

Thus, the formulation and results represent a promising step in furthering the understanding of the management of rental systems. Of course, more work remains to be done. There are several aspects of the allocation problem itself that merit additional analysis.

First, as we noted in §2, rental companies may have prior estimates of the expected duration of the rental period, and this information would be of value when deciding whether to admit a customer to the system. At the same time, the use of this information will also significantly complicate the analysis. For example, it will likely require expanding the state space of the system from numbers of pieces of equipment in use to estimates of the duration of the remaining rental period for every piece of equipment.

Similarly, our description of rental dynamics does not include the treatment of reservation systems, which may provide additional information about rental demand. Again, the inclusion of reservation systems should help to improve system performance, and it will also add an additional layer of complexity to the analysis.

One may also consider price, in addition to capacity allocation, as a mechanism for control. In particular, an interesting case exists in which one class of customers represents national accounts, whose prices are fixed by long-term contracts, while the other represents “rack-rate” customers, for whom price may be used as a short-term control. Then, the rental company may use capacity allocation to maintain service levels for national-account clients at the same time it uses prices to maximize profits from rack-rate customers.

The relationship between capacity allocation and fleet sizing can also be explored over a longer time horizon. For example, consider a longer-term, discrete-time problem in which each period represents a season. At the start of each season, rental capacity is adjusted by buying and selling units, and during the season tactical controls, such as the ones developed in this paper, are used to manage short-term capacity shortages. Then, if the season’s expected revenues are concave in the fleet size, it is not difficult to show that the optimal fleet-sizing policy is a “buy-up-to/sell-down-to” policy that is an analogue of “order-up-to” policies in the inventory literature (Heyman and Sobel 1984).

Finally, we recall that our formulation uses lump-sum penalty costs to capture the long-term cost of denying access to customers. An alternative would be to impose service-level constraints on the blocking probabilities of arrivals. While we believe that the current policies should be “nearly” feasible, particularly for large systems, a thorough analysis of the relationship between the two formulations would be of broad interest.

Appendix

The appendix is available in the online companion on the *Operations Research* website at <http://or.pubs.informs.org/Pages/collect.html>.

Acknowledgments

The authors acknowledge the referees and especially the associate editor, whose comments helped to improve the paper substantially. This research was supported by NSF grants SBR-9733739 and 540075-8001, and by the Fishman-Davidson Center for Service and Operations Management, The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania.

References

- Alstrup, J., S. Boas, O. B. G. Madsen, R. V. V. Vidal. 1986. Booking policy for flights with two types of passengers. *Eur. J. Oper. Res.* **27** 274–288.
- Altman, E., G. Koole, T. Jiménez. 2001. On optimal call admission control in a resource-sharing system. *IEEE Trans. Comm.* **49** 1659–1668.
- Belobaba, P. P. 1989. Application of a probabilistic decision model to airline seat inventory control. *Oper. Res.* **37** 183–197.
- Bitran, G. R., S. M. Gilbert. 1996. Managing hotel reservations with uncertain arrivals. *Oper. Res.* **44** 35–49.
- Heyman, D. P., M. J. Sobel. 1984. *Stochastic Models in Operations Research*, Vol. 2. McGraw-Hill, New York, 391–407.
- Ladany, S. 1977. Bayesian dynamic operating rules for optimal hotel reservation. *Z. Oper. Res.* **21** B165–B176.
- Liberman, V., U. Yechiali. 1978. On the hotel overbooking problem: An inventory system with stochastic cancellations. *Management Sci.* **24** 1117–1126.
- Lippman, S. 1975. Applying a new device in the optimization of exponential queueing systems. *Oper. Res.* **23** 687–710.
- Littlewood, K. 1972. Forecasting and control of passenger bookings. *AGIFORS Sympos. Proc.* **12** 95–117.
- Miller, B. 1969. A queueing reward system with several customer classes. *Management Sci.* **16** 234–245.
- Örmeci, L., A. Burnetas, J. van der Wal. 2001. Admission policies to a two-class loss system. *Stochastic Models* **17** 513–539.
- Porteus, E. 1982. Conditions for characterizing the structure of optimal strategies in infinite-horizon dynamic programs. *J. Optim. Theory Appl.* **36** 419–432.
- Puterman, M. L. 1994. *Discrete Stochastic Dynamic Programming*. John Wiley and Sons, New York.
- Ross, K. W., D. H. K. Tsang. 1989. The stochastic knapsack problem. *IEEE Trans. Comm.* **37** 740–747.
- Rothstein, M. 1974. Hotel overbooking as a Markovian sequential decision process. *Decision Sci.* **5** 389–404.
- Savin, S. V. 2001. Managing capital intensive rental businesses. Ph.D. thesis, The Wharton School, University of Pennsylvania, Philadelphia, PA.
- Serfozo, R. F. 1979. An equivalence between continuous and discrete time Markov decision processes. *Oper. Res.* **27** 616–620.
- Weatherford, L. R. 1995. Length of stay heuristics: Do they really make a difference? *Cornell Hotel Restaurant Admin. Quart.* **36** 70–79.
- Williams, F. E. 1977. Decision theory and the innkeeper: An approach for setting hotel reservation policy. *Interfaces* **7** 18–30.
- Williamson, E. L. 1992. Airline network seat inventory control: Methodologies and revenue impact. Ph.D. dissertation, MIT Flight Transportation Laboratory Report R92-3, Massachusetts Institute of Technology, Cambridge, MA.