OPERATIONS RESEARCH
DOI 10.1287/opre.1090.0745ec

e-companion

ONLY AVAILABLE IN ELECTRONIC FORM

Electronic Companion—"Capacity Rationing in Stochastic Rental Systems with Advance Demand Information" by Felix Papier and Ulrich W. Thonemann, *Operations Research*, DOI 10.1287/opre.1090.0745.

Online Appendix

A Table of Symbols

$\varnothing$	Empty set
$\mathbf{0}$	Null matrix
$\mathbf{1}(x)$	Indicator function of expression x
$\alpha_k(x)$	Auxiliary function
a_j, a_j^*	(Optimal) admission decision for customer j
A_j	Arrival time of customer j
$\{(A_k, CL_k)\}$	Arrival process with arrival times A_k and class indicators CL_k
$\mathbf{A}$, $\mathbf{C}$, $\widetilde{\mathbf{P}}$	Auxiliary matrices
c	Fleet size
CL_j	Class indicator of jth customer
δ_i	Normalizing constant of row i
$\Delta\Pi$	Delta in profit between two states or systems
e, A, B, C, D	Event types
$E[X]$	Expectation of random variable X
$\overline{F}_S(t)$	Complementary cumulative distribution function of rental time S
Γ, Γ^*	(Optimal) admission policy
$\mathbf{I}_i$	Identity matrix with dimension i
λ_I, λ_{II}	Arrival rates of ADI and non-ADI customers
MR	Maximum number of reservations
n	Number of subintervals in demand leadtime, number of ADI classes in Section F
v	Share of mean ADI demand in total demand
N, N^-, N^+	Number of busy cars, at events either directly before (N^-) or after (N^+) the corresponding event
(N, R)	State space characterization
p	Penalty cost

$p_{ij}(a)$, $p_{ij}^k(a)$	Transition probability from state i to state j under decision a
$p_{ij,n}$, p_{in}	Probability of event i to occur in subinterval n for customer j (index j is dropped for conciseness)
$p_{ij,\tau+}$, $p_{i,\tau+}$	Probability of event i to occur after demand leadtime for customer j (index j is dropped for conciseness)
$\boldsymbol{\pi}_n^-$, $\boldsymbol{\pi}_n^+$	State distribution right before ($\boldsymbol{\pi}_n^-$) or after ($\boldsymbol{\pi}_n^+$) the arrival of the nth ADI customer after the current arrival
$P_i(N^-, R)$	Probability of event i, depending on state (N^-, R)
$P_j(Event\ \mathbf{i})$	Probability of event i to occur after arrival of customer j
$\mathbf{P}(t)$	Transition matrix at time t
$\mathbf{P}_{\tau+}$	Event transition matrix after demand leadtime
Π^+	Future profit after corresponding event
$\Pi_k(i)$	Future profit after customer k if in state i
$\mathbf{Q}$, $\mathbf{Q}_m^{multiple}$	Infinitesimal generator matrices
r_I, r_{II}	Revenues from ADI and non-ADI customers
R	Set of reservations
R^2	Determination coefficient
RT	Run-Time
S, $\overline{S}$, $\overline{S}_i$	(Mean) rental time (of class i)
$\widehat{S}$	Mean rental time used for ADI heuristic (Subsection F.3)
$\mathcal{S}$	State space
τ	Demand leadtime
t_k	Reservation time
t, $\widehat{t}$	Arbitrary time epochs
$T(R)$	Threshold level, depending on R
$T_{\cdot\infty}$	First passage time into an absorbing state
U	Matrix of (right) eigenvectors of $\mathbf{Q}$
$V_i(\Gamma)$	Service rate of class i under policy Γ
w	Vector of eigenvalues of $\mathbf{Q}$
X, Y	Rental systems

B SMDP Characterization

The model described in Section 3 is a SMDP with finite action space and infinite and (multi-dimensional) continuous state space. The objective is maximizing average profit (which includes positive and negative rewards) over an infinite horizon.

The action space is $\{0\}$ for event A, $\{0,1\}$ if $N < c$ for events B and D, and $\{-1\}$ if $N > 0$ for event C. The state space $\mathcal{S}$ is defined as $\mathcal{S} = \mathbb{N} \times \mathbb{R}_+^{|R|} \times \{A,B,C,D\}$. The optimal profit (or value) function $\Pi(i)$ is given by

$$\Pi(i) = \max_a \left[C(i,a) + \sum_j p_{ij}(a)\Pi(j) \right] \tag{9}$$

for any initial state $i = (N,R,e) \in \mathcal{S}$ and $j = (N',R',e') \in \mathcal{S}$, where $C(i,a)$ is defined as

$$C(i,a) = \begin{cases} r_I & e = D \wedge a = 1 \\ -p & e = D \wedge a = 0 \\ r_{II} & e = B \wedge a = 1 \\ 0 & else \end{cases}, \tag{10}$$

and $p_{ij}(a)$ are the transition probabilities from i to j. Note that our SMDP can be transformed into an equivalent SMDP with positive rewards only, so all results for models with positive rewards (cf. Puterman (1994)) also apply to our SMDP.

The sojourn (transition) time $\widehat{S}_i$ from state i into the next state is given by cdf

$$F_i(t \mid a) = \begin{cases} \frac{1-e^{-\mu t}}{1-e^{-\mu t_1}} & 0 < t < t_1 \\ 1 & t_1 \leq t \\ 0 & else \end{cases}, \tag{11}$$

where $\mu = \lambda_I + \lambda_{II} + N + a$. If $R = \varnothing$, $F_i(t \mid a)$ reduces to $F_i(t \mid a) = 1 - e^{-\mu t}$.

The transition probabilities $p_{ij}(a)$ are given by

$$p_{ij}(a) = \begin{cases} \frac{\lambda_I}{\mu}\left(1-e^{-\mu t_1}\right) & R' = \{t_k - t \mid t_k \in R\} \cup \{\tau\} \wedge e' = A \\ \frac{\lambda_{II}}{\mu}\left(1-e^{-\mu t_1}\right) & R' = \{t_k - t \mid t_k \in R\} \wedge e' = B \\ \frac{N+a}{\mu}\left(1-e^{-\mu t_1}\right) & R' = \{t_k - t \mid t_k \in R\} \wedge e' = C \\ e^{-\mu t_1} & R' = \{t_k - t \mid t_k \in R\} \setminus \{t_1\} \wedge R' \neq \varnothing \wedge e' = D \\ 0 & else \end{cases}, \tag{12}$$

where generally $N' = N + a$ (all transitions with $N' \neq N + a$ have a probability of zero). Now, we show how the transition probability for event A ($N' = N+a \wedge R' = \{t_k - t \mid t_k \in R\} \cup \{\tau\} \wedge e' = A$) can be derived. The derivation of the other probabilities in Equation (12) is analogous. For having event A (a new ADI arrival is announced and added to the set of ADI, R), the transition times to other events must be greater than that to event A. The transition time to events B and C are exponentially distributed with a joint rate of $\mu' = \lambda_{II} + N + a$. The transition time to event D is fix with t_1 (assume for now that $R \neq \varnothing$). Thus, we have

$$\begin{aligned} P(A) &= P(\widehat{S}_i^A \leq \min(\widehat{S}_i^B, \widehat{S}_i^C, t_1)) \\ &= \int_0^{t_1} \left(1 - e^{-\lambda_I t}\right) \mu' e^{-\mu' t} dt + \left(1 - e^{-\lambda_I t_1}\right) e^{-\mu' t_1} \\ &= \frac{\lambda_I}{\lambda_I + \mu'}\left(1 - e^{-(\lambda_I + \mu')t_1}\right) = \frac{\lambda_I}{\mu}\left(1 - e^{-\mu t_1}\right). \end{aligned} \tag{13}$$

If $R = \varnothing$, the next transition cannot be into event D. The probabilities can be obtained by taking the limits with $t_1 \to \infty$.

C Algorithm for Computation of ADI Event Probabilities

The results of Section 5 allow the computation of the estimates for $P_1(N^-, R)$ and $P_2(N^-, R)$ for each non-ADI arrival. The estimates can then be used for the admission rule of Proposition 2.

Algorithm 1 *Computation of event probabilities for non-ADI arrival in state* (N^-, R)

0. Set $\widetilde{p} = 1$, $P_i(N^-, R) = 0$ *for* $i = 1, 2, 3$.

0b. Set $\left(\boldsymbol{\pi}_0^+\right)_k = 1$ *for* $k = N_j^-$ *and* $\left(\boldsymbol{\pi}_0^+\right)_k = 0$ *for all other* k.

1. Repeat for all reservations in R, use index $n = 1....$

1a. Compute $\boldsymbol{\pi}_n^-$.

1b. Compute p_{1n}, p_{2n}, and p_{3n} with Equations (5), (6), and (7).

1c. Compute $\boldsymbol{\pi}_n^+$.

1d. Update $P_i(N^-, R) = P_i(N^-, R) + \widetilde{p} \cdot p_{in}$ for $i = 1, 2, 3$.

1e. Update $\widetilde{p} = \widetilde{p} \cdot (1 - p_{1n} - p_{2n} - p_{3n})$.

2. Compute $\boldsymbol{\pi}_n^-$ and subinterval length $A_0 + \tau - A_{|R|}$.

3. Compute p_{2n} and p_{3n} with Equations (5) and (7) and set $p_{1n} = 0$.

4. Copy $\boldsymbol{\pi}_n^-$ to $\boldsymbol{\pi}_n^+$.

5. Update $P_i(N^-, R) = P_i(N^-, R) + \widetilde{p} \cdot p_{in}$ for $i = 1, 2, 3$.

6. Update $\widetilde{p} = \widetilde{p} \cdot (1 - p_{1n} - p_{2n} - p_{3n})$.

7. Update $P_i(N^-, R) = P_i(N^-, R) + \widetilde{p} \cdot p_{i,\tau+}$ for $i = 1, 2, 3$.

D Proofs

Proof of Proposition 1. Let $\Pi_j^+(k)$ denote the optimal future profit of a rental system with currently k busy cars and reservation information R right after arrival j. *Note that $\Pi_j^+(k)$ is, in this proof, a random variable.* We show that

$$\Pi_j^+(k+2) - \Pi_j^+(k+1) \leq_{st} \Pi_j^+(k+1) - \Pi_j^+(k),\ \forall k = 0...c-2, \tag{14}$$

from which follows that $E\left(\Pi_j^+(k)\right)$ is concave in $k = 0...c-2$. With $E\left(\Pi_j^+(k)\right)$ being concave in k, $E\left(\Delta\Pi_j^+(k)\right)$ is decreasing in k. Thus, if $E\left(\Delta\Pi_j^+(k)\right) < -r_j$ for some k, it would be optimal to reject customer k and it would also be optimal for all $k' > k$ (because $E\left(\Delta\Pi_j^+(k)\right)$ is decreasing). The smallest k for which we would reject the customer is the threshold level. This level may not be static but depends on the available ADI, R. Since the demand processes are stationary, it is obvious that if the system has the same ADI at two different time epochs, the threshold levels at the two time epochs will be identical. Thus, the threshold level can *only* depend on R.

Given a certain action a_{j+1}, we get

$$\Pi_j^+(k|a_{j+1}) = r_{j+1}a_{j+1} + \Pi_{j+1}^+(k - D_j + a_{j+1}) \tag{15}$$

with D_j as the departures in time period $[A_j, A_{j+1}]$. Note that we write r_j for the reward of accepting customer j. This is $r_I + p$ for an ADI customer in the two class model, because accepting a customer avoids the penalty p.

Now, we consider four systems $\mathcal{S}1$, $\mathcal{S}2$, $\mathcal{S}3$, and $\mathcal{S}4$ with identical arrival processes (and, thus, all systems have the same ADI). The systems have the respective system states $k+2$, k, $k+1$, $k+1$. The departure probability of the first additional customer in $\mathcal{S}1$ and the additional one in $\mathcal{S}3$ are identical. Therefore, we couple these two events and denote the corresponding random variable of the customer still being in the system with N' ($N' = 0$ corresponds to a departure of this customer, $N' = 1$ means that the customer is still in the system). We also couple the departure event of the second additional customer in $\mathcal{S}1$ and the departure event of the additional customer in $\mathcal{S}4$. We call the corresponding random variable N''. Let systems $\mathcal{S}1$ and $\mathcal{S}2$ behave optimally, while systems $\mathcal{S}3$ and $\mathcal{S}4$ take their actions based on the system states and the prior departure events.

Now assume we are at the last arrival j_{max}. In this period, the terminal value is identical for all system states (by assumption), i.e., we have

$$\Pi^+_{j_{\max}}(k+2) + \Pi^+_{j_{\max}}(k) - \Pi^+_{j_{\max}}(k+1) - \Pi^+_{j_{\max}}(k+1) = 0,\ \forall k = 0...c-2 \tag{16}$$

and Ineq. (14) holds for $j = j_{max}$.

Next, we proceed by backward induction on arrival j. Assume that Ineq. (14) holds for $j+1$ (as we have shown for j_{max}). A consequence from this is that the optimal decisions of $\mathcal{S}1$ and $\mathcal{S}2$ are related: $a_{j+1}(\mathcal{S}1) \le a_{j+1}(\mathcal{S}2)$. We assume an arbitrary system state k and distinguish between two cases:

Case 1. $a_{j+1}(\mathcal{S}1) = a_{j+1}(\mathcal{S}2) = a$. In this case, let systems $\mathcal{S}3$ and $\mathcal{S}4$ take the same actions at $j+1$ as $\mathcal{S}1$ and $\mathcal{S}2$. Then, we have for arrival j,

$$\begin{aligned}
& \Pi_j^+(k+2) + \Pi_j^+(k) - \Pi_j^+(k+1|a_{j+1}(\mathcal{S}3)=a) - \Pi_j^+(k+1|a_{j+1}(\mathcal{S}4)=a) \\
=_{st} \quad & r_{j+1}a + \Pi_{j+1}^+(k+N'+N''-D_j+a) + r_{j+1}a + \Pi_{j+1}^+(k-D_j+a) \\
& -r_{j+1}a - \Pi_{j+1}^+(k+N'-D_j+a) - r_{j+1}a - \Pi_{j+1}^+(k+N''-D_j+a) \\
=_{st} \quad & \Pi_{j+1}^+(k-D_j+a+N'+N'') + \Pi_{j+1}^+(k-D_j+a) \\
& -\Pi_{j+1}^+(k-D_j+a+N') - \Pi_{j+1}^+(k-D_j+a+N'') \\
\leq_{st} \quad & 0. \qquad (17)
\end{aligned}$$

The last inequality holds, because it reduces to 0 if $(N'=0) \vee (N''=0)$ and it is $\leq_{st}$ 0 by induction assumption if $(N'=1) \wedge (N''=1)$.

Case 2. $(a_{j+1}(\mathcal{S}1)=0) \wedge (a_{j+1}(\mathcal{S}2)=1)$. Within case 2, we distinguish two subcases:

Subcase 2a. $(N'=0) \wedge (N''=1)$. In this case, let $\mathcal{S}3$ take the same action at $j+1$ as $\mathcal{S}2$ and $\mathcal{S}4$ the same action as $\mathcal{S}1$. We end up with

$$\begin{aligned}
& \Pi_j^+(k+2) + \Pi_j^+(k) - \Pi_j^+(k+1|a_{j+1}(\mathcal{S}3)=1) - \Pi_j^+(k+1|a_{j+1}(\mathcal{S}4)=0) \\
=_{st} \quad & \Pi_{j+1}^+(k+N'+N''-D_j) + r_{j+1} + \Pi_{j+1}^+(k-D_j+1) \\
& -r_{j+1} - \Pi_{j+1}^+(k+N'-D_j+1) - \Pi_{j+1}^+(k+N''-D_j) \\
=_{st} \quad & \Pi_{j+1}^+(k+1-D_j) + \Pi_{j+1}^+(k-D_j+1) - \Pi_{j+1}^+(k-D_j+1) - \Pi_{j+1}^+(k+1-D_j) \\
=_{st} \quad & 0. \qquad (18)
\end{aligned}$$

Subcase 2b. $\neg((N'=0) \wedge (N''=1))$. In this case, let $\mathcal{S}3$ take the same action at $j+1$ as $\mathcal{S}1$ and $\mathcal{S}4$ the same action as $\mathcal{S}2$. We end up with

$$\begin{aligned}
& \Pi_j^+(k+2) + \Pi_j^+(k) - \Pi_j^+(k+1|a_{j+1}(\mathcal{S}3)=0) - \Pi_j^+(k+1|a_{j+1}(\mathcal{S}4)=1) \\
=_{st} \quad & \Pi_{j+1}^+(k+N'+N''-D_j) + r_{j+1} + \Pi_{j+1}^+(k-D_j+1) \\
& -\Pi_{j+1}^+(k+N'-D_j) - r_{j+1} - \Pi_{j+1}^+(k+N''-D_j+1) \\
=_{st} \quad & \Pi_{j+1}^+(k+N'+N''-D_j) + \Pi_{j+1}^+(k-D_j+1) - \Pi_{j+1}^+(k+N'-D_j) - \Pi_{j+1}^+(k+N''-D_j+1) \\
=_{st} \quad & 0. \qquad (19)
\end{aligned}$$

The last equality holds, because either $N'' = 0$ or $(N' = 1) \wedge (N'' = 1)$ which reduces the term to 0.

The case $(a_{j+1}(\mathcal{S}1) = 1) \wedge (a_{j+1}(\mathcal{S}2) = 0)$ is not covered, because it can never occur due to $a_{j+1}(\mathcal{S}1) \leq a_{j+1}(\mathcal{S}2)$ by induction assumption. We have shown that $\Pi_j^+(k+2) + \Pi_j^+(k) - \Pi_j^+(k+1|a_{j+1}(\mathcal{S}3)) - \Pi_j^+(k+1|a_{j+1}(\mathcal{S}4)) \leq_{st} 0$ for some pre-defined actions of $\mathcal{S}3$ and $\mathcal{S}4$. If the random variables are stochastically ordered, their means are ordered as well. The potential suboptimality of these actions ensures

$$E(\Pi_j^+(k+2)) + E(\Pi_j^+(k)) - E(\Pi_j^+(k+1)) - E(\Pi_j^+(k+1)) \leq 0 \tag{20}$$

which completes the proof for any $k = 0...c-2$ and any arrival j. The proof is related to the proof of Lemma 2 in Örmeci and Burnetas (2004). ■

Proof of Proposition 2. At first, we prove the observation that only events 1 to 3 (in the two class model, consequently $n+1$ events in the n class model) can occur after accepting a customer by using sample path arguments. Assume a threshold policy with dynamic threshold level T_j at arrival j. T_j depends on the class of customer j and the known ADI. In the two-class model with $r_I + p > r_{II}$, $T_j = c$ if $CL_j = 1$ and T_j is dynamic if $CL_j = 2$. The arrival process is given by $\{(A_j, CL_j)\}$. The system state of the rental system is defined as

$$N_j^- = N_{j-1}^+ - \sum_{i=1}^{N_{j-1}^+} X_{i,j-1} \text{ and } N_j^+ = \min(T_j, N_j^- + 1) \tag{21}$$

with $X_{i,j}$ as the departure event of customer i in subinterval $]A_j, A_{j+1}]$ ($X_{i,j} = 1$ means that customer i has left the system, $X_{i,j} = 0$ means that the customer is still in the system). Now consider the difference in system states between two rental systems, I_j. Let one system ($\overline{N}$) have accepted a customer $j = 0$ and the other one not ($\underline{N}$). Then, $I_0^+ = \overline{N}_0^+ - \underline{N}_0^+ = +1$. For $j > 0$ we get

$$I_j^- = \overline{N}_j^- - \underline{N}_j^- = \begin{cases} I_{j-1}^+ - X_{\overline{N}_{j-1}^+, j-1} & I_{j-1}^+ = +1 \\ 0 & I_{j-1}^+ = 0 \end{cases} \tag{22}$$

and

$$I_j^+ = \overline{N}_j^+ - \underline{N}_j^+ = \min\Big(T_j, \underline{N}_j^- + I_j^- + 1\Big) - \min\Big(T_j, \underline{N}_j^- + 1\Big) \tag{23}$$

We make two observations from Eqs. (22) and (23): (a) $I_j^- \in \{0, +1\}$ and $I_j^+ \in \{0, +1\}$ for all $j = 0, 1, ...$ and (b) the following statements hold $\left(I_j^- = 0\right) \Rightarrow \left(I_j^+ = 0\right)$ and $\left(I_j^+ = 0\right) \Rightarrow \left(I_{j+1}^- = 0\right)$. Therefore, the system states can only differ by one customer and once they become identical, they will remain identical. The events that may lead to identical system states can only be between arrival times with $X_{\overline{N}_{j-1}^+, j-1} = 1$ (this corresponds with event 3 i.e. $n+1$) or at arrival times with $\underline{N}_j^- + 1 = T_j$. In the two-class model with $r_I + p > r_{II}$, this could either be an ADI arrival (this corresponds with event 1) or a non-ADI arrival (this corresponds with event 2). Next, we show that in the two-class model the optimal decision is given by Eq. (2). For an acceptance to be optimal, it must hold that $r_j + p_j + E(\overline{\Pi}^+) - E(\underline{\Pi}^+) > 0$. Now consider the difference in the two systems. We can express this difference as $E(\overline{\Pi}^+) - E(\underline{\Pi}^+) = \Delta\Pi_1^+ P_1(\underline{N}^-, R) + \Delta\Pi_2^+ P_2(\underline{N}^-, R) + \Delta\Pi_3^+ P_3(\underline{N}^-, R)$, which leads to Eq. (2) after some algebra. ■

Proof of Proposition 3. We prove that the optimal solution of the program converges towards an upper bound of the optimal profit of our system if $\widehat{t} \to \infty$. Next, we explain why the terminal value of the dynamic programming problem does not affect the average reward per time unit if $\widehat{t} \to \infty$. Our argumentation is as follows: If the inter-arrival time $A_{j+1} - A_j$ between an arrival and its successor is large enough, accepting the customer becomes optimal for any system state k. In an infinitely long sample path such an event occurs with probability 1. Any decision before that event is independent of the terminal value. Since the sample path is infinitely long with $\widehat{t} \to \infty$ and assuming a finite terminal value, the impact of this value on the average reward per time unit vanishes. Thus, if $\widehat{t} \to \infty$ the optimal solution of the stochastic program becomes the optimal solution of a system with infinite time horizon in which all demands are known.

Note that the optimal policy of our ADI system (with only some ADI available) is a feasible policy for the case of completely known demand (the system manager could simply ignore some of the demand information). Thus, we have

$$\Pi(\Gamma^*) \leq \lim_{\widehat{t} \to \infty} \frac{\widehat{\Pi}(\widehat{t})}{\widehat{t}}. \tag{24}$$

■

Proof of Lemma 1. The proof of this lemma proceeds as follows. At first, we split the set of eigenvalues of **Q** into two groups. For the first group, we can explicitly determine the eigenvalues in closed-form. It can be seen that the eigenvalues from this group are distinct. For the second group, we show that the eigenvalues are the same eigenvalues as the eigenvalues of a submatrix of $\mathbf{Q}$. We prove that this submatrix is symmetric, which ensures that its eigenvalues are real and distinct. We finalize the proof by showing that, even if some eigenvalues in the first group overlap with eigenvalues in the second group, its corresponding eigenvectors are linearly independent.

We begin by taking a close look at the generator matrix $\mathbf{Q}$. Using the method of determinant expansion in cofactors (cf. Strang (1980), p. 165-167), we can write the determinant of $\mathbf{Q} - w\mathbf{I_c}$ as (let w be a real number and $\mathbf{I_c}$ be the unit matrix corresponding to the dimensions of $\mathbf{Q}$)

$$|\mathbf{Q} - w\mathbf{I_c}| = \sum_{j=0}^{c}(-1)^{j+c}\left(\mathbf{Q}-w\mathbf{I_c}\right)_{cj} M_{cj} \tag{25}$$

where M_{cj} is the determinant of a submatrix of $\mathbf{Q} - w\mathbf{I_c}$ with the last row and the jth column deleted. Since $(\mathbf{Q})_{cj} = 0$ for all j, the determinant reduces to

$$|\mathbf{Q} - w\mathbf{I_c}| = (-1)^{2c}(-w)M_{cc} = -w\left|\widetilde{\mathbf{Q}} - w\mathbf{I_{c-1}}\right| \tag{26}$$

with $\widetilde{\mathbf{Q}}$ as the submatrix of $\mathbf{Q}$ with the last row and the last column deleted. We can immediately see that $w = 0$ is an eigenvalue and that all eigenvalues of the reduced matrix $\widetilde{\mathbf{Q}}$ are also eigenvalues of $\mathbf{Q}$. The resulting matrix has only one non-zero entry in the last column, so we can again apply the method of determinant expansion. If $T < c$, we can repeat the same procedure for all columns $j \geq T$ and get

$$|\mathbf{Q} - w\mathbf{I_c}| = -w\left(\prod_{j=T}^{c-1}(-j-w)\right)\left|\widehat{\mathbf{Q}} - w\mathbf{I_T}\right| \tag{27}$$

with $\widehat{\mathbf{Q}}$ as a $T \times T$ submatrix of $\mathbf{Q}$ in the upper left corner (i.e., with the first T entries in each dimension). We can see that, $w = 0$, $w = -T, -(T+1), ..., -c+1$, and the T eigenvalues of $\widehat{\mathbf{Q}}$ are the $c+1$ eigenvalues of matrix $\mathbf{Q}$.

For the special case of $T = 0$, the eigenvalues of $\mathbf{Q}$ are $w = 0, -1, ..., -c+1$ with $w = 0$ occurring twice. However, the two eigenvectors resulting from $w = 0$ are linearly independent, as can be seen from the matrix setup: Assume that $\mathbf{u} = (u_0, u_2, ..., u_c)^{\mathbf{T}}$ is an eigenvector of $\mathbf{Q}$ corresponding to $w = 0$. Then, we have $u_0 = u_1 = u_2 = ... = u_{c-1}$ so that u_{c-1} and u_c can be chosen arbitrarily. We can construct two linearly independent eigenvectors.

For $T > 0$ we have to analyze the eigenvalues of $\widehat{\mathbf{Q}}$, which we do next. Let $\mathbf{A}$ be a diagonal $T \times T$ matrix with diagonal element $(\mathbf{A})_{ii} = \sqrt{\frac{\lambda_{II}^i}{i!}}$. Then, we have

$$\begin{aligned}
\left(\mathbf{A}\widehat{\mathbf{Q}}\mathbf{A}^{-1}\right)_{ij} &= \sum_{k=0}^{T-1} (\mathbf{A})_{ik} \left(\widehat{\mathbf{Q}}\mathbf{A}^{-1}\right)_{kj} = \sum_{k=0}^{T-1} (\mathbf{A})_{ik} \sum_{n=0}^{T-1} \left(\widehat{\mathbf{Q}}\right)_{kn} \left(\mathbf{A}^{-1}\right)_{nj} \\
&= (\mathbf{A})_{ii} \left(\widehat{\mathbf{Q}}\right)_{ij} \left(\mathbf{A}^{-1}\right)_{jj} = \sqrt{\frac{\lambda_{II}^i}{i!}\frac{j!}{\lambda_{II}^j}} \left(\widehat{\mathbf{Q}}\right)_{ij}. \qquad (28)
\end{aligned}$$

It follows that matrix $\mathbf{A}\widehat{\mathbf{Q}}\mathbf{A}^{-1}$ is symmetric, because the only non-zero, off-diagonal entries of $\widehat{\mathbf{Q}}$ are $\left(\widehat{\mathbf{Q}}\right)_{i,i-1} = i$ and $\left(\widehat{\mathbf{Q}}\right)_{i,i+1} = \lambda_{II}$, and so the following equation holds

$$\begin{aligned}
\left(\mathbf{A}\widehat{\mathbf{Q}}\mathbf{A}^{-1}\right)_{i,i-1} &= \sqrt{\frac{\lambda_{II}^i}{i!}\frac{(i-1)!}{\lambda_{II}^{i-1}}} \left(\widehat{\mathbf{Q}}\right)_{i,i-1} \\
&= \sqrt{\frac{\lambda_{II}}{i}} \left(\widehat{\mathbf{Q}}\right)_{i,i-1} = \sqrt{\lambda_{II} i} = \lambda_{II}\sqrt{\frac{i}{\lambda_{II}}} \\
&= \sqrt{\frac{\lambda_{II}^{i-1}}{(i-1)!}\frac{i!}{\lambda_{II}^i}} \left(\widehat{\mathbf{Q}}\right)_{i-1,i} = \left(\mathbf{A}\widehat{\mathbf{Q}}\mathbf{A}^{-1}\right)_{i-1,i}. \qquad (29)
\end{aligned}$$

The symmetric matrix $\mathbf{A}\widehat{\mathbf{Q}}\mathbf{A}^{-1}$ corresponds with an undirected graph whose T vertices are connected in one line (because its only non-zero entries are located in the diagonal, the sub-diagonal, and the super-diagonal). Thus, the graph complies with the definition of a tree. Leal-Duarte and Johnson (2002) have shown that symmetric matrices whose graphs form trees have at least as many distinct eigenvalues as vertices in the longest path of the tree. It follows that matrix $\mathbf{A}\widehat{\mathbf{Q}}\mathbf{A}^{-1}$ has at least T distinct eigenvalues. Since it can have at most T eigenvalues, we conclude that all eigenvalues of matrix $\mathbf{A}\widehat{\mathbf{Q}}\mathbf{A}^{-1}$ are distinct. They are also real, because symmetric matrices always have real eigenvalues (cf. Strang (1980), p. 222-224). Matrix $\mathbf{A}\widehat{\mathbf{Q}}\mathbf{A}^{-1}$ is similar to $\widehat{\mathbf{Q}}$, so $\widehat{\mathbf{Q}}$ has the same real and distinct eigenvalues. Eigenvectors of distinct eigenvalues are always linearly independent (cf. Strang (1980), p. 192). Note that the eigenvalues of $\widehat{\mathbf{Q}}$ are also non-zero due to the following argumentation: $\widehat{\mathbf{Q}}$ differs from the generator matrix $\mathbf{Q}_{Erlang}$ of an Erlang Loss system with $T-1$ servers, arrival rate λ_{II} and departure rate 1 only by the last entry, $\left(\widehat{\mathbf{Q}}\right)_{T-1T-1} = 1 - \lambda_{II} - T \neq 1 - T = (\mathbf{Q}_{Erlang})_{T-1T-1}$. Because $\mathbf{Q}_{Erlang}$ has a steady state solution, $\boldsymbol{\pi}\mathbf{Q}_{Erlang} = \overline{\mathbf{0}}$ holds for a certain row vector $\boldsymbol{\pi} = (\pi_0, ..., \pi_{T-1})$ with all $\pi_i > 0$. Then, we see by equating the last columns of matrices $\widehat{\mathbf{Q}}$ and $\mathbf{Q}_{Erlang}$ that $\boldsymbol{\pi}\widehat{\mathbf{Q}} = \overline{\mathbf{0}}$ could only hold if

$$T - 1 = \lambda_{II} + T - 1, \tag{30}$$

which is not possible due to $\lambda_{II} > 0$. Thus, $w = 0$ cannot be an eigenvalue of $\widehat{\mathbf{Q}}$.

Next, we show that even if some eigenvalues of $\widehat{\mathbf{Q}}$ overlap with the eigenvalues $w = -T, -(T+1), ..., -c+1$, the corresponding eigenvectors are linearly independent so that the algebraic multiplicity and the geometric multiplicity of these eigenvalues coincide. We do this by showing that, if $w \in \{-T, -(T+1), ..., -c+1\}$ is also an eigenvalue of $\widehat{\mathbf{Q}}$, $\mathbf{Q}$ has two linearly independent eigenvectors $\widetilde{u}$ and $\widetilde{v}$ corresponding to w.

Let $\mathbf{B}$ be a submatrix of $\mathbf{Q}$ from the lower right corner except the last row and column, i.e., $\mathbf{B}$ looks like

$$\mathbf{B} = \begin{pmatrix} -T & 0 & 0 & 0 \\ T+1 & -T-1 & 0 & 0 \\ ... & ... & ... & 0 \\ 0 & 0 & c-1 & -c+1 \end{pmatrix}. \tag{31}$$

Clearly, w is an eigenvalue of $\mathbf{B}$, because of $w \in \{-T, -(T+1), ..., -c+1\}$. Assume that vector $\mathbf{u}$ is an eigenvector of $\mathbf{B}$ that corresponds to w. Then, vector $\widetilde{\mathbf{u}} = \begin{pmatrix} \overline{\mathbf{0}} \\ \mathbf{u} \\ 0 \end{pmatrix} \in \mathbb{R}^{c+1}$ is an eigenvector of $\mathbf{Q}$, because from $\mathbf{B}\mathbf{u} = w\mathbf{u}$ follows $\mathbf{Q}\widetilde{\mathbf{u}} = w\widetilde{\mathbf{u}}$.

Now, assume that $\mathbf{v} = (v_0, v_1, ..., v_{T-1})^{\mathbf{T}}$ is an eigenvector of submatrix $\widehat{\mathbf{Q}}$ corresponding to the eigenvalue w so that $\widehat{\mathbf{Q}}\mathbf{v} = w\mathbf{v}$. The eigenvector of the complete matrix $\mathbf{Q}$ looks like $\widetilde{\mathbf{v}} = \begin{pmatrix} \mathbf{v} \\ \mathbf{x} \\ 0 \end{pmatrix} \in \mathbb{R}^{c+1}$ with $\mathbf{x} \in \mathbb{R}^{c-T}$ such that condition $\mathbf{Q}\widetilde{\mathbf{v}} = w\widetilde{\mathbf{v}}$ is met. The components x_i of $\mathbf{x}$ exist, because there must be a vector $\widetilde{\mathbf{v}}$ that satisfies $\mathbf{Q}\widetilde{\mathbf{v}} = w\widetilde{\mathbf{v}}$, and they can be derived recursively.

As we have shown, vectors $\widetilde{\mathbf{u}}$ and $\widetilde{\mathbf{v}}$ are eigenvectors of $\mathbf{Q}$ corresponding to eigenvalue w. Note that each of the subvectors $\mathbf{u}$ and $\mathbf{v}$ has at least one non-zero component, because both are eigenvectors by themselves. Then, vectors $\widetilde{\mathbf{u}}$ and $\widetilde{\mathbf{v}}$ are linearly independent, because

$$\alpha_1 \widetilde{\mathbf{u}} + \alpha_2 \widetilde{\mathbf{v}} = \overline{\mathbf{0}} \Leftrightarrow \begin{pmatrix} \overline{\mathbf{0}} + \alpha_2 \mathbf{v} \\ \alpha_1 \mathbf{u} + \alpha_2 \mathbf{x} \\ 0 \end{pmatrix} = \begin{pmatrix} \overline{\mathbf{0}} \\ \overline{\mathbf{0}} \\ 0 \end{pmatrix} \tag{32}$$

can hold only if $\alpha_1 = \alpha_2 = 0$. Scalar α_2 must be zero to have $\alpha_2 \mathbf{v} = \overline{\mathbf{0}}$ and this also requires $\alpha_1 = 0$ in order to have $\alpha_1 \mathbf{u} + 0\mathbf{x} = \overline{\mathbf{0}}$.

We have shown that all eigenvalues of $\mathbf{Q}$ are real and that there can be at most pairwise identical eigenvalues. For these eigenvalues, the eigenvectors must be linearly independent. Thus, all eigenvectors of $\mathbf{Q}$ are linearly independent. This ensures diagonalizability (cf. Strang (1980), p. 190). ■

Proof of Lemma 2. We proceed by calculating the probabilities of each path in the probability tree given in Figure 4. The probability of event $S \leq t_n - t_{n-1}$ with $t_n - t_{n-1}$ as the time between two consecutive ADI arrivals $n-1$ and n is conveniently given by $F_S(t_n - t_{n-1}) = 1 - e^{-(t_n - t_{n-1})}$. The probability that the CTMC is in state j at the end of the subinterval, $P(N^- = j)$, is given by $(\boldsymbol{\pi}_n^-)_j$. The cdf of the first passage time of the system into state c, $T_{\cdot c}$, given the initial distribution $\boldsymbol{\pi}_{n-1}^+$, is

$$\begin{aligned} P\left(T_{\cdot c} \leq t\right) &= \left(\boldsymbol{\pi}_{n-1}^+\right) \left(\mathbf{P}(t)\right)_c = \left(\boldsymbol{\pi}_{n-1}^+\right) \left(\mathbf{U}\mathbf{D}(t)\mathbf{U}^{-1}\right)_c \\ &= \sum_{k=0}^{c} e^{w_k t} \sum_{i=0}^{c} \left(\boldsymbol{\pi}_{n-1}^+\right)_i \mathbf{U}_{ik} \mathbf{U}_{kc}^{-1}. \end{aligned} \tag{33}$$

Note that Equation (33) holds because state c is absorbing. We begin by determining p_{3n}. We have that

$$
\begin{aligned}
p_{3n} &= P(S \le t_n - t_{n-1} \wedge N^- < c) + P(S \le t_n - t_{n-1} \wedge N^- = c \wedge S \le T_{\cdot c}) \\
&= P(S \le T_{\cdot c} \wedge S \le t_n - t_{n-1}) \\
&= P(S \le \min(T_{\cdot c}, A_n - A_{n-1})) \\
&= \int\limits_0^{+\infty} P(S = t)\left[1 - P(\min(T_{\cdot c}, t_n - t_{n-1}) \le t)\right] dt \\
&= \int\limits_0^{+\infty} e^{-t} dt - \int\limits_0^{t_n - t_{n-1}} e^{-t} P(T_{\cdot c} \le t) dt - \int\limits_{t_n - t_{n-1}}^{+\infty} e^{-t} P(t_n - t_{n-1} \le t) dt \\
&= 1 - \int\limits_{t_n - t_{n-1}}^{+\infty} e^{-t} dt - \int\limits_0^{t_n - t_{n-1}} e^{-t} P(T_{\cdot c} \le t) dt \\
&= 1 - e^{-(t_n - t_{n-1})} - \int\limits_0^{t_n - t_{n-1}} \sum_{k=0}^{c} e^{-t} e^{w_k t} \sum_{i=0}^{c} \left(\boldsymbol{\pi}_{n-1}^{+}\right)_i \mathbf{U}_{ik} \mathbf{U}_{kc}^{-1} dt \\
&= 1 - e^{-(t_n - t_{n-1})} - \sum_{k=0}^{c} \left(\sum_{i=0}^{c} \left(\boldsymbol{\pi}_{n-1}^{+}\right)_i \mathbf{U}_{ik} \mathbf{U}_{kc}^{-1} \right) \int\limits_0^{t_n - t_{n-1}} e^{(w_k - 1)t} dt \\
&= 1 - e^{-(t_n - t_{n-1})} - \sum_{k=0}^{c} \left(\sum_{i=0}^{c} \left(\boldsymbol{\pi}_{n-1}^{+}\right)_i \mathbf{U}_{ik} \mathbf{U}_{kc}^{-1} \right) \alpha_k(t_n - t_{n-1}) \qquad (34)
\end{aligned}
$$

and

$$
\alpha_k(x) = \int\limits_0^{x} e^{(w_k - 1)t} dt = \begin{cases} \frac{\left(e^{(w_k - 1)x} - 1\right)}{w_k - 1} & w_k \ne 1 \\ x & w_k = 1 \end{cases} . \qquad (35)
$$

Likewise, the probability of event 2 occurring first, p_{2n}, is given by

$$\begin{aligned}
p_{2n} &= P(S > t_n - t_{n-1} \wedge N^- = c) + P(S \le t_n - t_{n-1} \wedge N^- = c \wedge S > T_{\cdot c}) \\
&= P(S > t_n - t_{n-1} \wedge N^- = c \wedge S > T_{\cdot c}) + P(S \le t_n - t_{n-1} \wedge N^- = c \wedge S > T_{\cdot c}) \\
&= P(N^- = c \wedge S > T_{\cdot c}) \\
&= P(T_{\cdot c} \le t_n - t_{n-1} \wedge T_{\cdot c} \le S) \\
&= P(T_{\cdot c} \le \min(S, t_n - t_{n-1})) \\
&= \int_0^{+\infty} P(T_{\cdot c} \le \min(t, t_n - t_{n-1}))P(S = t)dt \\
&= \int_0^{t_n - t_{n-1}} P(T_{\cdot c} \le t)e^{-t}dt + \int_{t_n - t_{n-1}}^{+\infty} P(T_{\cdot c} \le t_n - t_{n-1})e^{-t}dt \\
&= \sum_{k=0}^{c} \left(\sum_{i=0}^{c} \left(\boldsymbol{\pi}_{n-1}^+\right)_i \mathbf{U}_{ik}\mathbf{U}_{kc}^{-1} \right) \alpha_k(t_n - t_{n-1}) + \left(\boldsymbol{\pi}_n^-\right)_c e^{-(t_n - t_{n-1})}. \qquad (36)
\end{aligned}$$

The probability of event 1 is easier to compute with

$$\begin{aligned}
p_{1n} &= P\left(S > t_n - t_{n-1} \wedge N^- = c - 1\right) \\
&= P\left(S > t_n - t_{n-1}\right) P\left(N^- = c - 1\right) \\
&= e^{-(t_n - t_{n-1})} \left(\boldsymbol{\pi}_n^-\right)_{c-1}, \qquad (37)
\end{aligned}$$

because the two random variables S and N^- are independent. ■

Proof of Proposition 4. The proof is straightforward and proceeds by induction. Start with the first decision for arrival $j = 0$. Define a policy $\widehat{\Gamma}_0$ which uses the decision logic of the ADI policy for arrival $j = 0$ and the decision logic of the optimal state-dependent threshold policy Γ^T for all future admission decisions. In this case, the computation of the event probabilities (p_{10}, p_{20}, and p_{30} for the two-class model) is exact, because the ADI policy also assumes a state-dependent threshold policy for future decisions. For this policy, it holds that

$$\Pi\left(\widehat{\Gamma}_0\right) \ge \Pi\left(\Gamma^T\right), \qquad (38)$$

because the decision of $\widehat{\Gamma}_0$ for $j = 0$ is optimal, i.e.,

$$\Pi\left(\widehat{\Gamma}_0\right) = \max(r_0 + \Pi_0^+\left(k+1\right), \Pi_0^+\left(k\right)) \geq \Pi\left(\Gamma^T\right). \tag{39}$$

Note that we write r_j for the reward of accepting customer j. This is $r_I + p$ for an ADI customer in the two class model, because accepting a customer avoids the penalty p.

Assume a policy $\widehat{\Gamma}_{j-1}$, which admits all customers after $j-1$ with a state-dependent threshold policy. Then, we can define a policy $\widehat{\Gamma}_j$ for arrival j, which takes the same decisions as $\widehat{\Gamma}_{j-1}$ up to arrival $j-1$, decides with the decision logic of the ADI policy for arrival j and behaves like a state-dependent threshold policy Γ^T for all future arrivals.

Then, we have

$$\begin{aligned} \Pi\left(\widehat{\Gamma}_j\right) &= \sum_{k=0}^{j-1} r_k a_k + \max(r_j + \Pi_j^+\left(k+1\right), \Pi_j^+\left(k\right)) \\ &\geq \sum_{k=0}^{j-1} r_k a_k + r_j a_j' + \Pi_j^+\left(k + a_j'\right) = \Pi\left(\widehat{\Gamma}_{j-1}\right). \end{aligned} \tag{40}$$

for an arbitrary decision a_j' of policy $\widehat{\Gamma}_{j-1}$. Note that the two policies differ only in the decision for arrival j. Assume that $\Pi\left(\widehat{\Gamma}_{j-1}\right) \geq \Pi\left(\Gamma^T\right)$ holds (which we have shown for $j-1=0$). From Equation (40) follows by induction that

$$\Pi\left(\widehat{\Gamma}_j\right) \geq \Pi\left(\widehat{\Gamma}_{j-1}\right) \geq \Pi\left(\Gamma^T\right) \tag{41}$$

for all $j = 0, 1, 2, \ldots$, even when $j \to +\infty$. It is obvious that

$$\Pi\left(\Gamma^{ADI}\right) = \lim_{j\to\infty} \Pi\left(\widehat{\Gamma}_j\right) \geq \Pi\left(\Gamma^T\right), \tag{42}$$

because we replace all decisions of policy Γ^T by the decisions of the ADI policy Γ^{ADI}. ■

Proof of Lemma 3.

This lemma can be directly derived by adapting the event structure shown in Figure 4 to a model with n customer classes. Note that between two consecutive ADI arrivals in interval m, only events $1 \ldots m$ or $n+1$ may occur because we know that no ADI customers not included in R of classes with longer demand leadtime arrive. In addition, only event $i = CL_j$ can occur at the end of a subinterval, because customer j arrives at this point in time. ■

E Comparison of ADI Policy with Optimal Discrete-Time Policy

We used an upper bound to evaluate the performance of the ADI policy and showed that the ADI policy performs close to the upper bound. Since we do not know the optimal policy, we cannot compare the performance of the ADI policy against the optimal policy. However, for small problems, we can use a discrete-time model and solve it optimally to evaluate the performance of the ADI policy against this solution. We can choose a short time interval, such that the performance of the discrete-time model approaches the performance of the continuous time model we consider. In this section, we present and solve such a discrete-time model and compare the performance of the ADI policy against this benchmark.

In the discrete-time model, we require that the demand lead time is equal to an integer number of periods. We assume that there exists a maximum number of ADI reservations, MR, that the system is allowed to keep at any time (we set MR large enough, such that the assumption is essentially not restrictive). Then, the state space of the system at the beginning of a period is finite and is described by the number of busy cars N, the number of reservations $0 \leq R \leq MR$, and the (integer) expiration time of each reservation. The sequence of events of each period is as follows: First, we accept ADI customers (up to the number of cars available). Then, we decide on how many non-ADI customers to accept. We then serve that number of non-ADI customers. If less non-ADI customers arrive than we would allow into the system, we serve all arriving non-ADI customers. Finally, customers (from both classes) leave the system.

The resulting discrete-time model has a finite state space and can be solved optimally. We use relative value iteration to solve the model and terminate the algorithm after the span of the value vector reaches 0.01 (e.g., Puterman (1994)). We consider an example with demand lead time $\tau = 7$ days, mean rental time $\bar{S} = 14$ days, demand rate $\lambda = 0.38$ cars/day, fraction of ADI demands $\nu = 30\%$, fleet size $c = 5$ cars, unit ADI revenue $r_I = 350$ Euro/rental, unit non-ADI revenue $r_{II} = 200$ Euro, and unit penalty cost $p = 850$ Euro/loss. The revenue and cost data of the example are based on the previous examples. The demand rate and fleet size are chosen

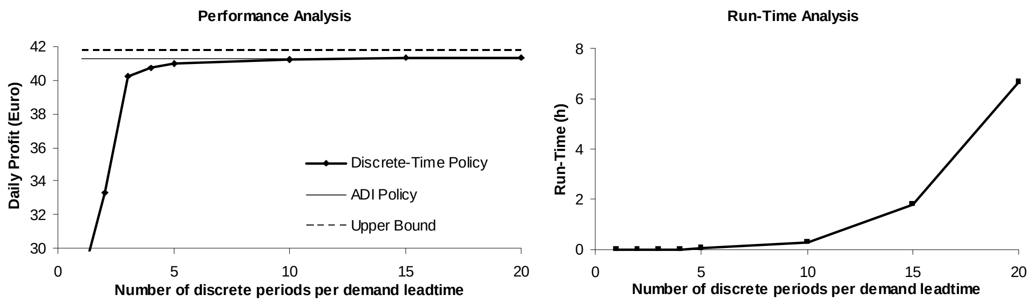

Figure 9: Performance and run-time analysis of discrete-time policy for varying numbers of periods, benchmarked against ADI policy and upper bound

small enough, such that solving the dynamic program is computationally feasible. In the numerical experiments, we use $MR = 4$ cars, such that the probability of ADI demand within the demand lead time exceeding MR is below 3%. Increasing MR from 4 to 5 did not affect the results but increased run times considerably.

The solution of the discrete-time model converges to the solution of the continuous time model as we decrease the length of the periods. We analyze the convergence by solving the discrete-time model for various period lengths. The results are shown in Figure 9 and indicate that the discrete-time model converges relatively quickly - the performance of the model with 15 periods per demand lead time is identical to the performance of the model with 20 periods per demand lead time. The numerical results also indicate that the ADI policy performs close to this benchmark: The expected profit of the discrete-time model (with 15 periods per demand lead time) is 41.3 Euro/day and the expected profit of the ADI policy is 41.2 Euro/day (0.2% below). Figure 9 also shows the run-times required for solving the discrete-time model on a 2 GHz PC with 1 GB of main memory.

Based on the analyses of this section, we conclude that the performance of the optimal policy in a discrete-time setting converges quickly to the performance of the continuous time model as we decrease the lengths of the periods. For small problems, we can build on this observation and use the discrete-time model to approximate very closely the performance of the optimal policy in continuous time. The results show that the performance of the ADI policy is very close to the performance of this policy, indicating that the performance of the ADI policy is very close to optimality.

F Extensions

In this section, we present several extensions to our main model: In Subsection F.1 we modify the main model to consider situations in which ADI customers receive a discount rather than pay a premium for providing ADI and receiving a service guarantee. This case applies when the rental company passes some of the savings from using ADI on to the customers who provide it. In Subsection F.2 we extend our model to an arbitrary number of ADI customer classes. For each extension, we show what results from the preceding sections carry over to the modified model and how the remaining results need to be modified in order to hold. In the last part of this section, we analyze the performance of our model in a setting where rental time distributions are class-specific.

F.1 ADI Price Discounts

So far, we have assumed that $r_I + p \geq r_{II}$. We next analyze the case when $r_I + p < r_{II}$.

The Optimal Policy Propositions 1 and 2 and their corresponding proofs hold without modification, because they do not depend on a specific cost structure. Therefore, the optimal policy is still a threshold policy, such that customers for a specific class are only accepted up to the threshold. However, if $r_I + p < r_{II}$, the non-ADI class has the higher unit revenues than the ADI class and non-ADI customers will always be accepted under the optimal policy. ADI customers will only be accepted up to a certain level, which depends on the set of reservations R. The event structure described in Section 4.2, which is general enough to cover both cases, $r_I + p \geq r_{II}$ and $r_I + p < r_{II}$, remains the same and the optimal decision is given by Equation (2). The lower and upper bounds are also valid, because the dynamic programs on which the bounds are based do not depend on a specific cost situation.

The ADI Policy The ADI policy, however, needs to be adapted to the modified model, because the preferred class is now the non-ADI class. Equation (3) still holds because it is independent of the cost structure. However, the modified policy differs from the original ADI policy, since now all non-ADI customers are accepted. We can model this situation by using the same CTMC representation as described in Subsection 5.2 but with a non-ADI threshold level equal to the fleet size, i.e., $T = c$. The policy differs further in the representation of ADI acceptances. While in the original model all ADI arrivals were accepted, we now accept ADI arrivals only if $N < T_{ADI}$ (in direct analogy to accepting non-ADI customers only if $N < T$ in the original model). T_{ADI} is the constant threshold level for the ADI class that can be derived from the benchmark policy described in Subsection 4.3. Event 1 occurs at the end of a subinterval between two consecutive ADI arrivals only if that customer is accepted in one scenario and rejected in the other scenario. This occurs only if $N^- = T_{ADI} - 1$. The event probability of event 1 within a subinterval (as described in Subsection 5.3 for the original model), p_{1n}, is now given by

$$\begin{aligned} p_{1n} &= P\left(S > t_n - t_{n-1} \wedge N^- = T_{ADI} - 1\right) \\ &= P\left(S > t_n - t_{n-1}\right) P\left(N^- = T_{ADI} - 1\right) \\ &= e^{-(t_n - t_{n-1})} \left(\boldsymbol{\pi}_n^-\right)_{T_{ADI}-1}. \end{aligned}$$

Note that the two random variables S as the rental duration of the additional customer and N^- as the number of busy cars are independent.

Equation (8) also needs to be modified and becomes

$$\begin{aligned} \left(\boldsymbol{\pi}_n^+\right)_j &= P\left(N^+ = j \mid N^- \neq \{T_{ADI} - 1, c\}\right) \\ &= \begin{cases} \frac{1}{1-\left(\boldsymbol{\pi}_n^-\right)_{T_{ADI}-1} - \left(\boldsymbol{\pi}_n^-\right)_c} \left(\boldsymbol{\pi}_n^-\right)_{j-1} & 0 < j < T_{ADI} \\ \frac{1}{1-\left(\boldsymbol{\pi}_n^-\right)_{T_{ADI}-1} - \left(\boldsymbol{\pi}_n^-\right)_c} \left(\boldsymbol{\pi}_n^-\right)_{j} & T_{ADI} < j < c \\ 0 & \textit{otherwise.} \end{cases} \end{aligned}$$

Similarly, the DTMC that models the event behavior after demand leadtime τ must be adapted by changing the transition matrix $\mathbf{P}_{\tau+}$ to

$$(\mathbf{P}_{\tau+})_{ij} = \begin{cases} \frac{\lambda_I + \lambda_{II}}{\delta_i} & i < T_{ADI} - 1;\ j = i+1 \\ \frac{\lambda_{II}}{\delta_i} & T_{ADI} - 1 \leq i < c;\ j = i+1 \\ \frac{\lambda_I}{\delta_i} & i = T_{ADI} - 1;\ j = c+1 \\ \frac{i}{\delta_i} & 0 < i < c;\ j = i-1 \\ \frac{1}{\delta_i} & i < c;\ j = c+2 \\ 1 & i \geq c;\ j = i \\ 0 & \textit{otherwise.} \end{cases}$$

The algorithm in the appendix virtually remains identical, but must be executed for every ADI arrival instead of for every non-ADI arrival. Proposition 4 holds without modification, because it does not rely on a specific cost structure, implying that the ADI policy dominates the best policy without any ADI also for $r_I + p < r_{II}$.

F.2 Extension to Multiple ADI Classes

So far, we have restricted the analysis to two customer classes, one ADI class and one non-ADI class. In this subsection, we extend the model to multiple ADI classes. We show the our approach can be extended to cover this extension but not that actually computing the policies can be computationally expensive. We consider a model with n classes with each class i having a class-specific demand rate λ_i, demand leadtime τ_i, and contribution margin r_i. The non-ADI class is class 1, i.e., $\tau_1 = 0$ and customers from this class do not receive a compensation if their demand cannot be met ($p_1 = 0$). Without loss of generality, let the ADI classes $i = 2, ..., n$ be ordered by increasing demand leadtime, such that $\tau_i \leq \tau_{i+1}$ for $i = 2, ..., n-1$. We have seen at the main model that a longer demand leadtime leads to higher profit, and we therefore require that the unit revenues r_i are decreasing in the demand leadtime, i.e., classes with a longer demand leadtime have lower unit revenues ($r_i \geq r_{i+1}$ for $i = 2, ..., n-1$). The non-ADI class $i = 1$ does not pay a premium for a service guarantee and $r_1 < r_n$. Customers from the ADI classes (classes $2...n$) receive a compensation $p_i = p$ if their announced demand cannot be met. Note that we could also model different penalties per class as long as $r_i + p_i \geq r_{i+1} + p_{i+1}$ is ensured. We normalize rental times such that $\overline{S} = 1$. The profit function becomes

$$\Pi(\Gamma) = \sum_{i=1}^{n} r_i V_i(\Gamma) - p(\sum_{i=2}^{n} (\lambda_i - V_i(\Gamma))).$$

The Optimal Policy In the n class model, the state space is again characterized by the number of busy cars N and the set of ADI R. R, however, does not only contain the arrival epochs of ADI customers but also their class origin, i.e., $R = \{(t_1, CL_1), (t_2, CL_2), ...\}$ where $CL_j \in \{2, ..., n\}$ denotes the class of customer j. The proof of Proposition 1 is general enough to cover a model with n classes. Proposition 1 therefore holds without modification and the optimal policy is still a threshold policy. In the n class model, however, more than three events may occur after acceptance of a customer:

Event 1 The accepted customer prevents a non-ADI customer (class 1) from entering the system

Event 2 The accepted customer prevents a class 2 ADI customer from entering the system

...

Event n The accepted customer prevents a class n ADI customer from entering the system

Event n+1 The accepted customer leaves without preventing other customers from entering the system

Since the optimal policy of the n class model is a threshold policy, we can again formulate the optimal decision based on $n+1$ different events that may occur. Such an optimal decision is described in Corollary 1, which is an extension of Proposition 2 of the two class model to an n class model.

Corollary 1 *For a model with n customer classes and any threshold policy, the optimal decision has the following property*

$$a_j^* = \begin{cases} \mathbf{1}(N^- < c), & r_{CL_j} + p_{CL_j} - \sum_{i=1}^{n} (r_i + p_i) P_i(N^-, R) > 0 \\ 0, & else. \end{cases}$$

$P_i(N^-, R)$ *denotes the probability that event i is the next event after the current event (define* $p_1 = 0$*).*

The proof of Corollary 1 follows directly from the proof of Proposition 2. As a lower bound we use again the best non-ADI policy.

Because of the cost structure of the n class model, it follows that $T_i^* \geq T_{i+1}^*$ (for $i > 1$) with $T_2^* = c$ and $T_1^* \leq T_n^*$. To keep the notation concise, we abbreviate T_i^* by T_i. The upper bound of the n class model is similar to the upper bound of the two class model but with an optimal profit (or value) function $\Pi_k(i)$ for arrival k of

$$\Pi_k(i) = \max_{a \in \{0,1\}} \left[r_{CL_j} a - p_{CL_j} (1-a) + \sum_j p_{ij}^k(a) \Pi_{k+1}(j) \right]$$

for any states $i, j \in S$.

The ADI Policy The ADI policy can be extended to the n class model. We use the idea of assuming that all future arrivals are admitted based on the best non-ADI policy, the threshold policy. For the n class model we use the threshold T_i for class i, given by the optimal policy without ADI. We had two different CTMC setups in the two class model (within demand leadtime and after demand leadtime) and now we have n different CTMC setups. Consider the future time line after a customer arrival as a sequence of intervals $]A_0 + \tau_i, A_0 + \tau_{i+1}]$ for $i = 1...n-1$ and finally $]A_0 + \tau_n, +\infty[$. For each of these intervals, a series of CTMCs can be constructed. During the first interval, all ADI arrivals, i.e., all arrivals of classes $2...n$ are known. In the next interval, only arrivals of classes $3...n$ are known while arrivals of classes 1 and 2 are not known. In the interval before the last one, $]A_0 + \tau_{n-1}, A_0 + \tau_n]$, ADI exist only for class n customers, and in the last interval no arrivals are known. In each interval, we start with the initial state distribution given by the end-of-interval distribution from the previous interval. Then, we proceed by the same method that we used in Subsection 5.1 for the two-class model, i.e., we split each interval m into subintervals which are separated by the known ADI customer arrival epochs. The dimension of generator matrix $\mathbf{Q}_m^{multiple}$ in interval m depends on m, because in each subinterval only events $1...m$ and event $n+1$ can occur (event i corresponds to blocking a ADI customer of class i, event $n+1$ corresponds to no blocking) and arrivals of ADI customers within the demand leadtime of their respective class only occur at time epochs included in R. The generator matrix $\mathbf{Q}_m^{multiple}$ is given by

$$\left(\mathbf{Q}_m^{multiple}\right)_{ij} = \begin{cases} \sum_{\{s|i<T_s-1\}} \lambda_s & j = i+1; s = 1...m \\ \lambda_s & i = T_s - 1;\ j = c+s-1; s = 1...m \\ i & 0 < i < c;\ j = i-1 \\ -\sum_{\forall a \neq i} (\mathbf{Q})_{ia} & j = i \\ 0 & otherwise, \end{cases}$$

which is an extension of Equation (4).

The individual event probabilities are connected by adding up the probabilities of each path of the probability tree, i.e., Equation (3) generalizes to

$$P_i(N_j^-, R_j) = \sum_{l=1}^{|R|+1} p_{ij,l} \prod_{k=1}^{l-1} \left(1 - \sum_{e=1}^{n+1} p_{\mathbf{e}j,k}\right) + p_{ij,\tau_{II}+} \prod_{k=1}^{|R|+1} \left(1 - \sum_{e=1}^{n+1} p_{\mathbf{e}j,k}\right).$$

Once a CTMC setup for an interval m is constructed, the event probabilities within that interval can be determined as shown in Subsection 5.3. Some equations, however, need to be generalized:

Lemma 3 *The probability of event i occurring first within or at the end of a subinterval before customer j in interval m are given by*

$$\begin{aligned}
P_j(Event\ \mathbf{CL}_j) &= P\left(S > t_j - t_{j-1} \wedge N^- = T_{CL_j} - 1\right), \\
P_j(Event\ \mathbf{i}) &= P\left(S > T_{\cdot\infty} \wedge N^- = c + i - 1\right) \text{ for } i = 1...m, \\
P_j(Event\ \mathbf{i}) &= 0 \text{ for } i = m+1...n,\ i \neq CL_j, \text{ and} \\
P_j(Event\ \mathbf{n+1}) &= P(S \leq \min(T_{\cdot\infty}, t_j - t_{j-1}))
\end{aligned}$$

with $T_{\cdot\infty}$ as the first passage time from any state into a state $i \geq c$ (i.e., an absorbing state) of a CTMC with generator $\mathbf{Q}_m^{multiple}$, S as an exponentially distributed random variable with $E(S) = 1$, and $t_j - t_{j-1}$ as the time between two consecutive ADI customers j and $j-1$.

At the end of each subinterval, the end-of-interval distribution needs to be updated. Equation (8) can be generalized to

$$\begin{aligned}
\left(\boldsymbol{\pi}_{n,m}^+\right)_j &= P\left(N^+ = j \mid N^- \neq T_{CL} - 1 \wedge N^- < c\right) \\
&= \begin{cases} \frac{1}{1 - \left(\boldsymbol{\pi}_n^-\right)_{T_{CL}-1} - \sum_{s=c}^{c+m-1} \left(\boldsymbol{\pi}_n^-\right)_s} \left(\boldsymbol{\pi}_n^-\right)_{j-1} & 0 < j < T_{CL} \\ \frac{1}{1 - \left(\boldsymbol{\pi}_n^-\right)_{T_{CL}-1} - \sum_{s=c}^{c+m-1} \left(\boldsymbol{\pi}_n^-\right)_s} \left(\boldsymbol{\pi}_n^-\right)_j & T_{CL} < j < c \\ 0 & otherwise, \end{cases}
\end{aligned}$$

denoting with CL the class of the customer that arrives at the end of the subinterval. After time $A_0 + \tau_n$, no further arrivals are known, and, similar to the case with only two customer classes, the future behavior must be modelled. Any of the $n+1$ events may occur. The DTMC from Subsection

5.4 must be adapted by modifying the transition matrix $\mathbf{P}_{\tau+}$ to $\mathbf{P}_{\tau_n+}$ with

$$(\mathbf{P}_{\tau_n+})_{ij} = \begin{cases} \frac{\sum_{\{s|i<T_s-1\}} \lambda_s}{\delta_i} & j = i+1 \\ \frac{\lambda_s}{\delta_i} & i = T_s - 1;\ j = c+s-1 \\ \frac{i}{\delta_i} & 0 < i < c;\ j = i-1 \\ \frac{1}{\delta_i} & i < c;\ j = c+n \\ 1 & i \geq c;\ j = i \\ 0 & otherwise, \end{cases}$$

where the δ_i are the corresponding weights that normalize each row sum of the matrix to 1. The absorbing matrix $\mathbf{P}_{\tau_n+}$ can be rewritten as before such that $\mathbf{A} = (\mathbf{I_c} - \widetilde{\mathbf{P}})^{-1}\mathbf{C}$ and $p_{i,\tau_n+} = \boldsymbol{\pi}_{|R|+1} \cdot (\mathbf{A})_i$ for event i.

As with the two class model, the ADI policy maximizes the expected return under the assumption that future decision makers do not use ADI information. Therefore, Proposition 4 holds also for the n class model. The general form of the proof of Proposition 4 covers the extended model.

F.3 Non-Identical Rental Time Distributions

So far, we have assumed that the rental time distributions of ADI and non-ADI customers are identical. In this subsection, we relax this assumptions and propose a heuristic for solving problems with class-dependent rental time distributions.

Let $\bar{S}_I$ be the mean rental time of ADI customers and $\bar{S}_{II}$ be the mean rental time of non-ADI customers. Under the heuristic admission policy we propose, we execute the ADI policy using a mean rental time of $\hat{S}, \min(\bar{S}_I, \bar{S}_{II}) \leq \hat{S} \leq \max(\bar{S}_I, \bar{S}_{II})$. More precisely, we decide on which customers to accept and which customers to reject using an ADI policy that assumes a mean rental time $\hat{S}$. To find the best value for $\hat{S}$, we simulate a system with class dependent rental times that operates under the ADI policy with mean rental time of $\hat{S}$. We repeat the simulation for various values of $\hat{S}$ and choose the value of $\hat{S}$ that results in the highest simulated expected profit as the parameter for the ADI policy.

To evaluate the performance of the heuristic, we follow the approach of Section E: We use a discrete-time model with short period lengths to approximate the continuous-time model and solve

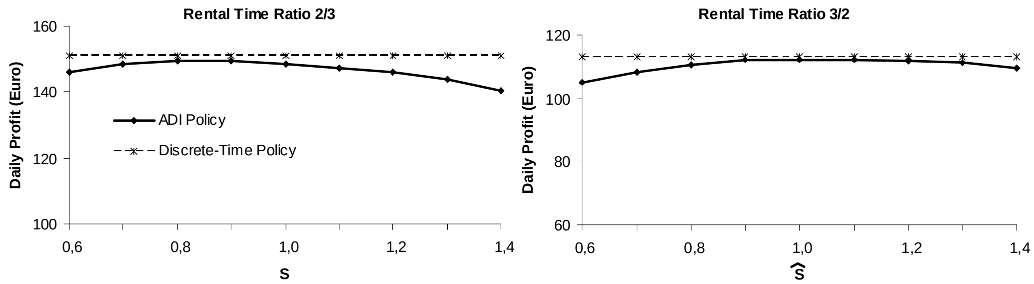

Figure 10: Expected daily profit for ADI policy heuristic under varying $\widehat{S}$, benchmarked against the discrete-time policy

the discrete-time model with dynamic programming. Due to the large state space of the two-class model, we can solve only problems with fleet sizes of up to $c = 3$ cars within reasonable time.

We analyze two examples. In Example 1, the mean rental time of ADI customers is $\bar{S}_I = \frac{20}{27}$ and is less than the mean rental time of non-ADI customers, $\bar{S}_{II} = \frac{30}{27}$. In Example 2, $\bar{S}_I = \frac{30}{23}$ and is greater than $\bar{S}_{II} = \frac{20}{23}$. In both examples, the overall mean rental time is $\bar{S} = 1$ day. All other parameters are the same for both examples: $\lambda = 3.1$ cars/day, $v = \lambda_I/(\lambda_I + \lambda_{II}) = 0.3$, $c = 3$ cars, $r_I = 100$ Euro/rental, $= r_{II} = 100$ Euro/rental, and $p = 500$ Euro/rental.

The numerical results are shown in Figure 10. The dotted line is the solution of the discrete-time model and serves as a benchmark. For Example 1 ($S_I < S_{II}$), the heuristic performs best with $\hat{S} = 0.9$, resulting in expected daily profits of 149.6 Euro, which are 1.0% below the benchmark of 151.1 Euro. For Example 2 ($S_I > S_{II}$), it performs best with $\hat{S} = 1.0$, resulting in expected daily profits of 112.4 Euro, which are 0.8% below the benchmark of 113.1 Euro.

Leal-Duarte, A., C. Johnson. 2002. On the minimum number of distinct eigenvalues for a symmetric matrix whose graph is a given tree. Mathematical Inequalities and Applications 5 175-180.