The steady-state mobile personnel booking problem

Edison Avraham, Tal Raviv

Department of Industrial Engineering, Tel Aviv University

October 2021

Abstract

We present a model for policy optimization of the online booking of mobile personnel over a multiday horizon with a different cutoff for each day, where the goal is to maximize the expected ratio of accepted requests at steady-state. This model fits the practice of many service providers who allow booking of time slots over a horizon of multiple days and use availability control of the demand. Since the planning horizon is indefinite and the service horizon of each day overlaps the horizon of subsequent days, the objective is defined in terms of the steady-state performance. The interactions with the customers are performed in a single step: The system offers an assortment of time slots covering the next few days, and the user either chooses one of them or abandons the system. Upon the arrival of a service request, the provider estimates the opportunity cost of serving the request at each of the available time slots. We model this cost as a linear function and a Cobb–Douglas function of features that concisely represent the current system state. The assortment of time slots following each service request is constructed by maximizing the expected net gain from the assortment. The parameters of the opportunity cost functions are fitted using a simulation framework. The proposed method is benchmarked based on randomly generated datasets in various demand scenarios and geographies. The method is shown to outperform more straightforward baseline policies significantly.

Keywords: Field Service Management, Multiday Booking, Stochastic dynamic optimization, assortment optimization, Reinforcement learning.

1 Introduction

Many companies provide services at their customers' locations, e.g., maintenance and repair of home appliances, grocery delivery, and medical care. Customers call the contact center of the providers or log in to their websites to book a service. In a typical situation, each customer is offered a set of possible time slots in which the service can be provided, and the customer may choose one or abandon the system. The actual scheduling and routing of the service personnel are determined before the beginning of each working day based on the service requests that should be served on that day. Requests may be booked a few days in advance, and therefore, the demand for time slots on different days is interrelated. That is, service on a particular day is a substitute for service on other days. We use the terms *period* and *day* interchangeably.

The first stage of the service provision process involves an interaction between the service provider and customer regarding the time slot in which the service should begin. For example, if the time slot

agreed upon is 8:00-10:00 tomorrow, the service personnel would arrive at the customer's location at any time during that slot. The provider should take into account spatial and temporal considerations when agreeing to certain time slots, while the customer has her considerations that are unknown to the provider. In this work, we refer to this interaction as the *booking process*. From the provider's side, this process is, in many cases, fully automated.

Attended home service booking problems are characterized by the arrival of requests at different points in time according to some stochastic process and during a *booking horizon*. Each request is accepted or rejected *online* upon arrival, following a short interaction between the customer and the service provider. The booking horizon ends at a certain *cutoff* time that precedes the *service horizon*, during which the requests are served. Deciding whether to accept or reject a request may be done based solely on previously accepted requests, that is, *myopically*. In a more farsighted approach, the decision process may incorporate some information related to future requests; this framework is referred to as *anticipatory*.

This paper is the first to study a multiday booking problem in which requests are accepted several working days in advance, where each day has its own cutoff time, and the customers may book slots in one of the following days. The service horizon is rolled forward by one day at the beginning of each working day, and requests that were already booked are kept in the plan but "shifted" one day backward. In contrast to the single cutoff case, there are inherent interactions between the various days in the multiday setting because requests that were accepted in previous days may affect the availability of future ones. These interactions require optimizing the steady-state performance of the system rather than optimizing each working day separately. This paper introduces an effective policy optimization method and performs a numerical experiment that demonstrates its merits. Moreover, we demonstrate that the additional flexibility obtained from the multiday setting results in a more efficient operation of the service crew.

We define the steady-state mobile personnel booking problem (SSMPBP) as follows: the provider is the leader in the booking process. That is, she offers several possible time slots to the customer. The offered time slots span several subsequent working days. The customer decides whether to accept one of the offered slots and place a service request or to abandon the system altogether. If the customer selects one of the offered slots, the provider is committed to serving her within the slot. To meet this commitment, the provider can offer only slots that allow serving the current customer without violating her commitments to previously booked customers. We refer to these slots as *feasible slots*.

The customers have a predefined preference over the slots. This preference consists of two components: a component common to all the customers, which can be estimated by the provider based on historical data, and an individual component modeled as a random variable with zero mean. We use a discrete choice model to describe the customer decision. The multinomial logit (MNL) model is used for demonstration in our numerical study. Similarly, the service time required by all customers is assumed to be identical. The customer uniformity assumptions can be justified because, during the booking process, the provider possesses relatively little information about the customer and the required service.

A solution to the SSMPBP is a *policy* by which the provider constructs an assortment of time slots upon the arrival of each request. The service provider's objective is to maximize the ratio of accepted requests in the long run. We note that even though different customers may represent different values for the provider, this value may be unknown at the booking's time.

An alternative modeling approach could be solving the problem under a relatively long finite horizon (say 50 days). In such a case, the optimal policy is affected by the day within the horizon. E.g., the optimal policy for the 50th day is different than the one of the 25th day. For a business working under similar conditions for an extended horizon, the policy of days in the "middle" is a good approximation for the optimal policy. However, in this paper, we adopted a different approach in which the conditions under which the system is working are assumed stationary over time, and thus, a simpler stationary policy is searched. In this section, we present the notation needed for the formulation of our mathematical model.

During the booking process, the service provider should consider the following trade-off. By offering many possible time slots, the provider can increase the customer's probability of accepting one of them but at the cost of a greater chance that the customer selects a slot that requires a longer travel time. Consequently, the availability of the service provider for future customer requests is reduced. Note that we do not explicitly include the travel time of mobile personnel in the objective function. However, a policy that maximizes the number of served customers within the predefined time slots induces efficient routes so that the mobile personnel spend a shorter time on roads and have enough time to serve many customers.

We use the term technician to describe mobile personnel (or employee) that visits customers and provides the service. However, the SSMPBP arises in many practical situations where service requests can be placed a few days in advance, for example, in attended home delivery. The provider cannot control the demand by dynamically setting the service fees but can control the availability of the time slots for service. In Figure 1, we present the web page used to book an attended home delivery service provided by a large e-grocer in the UK. Note that the service fees are fixed regardless of the selected slot, and the customer may select a slot in one of several working days.

Our solution method is based on the following idea. For each time slot, we estimate the *opportunity cost* of serving the current request in terms of the expected number of future requests that will be lost if the current one is accepted. This estimate is obtained by fitting the parameters of an approximated value function of the following arguments: (1) the remaining idle time in the slot of the technician after the customer request is inserted in the slot; (2) the remaining idle time in the daily route of the technician after the customer request is inserted in the slot; and (3) the additional travel time required to serve the request. We consider two families of value functions, namely linear and Cobb-Douglas. While linear function is commonly used in the literature, the advantage of using a Cobb–Douglas function to represent the opportunity cost lies in the fact that it allows capturing the interactions between the above arguments.

When a customer accepts an offered slot, the expected net gain is one minus its estimated opportunity cost. The assortment of slots offered by the provider is a subset of the feasible slots, which maximizes the expected net gain based on the customers' discrete choice model. We use a simulation to

estimate the system performance with the estimated opportunity cost function and iteratively fit its parameters.

	SUN 26	MON 27	TUE 28	WED 29	THU 30	FRI 01	SAT 02
08:00 - 09:00	Unavailable Collect?	Unavailable Collect?	Unavailable Collect?	Unavailable Collect?	Unavailable Collect?	Unavailable Collect?	Unavailable Collect?
09:00 - 10:00	Unavailable Collect?	Unavailable Collect?	Unavailable Collect?	Unavailable Collect?	£ 4.50	Unavailable Collect?	Unavailable Collect?
10:00 - 11:00	Unavailable Collect?	Unavailable Collect?	Unavailable Collect?	£ 4.50	Unavailable Collect?	Unavailable Collect?	Unavailable Collect?
11:00 - 12:00	Unavailable Collect?	Unavailable Collect?	£ 4.50	£ 4.50	Unavailable Collect?	Unavailable Collect?	Unavailable Collect?
12:00 - 13:00	Unavailable Collect?	Unavailable Collect?	Unavailable Collect?	Unavailable Collect?	Unavailable Collect?	Unavailable Collect?	Unavailable Collect?
13:00 - 14:00	Unavailable Collect?	Unavailable Collect?	Unavailable Collect?	Unavailable Collect?	Unavailable Collect?	Unavailable Collect?	Unavailable Collect?
14:00 - 15:00	Unavailable Collect?	Unavailable Collect?	£ 4.50	£ 4.50	£ 4.50	Unavailable Collect?	Unavailable Collect?
15:00 - 16:00	Unavailable Collect?	Unavailable Collect?	£ 4.50	£ 4.50	£ 4.50	Unavailable Collect?	£ 4.50

Figure 1: Booking page of a large e-grocer in the UK

The optimal assortment can be a strict subset of the feasible slots since there is an intricate interplay between service quality and efficiency. The fact that the provider withholds some of the feasible slots from customers may be perceived as lowering the service level. However, in the long run, it may give the provider additional flexibility that allows her to improve the service level by rejecting fewer customers. When bookings are allowed for several days in advance, postponing a service to one of the following days can further increase the acceptance rate of requests. Finally, we observe that the Cobb–Douglas function can approximate the opportunity cost somewhat better than a linear function used in previous studies. These insights are well demonstrated in our numerical study.

Section 2 provides a review of the state of the art in the domain of the online booking of mobile personnel and highlights the unique contribution of this study. In Section 3, we formulate the problem. Section 4 describes the proposed booking policy optimization method. Section 5 presents a numerical experiment that compares the performances of the optimized policy to other policies and explores the advantages of the multiday compared to single cutoff settings. Section 6 presents the final remarks and offers directions for future research.

2 Literature review

This paper is the first to study a booking problem in a multiperiod context. However, there are closely related previous studies on single-period booking and multiperiod dynamic routing. In this section, we

review the existing booking literature and compare each of the papers relevant to the current study along the following seven dimensions:

1. The *leader* in the interaction between the customer and the provider. Some studies assume that the customer requests one or more particular time slots, and then the provider either selects one of them or rejects the request. Other studies assume that the provider presents options to the customer, and then the customer selects one of them or abandons.

2. The characteristics of the process that governs customer selection when faced with several choices (e.g., several possible time slots) is the *user choice model*. The papers reviewed below differ by their assumptions regarding this model and the service provider's knowledge about it.

3. The *demand control method* is how the provider manipulates the requests. For the case where the provider is the leader in the negotiation processes, the literature presents two methods, either by *pricing*, i.e., setting a price for each slot, or by *availability control*, i.e., offering only a subset of the feasible slots. If the customer is the leader, the provider controls the demand by selecting one of the options offered by the customers or rejecting the request. For reviews of choice models and demand control in booking systems, see Strauss et al. (2018) and Klein et al. (2020).

4. The decision made by the provider at the time of booking may be based on a *tentative detailed routing* plan or by more abstract considerations that are based on the capacity at each service zone.

5. In practice, service providers accept booking requests for several periods (typically days) in advance, and their objective is to optimize performance in the long run. Most of the existing literature focuses on a single period. Few papers look at a small finite set of service days. This study is the first to optimize the steady-state performance of the system in a setting where bookings are allowed for several periods in parallel. We note that in the single-period case, optimizing each period separately results in optimizing the steady-state performances of the system, but this is not true for the multiperiod case. Hence, we compare the previous studies in terms of their approach to *periodicity*.

6. The existing literature can be classified by the *farsightedness* of the models presented. Some models do not use information about future requests, while others make use of forecasted information.

7. A variety of solution methods have been presented to solve booking problems. The *methodology* is highly related to the modeling assumptions, yet some families of solution methods are relevant for various modeling assumptions. E.g., approximate dynamic programming (ADP) with value function approximation (VFA) is commonly used for various demand control methods.

The main differences between the current study and the previous papers on booking problems are summarized in Table 1. Next, we discuss the main contributions in this stream. Two early studies model the booking problem with the customer as a leader. Campbell and Savelsbergh (2005) study a problem where the customers offer several possible slots, and the provider chooses one. Ehmke and Campbell (2014) study an online booking problem where the customer offers time slots sequentially, and the provider can accept or reject each offer. The objective in both studies is to maximize the accepted number of requests .

Campbell and Savelsbergh (2006) are the first to study an online booking problem where the provider presents a different price (or discount) dynamically for each slot with the goal of maximizing

the expected periodic revenue. Tentative routes are determined by a GRASP heuristic, and the pricing problem is solved myopically, i.e., without considering future requests.

Köhler et al. (2020) study an online booking problem with availability control where the service provider can present a customer with a set of feasible time slots with two possible lengths. A specific (axiomatic) user choice model that favors shorter time slots is assumed. The objective is to maximize the number of accepted customers. A heuristic decision criterion is applied for constructing an assortment of slots for each customer.

Asdemir et al. (2009) study an online booking problem where each request is associated with a certain revenue and capacity consumption. The provider presents the customer with a set of time slots that do not violate the capacity constraint and determines a price for each presented slot. The customer chooses one of the proposed time slots or abandons the process according to an MNL discrete choice model. The utility in the MNL model is determined by the presented price and each slot's particular characteristic. The authors devised a policy that determines the prices to maximize the expected total profit of the provider. The problem is formulated as a Markov decision process (MDP) that identifies the opportunity cost of serving a request in each slot, i.e., the future profit obtained if the newly arrived request is not served minus the future profit obtained if it is .

Yang et al. (2016) study a problem similar to that of Asdemir et al. (2009) but in a more general setting, where each request is associated with a location and the tentative routes are constructed using a GRASP heuristic inspired by Campbell and Savelsbergh (2006). A heuristic method is applied to estimate the opportunity cost of each arriving request. These estimations are used to optimize the prices. Yang and Strauss (2017) are the first to apply an offline ADP framework using linear VFA to estimate the value and calculate the prices. The routing costs are approximated heuristically without constructing tentative routes. Koch and Klein (2020) study a similar problem. Dummy customers, representing future expected demand, are initially included in the routing plan for a more accurate estimation of the opportunity cost. An ADP is used to set the prices of the slots .

Strauss et al. (2020) study an online booking problem where customers can express their possible degree of flexibility regarding the desired service time. The customer choice model is the nested MNL. Upon the arrival of a request, a customer is presented with a list of bundles of time slots. Each bundle is associated with a fee. The customer may choose one of the offered bundles or abandon the system. The service will be provided at one of the slots in the selected bundle at the firm's discretion and without a commitment to the customer in advance. The problem is to dynamically price the bundles with the objective of maximizing the expected gross profit plus the fees minus the routing costs .

Mackert (2019) studies an e-grocer delivery booking problem with availability control. The customers are offered a list of delivery slots to choose from with the goal of maximizing the retailer profit net of the delivery cost by selecting the best assortment of time slots. The decisions made by customers (choose one of the slots or abandon the system) are modeled by the generalized attraction model (GAM) proposed by Gallego et al. (2015). The author advocates the GAM instead of its more commonly used special case MNL since it allows avoiding overestimating the demand by reflecting customers' dissatisfaction with small assortments of offered slots. The approximation of the opportunity cost of each request is inspired by a method presented by Klein et al. (2018). A linear program is solved

to find an optimal assortment of time slots to offer. Mackert (2019) studies a model similar to ours and differs only in the periodicity and solution method .

Ulmer and Thomas (2019) study an online booking problem in which each delivery request is associated with the potential revenue, location, and vehicle capacity consumption. A request is either served at any time in the service horizon or rejected. The objective function is to maximize the sum of the revenues. Routing is done using an insertion heuristic .

With respect to the literature on booking problems, we extend the state of the art mainly in terms of the periodicity and the steady-state performance objective. Indeed, the multi-period service horizon with multiple cutoff times results in interactions between the various periods and is thus more suitable for steady-state rather than transient analysis. Following previous studies, we use a general user choice model where the service provider is the leader in the negotiation process. The demand is manipulated using availability control, and the decisions about the offered assortments incorporate farsighted considerations by estimating the opportunity cost of each offered slot. Moreover, we extend previous studies by using a nonlinear function for approximating the opportunity cost and explore its merits compared to the commonly used linear function. In Table 1, we summarize the booking literature discussed above along the seven dimensions that we identified.

Table 1: Summary of the booking literature

Paper	Leader	Choice Model	Demand control	Routing	Period-icity	Farsig-htedness	Solution method
Campbell and Savelsbergh (2005)	Customer	Fixed and known choices	Availability	Solved with GRASP	No	Yes	Heuristic criteria
Ehmke and Campbell (2014)	Customer	Fixed random choices	Availability	Adaptation of Solomon's (1987) heuristic	No	No	Adapted from Campbell and Savelsbergh (2005)
Campbell and Savelsbergh (2006)	Provider	A linear function that maps the slot prices to selection probabilities	Pricing	Solved with GRAPS	No	No	Myopically optimal pricing is calculated
Köhler et al. (2020)	Provider	Axiomatic, defined by this paper	Availability control	Insertion heuristic	No	Yes	Heuristic decision criteria
Asdemir et al. (2009)	Provider	MNL	Pricing	None (capacity allocation only)	No	Yes	MDP
Yang et al. (2016)	Provider	MNL	Pricing	GRASP inspired by Campbell and Savelsbergh (2006)	No	Yes	Opportunity cost is approximated heuristically. Pricing is solved using mathematical programming.
Yang and Strauss (2017)	Provider	MNL	Pricing	No detailed routing. The routing cost is approximated.	No	Yes	ADP
Koch and Klein (2020)	Provider	Reservation price model (a general-ization of MNL)	Pricing	Insertion heuristic	No	Yes	ADP

Strauss et al. (2020)	Provider	Nested MNL	Pricing	No detailed routing. The routing cost is approximated.	No	Yes	Opportunity cost is approximated based on forecasted future requests using a linear program.
Mackert (2019)	Provider	GAM	Availability control	Insertion heuristic	No	Yes	Opportunity cost is approximated using a mixed integer linear program and the assortment problem is solved using a linear program.
Ulmer and Thomas (2019)	Customer	No choice. A request is either accepted or rejected	Availability control	Insertion heuristic	No	Yes	ADP
This paper	Provider	Generic (MNL as an example)	Availability control	Insertion heuristic	Yes	Yes	Value function approximation using reinforcement learning. Assortment problem is solver using math programming

The much more developed literature on the dynamic vehicle routing problem (VRP) with stochastic requests shares some properties with the booking literature. In both types of problems, a customer calls a provider to set up a visit of mobile personnel. However, in the dynamic VRP setting, there is no notion of a booking horizon or cutoff times, and typically it is not required to model the customer choice mechanism or to solve a pricing or an assortment problem. Moreover, this stream of literature focuses on a finite horizon and thus does not optimize the steady-state performance as we do in the current paper. For comprehensive surveys on single-day dynamic vehicle routing problems, readers may refer to Pillac et al. (2013) and Psaraftis et al. (2016). For studies on dynamic routing with stochastic requests, see Bent and Van Hentenryck (2004), Ghiani et al. (2009), Ghiani et al. (2012), Soeffker et al. (2017), Ulmer (2019), and Voccia et al. (2019.(

Ulmer (2020) studies a dynamic VRP where delivery requests are served on their arrival day (same-day delivery). The customer can select the delivery deadline from several alternatives with different prices. The provider assumes a willingness-to-pay choice model. The problem is to dynamically price the deliveries to maximize the expected revenue. The problem is formulated as an MDP, and a VFA scheme is used to estimate the opportunity costs of each delivery option. Routing is carried out using an insertion heuristic. The numerical experiments show the superiority of the presented framework over fixed pricing and other benchmark policies. While Ulmer (2020) shares many similarities with the booking literature, it differs from our study in several dimensions mentioned above. In particular, it optimizes the revenue during a finite planning horizon rather than the steady-state performance of the system.

In most of the literature on the multiday dynamic VRP, customers place requests, and the provider serves them at some time during the service horizon directly or by outsourcing at a higher cost. The objective is typically to minimize the cost or maximize the requests served on their arrival day within a finite planning horizon. For some papers in this stream, see Angelelli et al. (2009), Wen et al. (2010),

Albareda-Sambola et al. (2014), and Ulmer et al. (2018b). Pérez Rivera and Mes (2017) study a multi-period routing problem in the domain of freight transportation. Chen et al. (2017) solve a multiperiod technician scheduling problem where the service time of each technician is reduced as a result of increasing previous experience with similar tasks.

Finally, several authors have studied tactical planning in the context of online booking, i.e., the offline process of determining service zones, the allocation of personnel to zones, and service time slots. See, for example, Agatz et al. (2011), Cleophas and Ehmke (2014), Hernandez et al. (2017), Bühler et al. (2016), Restrepo et al. (2019), and Klein et al. (2019).

3 Problem definition

The SSMPBP is stated as follows: service requests arrive according to a nonhomogeneous periodic Poisson process. Each request is characterized by its location. The service time of all the requests is assumed to be identical and denoted by τ. The horizon is divided into working days, and each day is divided into several non-overlapping service windows (e.g., 8–10, 10–12, 12–14, 14–16, 16–18). A combination of a service day and a window is referred to as a *slot*. Upon the arrival of a request on day h, it can be either rejected or scheduled for service in a slot on one of the following days $\{h + 1, h + 2, .., h + H^p\}$; this set of days is referred to as the *service horizon*. From the description above, it follows that a cutoff time associated with day $h + 1$ is the end of day h.

The accepted requests are served by a crew of M homogenous mobile personnel (technicians) available on each working day. The technicians depart from a depot, not before a predefined time, and return to the depot at most L time units later. The travel time between any two locations j and k in the service area, t_{jk}, is known and deterministic.

Each newly arrived service request is initiated by a customer and triggers the execution of the following protocol.

- The service provider's system displays an assortment of time slots in the service horizon for the service's possible starting times.
- The customer may select one of the proposed time slots or abort the process without booking his request. The customer's decision is modeled using a discrete choice model, and the parameters of this model are assumed to be known to the service provider.
- Once the customer selects a slot, the service request is booked in the system and inserted into the selected time slot's tentative schedule.

The assortment is a subset of the *feasible* slots in which the service of a customer may begin. A slot is considered feasible if assigning the request to that slot causes no late arrivals at the locations of the previously accepted service requests and allows the technicians to return to the depot by the end of the working day. The assortment may be empty, in which case the request is rejected immediately.

It should be noted that the actual schedules and routes of the technicians for each working day can be reoptimized after the cutoff time by solving a static vehicle routing problem with time windows (VRP-TW). Moreover, in a typical situation, there are several hours between the cutoff time (e.g., at midnight) and the beginning of the working day (e.g., at 7 am). During these hours, the system can exert additional computational resources to achieve a plan with a shorter total travel time compared to the

9

original tentative plan. We note that there is a vast literature on static vehicle routing problems with time windows, which is out of the scope of this study. We refer the reader to a recent survey on this topic (Braekers et al., 2016) and a paper on stochastic routing and scheduling with time slots (Avraham and Raviv, 2020).

A solution for the SSMPBP is a policy for creating an assortment for each customer request online with the objective of maximizing the expected acceptance rate of requests in the steady-state.

The description of a discrete choice model is included in the input of the SSMPBP. We assume that all the customers follow the same choice model. In the numerical section of this paper, we use an MNL model as an example; however, any discrete choice model could be used. An advantage of the MNL model is that it is possible to estimate the model's parameters from historical data efficiently; see Yang et al. (2016).

We formulate the problem of maximizing the expected number of accepted requests per day in the steady-state as an MDP. To this end, we first define notations of the possible states, actions, probabilities, and other parameters in Section 3.1. Next, in Section 3.2, we define the transition probability as a function of the current state and the action taken. In Section 3.3, we define the objective function and the solution as a policy optimization problem

3.1 Notation

In this section, we present the notation needed for the formulation of our mathematical model.

$\mathcal{A}$ The set of time slots in the service horizon. Each subset of slots (an assortment) can be a feasible action in a particular state.

w_0 The default alternative customer choice of abandoning the system.

$\mathcal{L}$ Set of locations of all potential customers. For simplicity of presentation, we assume here that this is a discrete set.

T The number of time intervals during a single day. Each interval is short enough to neglect the probability that the number of requests arriving during the interval is greater than 1. The intervals are indexed as $t = 1, \dots, T$.

λ_t The arrival probability of a request during interval t.

p_{lt} The probability that a request from location $l \in \mathcal{L}$ arrives at time interval $t \in \{1, \dots, T\}$. Note that $\sum_{l \in \mathcal{L}} p_{lt} = \lambda_t$.

$\mathcal{P}(A, w)$ The probability that a customer will choose time slot $w \in \mathcal{A}$ when presented with an assortment of time slots $A \subseteq \mathcal{A}$. Recall that the alternative of abandoning the system is denoted as w_0. Therefore, $\mathcal{P}(A, w_0)$ denotes the probability of not selecting any slot when the assortment A is presented. We use the convention that $\mathcal{P}(\emptyset, w_0) = 1$. Note that $\mathcal{P}(A, w) = 0$ if $w \notin A$.

S_{ht} The state of the system at time interval t on working day h. A state includes information about the number of accepted customer requests at each time slot in the service horizon, $w \in \mathcal{A}$, at each location, $l \in \mathcal{L}$, and at the current time of the day t. The state can be described as a tuple of the current time and a sparse matrix with one column for each possible location and one row for each time slot. The entries of the matrix are the number of booked customers at the corresponding slot and location.

$\mathcal{S}$	The set of all feasible states.
$\pi(S, l)$	A policy π is a mapping of $\pi: (S, l) \to 2^{\mathcal{A}}$. The policy prescribes a set of feasible slots to offer to a customer from location l that places a service request while the system's state is S. $2^{\mathcal{A}}$ denotes the collection of all subsets of time slots in the service horizon. We use the convention that $\pi(S, l_0) = \emptyset$ to refer to situations where no request arrives.
Π	The set of all policies.
$N(S_{ht})$	The set of states that can be reached from the current state in a one-time interval resulting from booking at most one customer request. For the last interval in a day, $N(S_{hT})$ consists of states $S_{h+1,1}$. Note that in the transition from time interval T to time interval 1, the cutoff time is crossed. After the cutoff time, the bookings of the first day of the previous service horizon are removed, the rest of the bookings are moved to one day earlier, and new empty slots for the last day of the new service horizon are added.
$I(S_{ht})$	For $t < T$, the state at time interval $t+1$ if no request was accepted at time t. For the last interval $I(S_{hT})$ is the state $S_{h+1,1}$ if no request was accepted at time T or a request was accepted for the first day of the service horizon.
$\Delta(S_{ht}, S_{h,t+1})$	For $S_{h,t+1} \in N(S_{ht})$ and $t < T$, the location of the arriving request at time t. For $S_{h+1,1} \in N(S_{hT})$, the location of the arriving request at time T.
$w(S_{ht}, S_{h,t+1})$	For $S_{h,t+1} \in N(S_{h,t})$ and $t < T$, the slot of the arriving request at time t. For $S_{h+1,1} \in N(S_{h,T})$, the slot of the arriving request at time T.
W_1	The set of slots on the first day of the service horizon.

3.2 Transition probability function

Using the notation presented in section 3.1, we can define the transition probability between any pair of states in consecutive time intervals given a request at location l for $t = 1, \dots, T-1$, as follows.

$$P(S_{ht}, S_{h,t+1}) = \begin{cases} 1 - \lambda_t + \sum_{l \in \mathcal{L}} p_{lt} \mathcal{P}(\pi(S_{ht}, l), w_0), & S_{h,t+1} = I(S_{ht}) \\ p_{\Delta(S_{ht}, S_{h,t+1}), t} \mathcal{P}(\pi(S_{ht}, \Delta(S_{ht}, S_{h,t+1})), w(S_{ht}, S_{h,t+1})), & S_{h,t+1} \in N(S_{ht}) \\ 0, & otherwise \end{cases} \quad (1)$$

The first case refers to transitions resulting from progressing from time interval t to $t+1$ without accepting any request. Note that $S_{h,t+1} = I(S_{ht})$ indicates that the only difference between the states S_{ht} and $S_{h,t+1}$ is the time. The probability of this transition is the sum of two probabilities of two mutually exclusive cases. Either no customer request arrives during interval t (with a probability of $1 - \lambda_t$) or a request arrives, an assortment $\pi(S_{ht}, l)$ is offered by the operator and the customer rejects it. The second case refers to transitions resulting from progressing from time interval t to $t+1$ where a request is accepted. The location of the accepted request is encoded in the information regarding the states S_{ht} and $S_{h,t+1}$ as $\Delta(S_{ht}, S_{h,t+1})$ and the slot as $w(S_{ht}, S_{h,t+1})$. The third case refers to all the other transitions, which are infeasible in a one-time interval. For time interval T, just before the cutoff time, the transition probabilities are

$$P(S_{hT}, S_{h+1,1}) = \begin{cases} 1 - \lambda_T + \sum\limits_{l \in \mathcal{L}} p_{lT} \sum\limits_{w \in W_1 \cup \{w_0\}} \mathcal{P}(\pi(S_{hT}, l), w), & S_{h+1,1} = I(S_{hT}) \\ p_{\Delta(S_{hT}, S_{h1}), t} \mathcal{P}(\pi(S_{ht}, \Delta(S_{hT}, S_{h1})), w(S_{hT}, S_{h1})), & S_{h+1,1} \in N(S_{hT}) \\ 0, & otherwise \end{cases} \tag{2}$$

The first case refers to transitions from time interval T to 1 without accepting any request or when the accepted request is for the first day of the service horizon and hence is deleted from the state at the cutoff time at the end of time interval T. The first term is as shown in (1), and the second term is a sum over the events where the customer either chooses to abandon the system or chooses a slot on the first day of the service horizon. The second and third cases are the same as in (1), but note the unique definitions of $N(S_{hT}), \Delta(S_{hT}, S_{h+1,1})$, and $w(S_{hT}, S_{h+1,1})$.

3.3 Policy optimization

When a policy π is known, the transition probability above defines a Markov chain over the states $S \in \mathcal{S}$. We denote the steady-state probability of this chain at time interval t by $Q_t(S, \pi)$. An optimal solution is a policy that satisfies

$$\pi^* = \underset{\pi \in \Pi}{\operatorname{argmax}} \sum_{t=1}^{T} \sum_{S \in \mathcal{S}} Q_t(S, \pi) \cdot \sum_{l \in \mathcal{L}} p_{lt} \sum_{w \in \pi(S, l)} \mathcal{P}(\pi(S, l), w)$$

Recall that we are interested in the steady-state performance of the system when the demand is assumed to be stationary. Therefore, it is enough to consider the expected number of accepted requests in one typical day.

The expected number of accepted requests per day in the steady-state is calculated as the sum of the probabilities of accepting a request in each time interval. For each particular state S, the probability of accepting a request in the corresponding interval is calculated by the sum of the probability of the arrival events multiplied by the probability that one of the offered slots will be accepted.

Unlike previous authors in the booking literature, we used a general MDP rather than a sequential set of Bellman equations to define our problem formally. The MDP framework is suitable because we are interested in the expected number of accepted requests in a steady-state rather than the number of accepted requests within a finite booking horizon.

In practice, customers differ in their characteristics, such as expected service time and choice model. If the service provider is aware of the particular characteristics of the customer when a request is booked, then a more detailed model can be used to achieve better service. To this end, the state space may include additional properties such as the customer type or the arrival time of the request. The method for approximating the value function presented in the next section can be adapted for such a detailed state description, but this would require fitting a higher-dimensional function that may be much more difficult to estimate. In this paper, we focus on the case where all the customers follow the same choice model.

4 Methodology

A direct solution of the MDP presented in Section 3 is not practical for realistic settings since the set of states is too large. Instead, we encode each state and the information regarding each new request in a

compact way and use simulation results to fit a function that approximates the opportunity cost of accepting a request from each possible location in each time slot in the service horizon. Upon the arrival of a service request, an assortment of feasible slots that approximately maximizes the expected gain from the customer is constructed and offered to him. This assortment is obtained by selecting one that maximizes the immediate (unit) gain minus the approximated *opportunity cost*. At any particular moment, the booking system maintains a tentative set of routes and schedules for all technicians within the entire service horizon. Once a customer selects a time slot, his request is inserted into one of the tentative routes, and the schedule is updated. The tentative plan and the insertion procedure enable us to quickly evaluate the feasibility of inserting the request in each slot.

Our solution method can be classified as reinforcement learning based VFA (see Powell (2019)). The opportunity cost is the expected number of customers that will be lost from consistently allocating a particular slot to a customer at a given state of the system in terms of the expected steady-state performance. In previous studies that analyzed finite horizon booking problems, the opportunity cost of a slot was regarded simply as the difference between the value of the states of the system with and without the booked customer; see, for example, Mackert (2019), Koch and Klein (2020), and Ulmer (2020). In our settings, the relation between the opportunity cost and the value of the policy is more intricate since the values of the states represent the daily rate of accepted customers. The opportunity cost can be viewed as the difference in the steady-state acceptance rate if the slot would be allocated to any identical customer that will arrive in the same state.

The idea of estimating the opportunity cost is well established in the booking literature, as discussed in Section 2. However, we present the first application of this idea for optimizing the long-term performances of a booking process rather than focusing on a single period with a single cutoff.

Section 4.1 presents the approximated opportunity cost function and explains the modeling considerations. Section 4.2 presents the routing and scheduling oracle that we implement to aid in the cost estimation for each request in each slot. Section 4.3 presents the optimization model used to construct the list of offered slots. Finally, Section 4.4 presents the proposed procedure for fitting the parameters of the opportunity cost function.

4.1 The approximated opportunity cost function

The approximated opportunity cost function represents the implications of accepting a particular customer request at a particular slot. The function estimates the number of requests that will either be rejected or abandon the system due to the resources being allocated to the current one.

Consider a route r' that represents the set of requests that a technician visits on a given day and the starting times at each customer location, and the departure and arrival times at the depot. Let i be a newly arrived request that is a candidate to be added to route r' in time slot s. A new route r, consisting of all the requests in r' and request i, is feasible if the slots of all the previous requests in r' are not changed, and the service of request i starts in time slot s. Based on (r, r', i, s), one can calculate the following three measures:

$RTS(r,s)$ Remaining *idle* time in slot s (after the insertion of i into the route in slot s and assuming that the technician departs from the previous customer as early as possible). The technician is considered idle when no service or travel is carried out.

$RTR(r,s)$ Remaining *idle* time for route r during the entire working day (after the insertion of i into route r in slot s).

$TT(r',r)$ Increase in the travel time incurred by extending r' to r.

The calculation of RTS, RTR, and TT is discussed in detail in Section 4.2. Note that the values of the above three measures are nonnegative. Next,

$$f(r,r',s) = \begin{cases} RTS(r,s)^{\alpha} RTR(r,s)^{\beta} TT(r',r)^{\gamma}, & if \ route \ r \ is \ feasible \\ \infty, & otherwise \end{cases} \quad (3)$$

where α, β, γ are parameters that characterize the opportunity cost function f and may be positive or negative. The Cobb–Douglas function $f(r,r',s)$ represents the expected number of future requests that will be lost if the system accepts the current request in route r in slot s. Therefore, $f(r,r',s) > 1$ implies that serving the current request in route r in slot s is not worthwhile. Economists often use the Cobb–Douglas function to represent production and utility with respect to the inputs. We use it to approximate the opportunity cost because it is one of the simplest functional forms that can capture various interactions between inputs.

The measures RTS, RTR, and TT concisely capture sufficient information about the system state to estimate the opportunity cost effectively. The case for TT is straightforward: a higher marginal increase in the travel time corresponds to a higher likelihood that future requests will be lost. Thus, γ is likely to be positive. Next, high values of RTS and RTR indicate that the candidate route (and in the case of RTS, the candidate slot as well) is sufficiently vacant to serve additional requests after accepting request i. Thus, accepting i is unlikely to adversely affect the offers made in response to future customer requests. Hence, α and β are likely to be negative. We note that Ulmer et al. (2018a) and Koch and Klein (2020) use similar measures to aggregate the state space in their ADP frameworks. In Section 4.4, we present a method for fitting appropriate values of α, β, γ.

Note that RTS, RTR, and TT may be zero. In this case, since the exponent may be negative, to ensure that the Cobb–Douglas cost function is defined, we correct (3) by adding a small $\varepsilon > 0$ to each measure. Specifically,

$$f(r,r',s) = \begin{cases} (RTS(r,s) + \varepsilon)^{\alpha}(RTR(r,s) + \varepsilon)^{\beta}(TT(r',r) + \varepsilon)^{\gamma}, & if \ route \ r \ is \ feasible \\ \infty, & otherwise \end{cases} \quad (3')$$

Previous authors in the booking literature used a linear function to approximate the opportunity cost based on similar measures. For the sake of compression, we also introduce here a linear approximation cost function (with different coefficients) in the form

$$f(r,r',s) = \begin{cases} \max(0, \alpha RTS(r,s) + \beta RTR(r,s) + \gamma TT(r',r)), & if \ route \ r \ is \ feasible \\ \infty, & otherwise \end{cases}$$

In Section 5, we report the outcomes of using these two functions.

14

4.2 The routing and scheduling oracle

This section presents a procedure to insert new requests into the existing plan at particular routes and time slots. This procedure is used to construct the tentative routes and schedules and approximate the corresponding opportunity costs. These approximate values are later used to construct assortments offered to the customers, as described in Section 4.3. The opportunity cost of a time slot is defined as the opportunity cost of assigning the request to the technician that minimizes it.

Algorithm 1 presents a pseudocode of a program that finds the set of time slots with opportunity costs smaller than 1. For each such time slot, the program returns the technician to which the request should be tentatively assigned, the tentative position of the request along his route, and the approximate opportunity cost of the time slot. It should be noted that the routing and scheduling problem is not the focus of this study and that any other fast routing procedure can replace the proposed oracle.

Input:

    Set of service time slots at the next H^p working days, denoted by S

    Set of technicians, denoted by K

    Tentative routes of all the technicians in the following H^p working days

    Parameters of the opportunity cost function α, β, γ

    Identification of the current request, i, and its properties

for $s \in S$

  $C[s].f \leftarrow \infty$

  for $k \in K$

    $r' \leftarrow$ current route of k in the day of slot s

    for $p\ = 0\ to\ |r'|$ // *where $|r'|$ is the number of customers in route r'*

      $r \leftarrow$ route of k after inserting request i after position p to r'

      if route r is feasible for slot s

        calculate $RTS(r,s), RTR(r,s), TT(r',r),\ f(r,r',s)$

        if $f(r,r',s) < C[s].f$

          $C[s].tech \leftarrow k$

          $C[s].f \leftarrow f(r,r',s)$

          $C[s].route \leftarrow r$

$S' = \{s \in S | C[s].f \leq 1\}$

Return: S', C

Algorithm 1: Pseudocode for finding a set of candidate slots

For each slot s, we define an object $C[s]$ that stores information regarding the best-found route to serve request i in slot s. In particular, $C[s].tech$ is the identity of the most appropriate technician, and $C[s].route$ is the route of this technician, including the newly arrived request i. Recall that a route is defined by a set of requests and their service times. In $C[s].f$, we store the approximated opportunity cost for the particular request and slot. We initialize $C[s].f$ to infinity.

15

Our algorithm updates C by looping through all the slots, technicians, and insertion positions along the routes. Note that inserting a request at position p means that the request is inserted immediately after the p^{th} customer in the sequence, with the convention that insertion at position 0 means inserting at the first position along the route. After the insertion, the starting times of all requests starting from the position of the inserted request are set to the earliest possible times in their pre-committed slots. Recall that the obtained route is feasible if all its requests start at their required time slots.

The technician's route and schedule with the minimal opportunity cost are found for each slot. Finally, a list with all the slots with an expected opportunity cost that is not greater than one, S', and the best-found technicians, routes, and their opportunity costs (stored in C) are returned. A subset of the slots in this list is offered to the customer, as discussed in Section 4.3.

Next, we elaborate on the calculation of RTS, RTR, and TT. To calculate TT, let j and k denote the previous and next customers of i along the new route r, respectively. Thus, k follows j along the old route r'. Now, $TT = t_{ji} + t_{ik} - t_{jk}$.

Let u_j denote the service starting time at the location of request j, and let a_s and b_s denote the starting and ending times of slot s, respectively. The routing and scheduling oracle inserts all the customers at the earliest possible starting time u_j that maintains their time slot constraints. Specifically, if customer k is the immediate successor to customer j and k is assigned to slot s, $u_k = \max(a_s, u_j + \tau + t_{jk})$, where τ is the service time. To calculate the RTS for slot s, let $\hat{j}$ denote the customer whose service starting time is the latest in slot s. That is, $\hat{j} = argmax_{j \in R}\{u_j | a_s \leq u_j \leq b_s\}$, where R is the set of customers along route r. Next, let $\hat{k}$ denote the immediate successor of $\hat{j}$. Hence,

$$RTS = b_s - \min(u_{\hat{j}} + \tau + t_{\hat{j}\hat{k}}, b_s).$$

After serving customer $\hat{j}$, the technician is ready to start the service of $\hat{k}$ at time $u_{\hat{j}} + \tau + t_{\hat{j}\hat{k}}$. If this time is still within slot s, i.e., before b_s, the remaining time, RTS, is $b_s - (u_{\hat{j}} + \tau + t_{\hat{j}\hat{k}})$; otherwise, it is zero.

Finally, the RTR is calculated as the difference between the length of the working day, L, and the total travel and service time allocated to the technician.

4.3 Constructing an assortment

After calculating the approximated opportunity cost of scheduling the request in each slot, the planner can construct an assortment of slots to offer that will approximately maximize the expected number of current and future requests that will be served. Recall that the net gain from serving a customer in a particular slot is one minus the expected number of future requests that will be lost due to the resources allocated for the current request. The problem of selecting a subset of slots from the candidate slot list, S', which is created as described in Section 4.2, to maximize the expected net gain can be formulated as a nonlinear mixed-integer program (4)–(6). For each $s \in S'$, we define a binary decision x_s that indicates whether the slot is included in the assortment and an auxiliary variable P^s that holds the probability that the customer will select slot s.

$$\text{maximize} \sum_{s \in S'} P^s(1 - C[s].f) \tag{4}$$

subject to

$$P^s = g_s(\mathbf{x}) \qquad \forall s \in S' \tag{5}$$

$$x_s \in \{0,1\} \quad \forall s \in S' \tag{6}$$

The objective function maximizes the expected net gain from the offer by weighting the expected gain from each slot with the probability that the customer will select it. Constraint (5) relates the auxiliary variables P^s to the offered slots represented by x. The function $g_s(\mathbf{x})$ represents the discrete choice model and calculates the probability of selecting slot s when the assortment is represented by the characteristic vector **x**. We also comment that

1. For a slot s that is not included in the offer (i.e., $x_s = 0$), $P^s = 0$.
2. The sum $\sum_{s \in S'} P^s \leq 1$. This inequality is strict if the probability of customer abandonment is positive.

If the choice model is MNL, (5) can be replaced with

$$P^s = \frac{e^{u_s} x_s}{\sum_{q \in S'} e^{u_q} x_q + 1} \qquad \forall s \in S' \tag{5'}$$

In the denominator, we add the term 1 to represent the abandonment alternative. Note that the utility to a customer from abandoning the system is 0, and $e^0 = 1$.

In the numerical experiment, the optimization problem (4)–(6) is linearized according to an idea presented by Mackert (2019). A commercial solver solves the resulting linear program in fractions of a second. Note that this linearization is applicable to the GAM, which is a generalization of the MNL model. However, solving the assortment model for other choice models may be NP-hard. In these cases, one may either use enumeration of all the possible assortments or use heuristic methods.

Once the optimal assortment is presented to the customer, he may either select one of the slots or abandon the system. If the customer selects a slot s, the request is inserted into the route of technician $C[s].tech$, as prescribed by Algorithm 1. The resulting plan is tentative, and the actual routes and schedule may be reoptimized later.

4.4 Fitting parameters for the opportunity cost function

In this section, we present a procedure for fitting values for the opportunity cost model parameters. For any combination of the parameters α, β, and γ, the opportunity cost function f defines a policy with a (possibly) different expected number of accepted requests per day. The objective is to determine the model parameters that maximize this number. The procedure described below converges to a local maximum of this optimization problem, but convergence to a global optimum cannot be ensured. In Section 5, we provide numerical evidence of the effectiveness of the proposed heuristic.

The proposed procedure relies on a variant of the steepest descent method. The procedure has three phases that are applied repeatedly: 1) estimation of the steady-state performance of a current solution;

2) estimation of the directional derivatives for the three parameters; and 3) improvement of the current solution by applying a line search in the most promising direction. The process is repeated until a stopping criterion is met.

In Phase 1, we estimate the expected number of accepted requests for a given set of parameters using a discrete event simulation. Specifically, we simulate multiple working days and calculate an estimator for the expected number of accepted requests in the steady-state by ignoring the first few (warmup) days. We note that the optimal policy for our steady-state model remains approximately the same after the warmup period.

We use common random numbers to create the same stream of orders in each iteration during the optimization process. This technique reduces the variability in the simulation results. We run a relatively long single replication in each iteration. The accuracy of the simulation is analyzed and controlled by dividing the long replication into multiple shorter batches. By applying the batch means method, we can extract multiple observations from a single replication without running the warmup period multiple times, see Law (2013), Section 9.5.3. For a further discussion of variance reduction techniques and the statistical analysis of simulation results, see, for example, Ross (2013).

In Phase 2, we estimate the expected number of accepted requests for 26 neighboring combinations of α, β, and γ. In our framework, the neighborhood of a current solution $\boldsymbol{x}^{(n)} = (\alpha^{(n)}, \beta^{(n)}, \gamma^{(n)})$, denoted by $\mathcal{N}$, is obtained as the set of all the solutions generated from increasing, fixing, or decreasing the value of the three parameters by a small nonnegative constant Δ. The neighborhood is denoted as

$$\mathcal{N}(\alpha^{(n)}, \beta^{(n)}, \gamma^{(n)}) = \{\alpha^{(n)}, \alpha^{(n)} + \Delta, \alpha^{(n)} - \Delta\} \times \{\beta^{(n)}, \beta^{(n)} + \Delta, \beta^{(n)} - \Delta\} \times \{\gamma^{(n)}, \gamma^{(n)} + \Delta, \gamma^{(n)} - \Delta\}$$

The expected number of accepted requests of each solution in $\boldsymbol{x} \in \mathcal{N}$ is evaluated by simulation.

In Phase 3, if one or more of the neighboring solutions is better than $\mathbf{x}^{(n)}$, let $\mathbf{x}^{(n')}$ be the best solution in $\mathcal{N}$ and thus $\mathbf{x}^{(n')} - \mathbf{x}^{(n)}$ be the best *improving direction.* Next, we initiate a line search along a ray segment from $\mathbf{x}^{(n)}$ in the direction of $\mathbf{x}^{(n')}$, that is, in

$$\mathcal{M} = \{(\alpha, \beta, \gamma) | \mathbf{x}^{(n)} + l(\mathbf{x}^{(n')} - \mathbf{x}^{(n)}), \ l \in \mathfrak{R}_+\}$$

We use the golden section method to find a locally optimal solution $\boldsymbol{x}^{(n+1)}$ in ray $\mathcal{M}$. Note that since $\boldsymbol{x}^{(n')} \in \mathcal{M}$, the best solution obtained by the line search is at least as good as $\boldsymbol{x}^{(n')}$ and thus strictly better than $\boldsymbol{x}^{(n)}$.

If no improving solutions are found in $\mathcal{N}$, then $\boldsymbol{x}^{(n+1)} \leftarrow \boldsymbol{x}^{(n)}$, and Δ is updated as $\Delta \leftarrow \Delta/\omega$, where $\omega > 1$ represents an update factor. Note that reducing Δ redefines the neighborhood with more similar solutions to the current one. This update can be beneficial in later stages of the search where the current solution is already close to a local optimum.

The search procedure starts from some initial solution $\boldsymbol{x}^{(0)} = (\alpha^{(0)}, \beta^{(0)}, \gamma^{(0)})$ and iterates to generate a series of solutions with nondecreasing values. The value of Δ is reset to its initial value after each improvement. If no improving solution is found for a predefined number of iterations, the search process is terminated.

5 Numerical experiments

This section presents the numerical experiments performed to benchmark the proposed method and to derive some insights. Section 5.1 describes the benchmark problem instances. Section 5.2 presents a performance analysis of the proposed method. In Section 5.3, the importance of each of the model parameters is tested. In Section 5.4, we demonstrate the merits of a long service horizon with multiple cutoff times. Finally, in section 5.5, we study and compare the two opportunity cost approximation functions – Cobb-Douglas and linear.

5.1 Description of problem instances

This section describes the settings of the benchmark instances used in the main experiment reported in Section 5.2. Some of the parameter values are diversified in Section 5.4 to gain additional insights.

The service crew in our benchmark instances consists of $M = 6$ technicians that are available on each working day. The working day begins at 8:00 and ends at 18:00; thus, $L = 600$ minutes. Each working day is divided into five non-overlapping service windows of two hours each, i.e., 8:00 – 10:00, 10:00 – 12:00, 12:00 – 14:00, 14:00 – 16:00, and 16:00 – 18:00. All the slots of the same working day have the same cutoff time, which is the beginning of the corresponding working day.

The length of the service horizon is three days, $H^p = 3$. Therefore, 15 potential time slots may be available for the service of a newly arrived request. The customers are assumed to follow the MNL choice model. The expected utilities of obtaining the service in each of the five slots on the next day, i.e., the first day of the service horizon, are presented in Table 2.

Table 2: Expected utility of each possible slot

	8:00 – 10:00	10:00 – 12:00	12:00 – 14:00	14:00 – 16:00	16:00 – 18:00
Next Day	3	2	1	2	3

We selected these utility values for the experiment since we assume that most customers will prefer to be served at the beginning or the end of their working days instead of breaking the day in the middle. Each additional day of delay in the service results in a 20% reduction in the utility associated with each slot. The utility obtained from abandoning the process is 0. Recall that the expected utilities are assumed to be the same for all customers. This preference model reflects the fact that the customers desire to be served as soon as possible and that the time slots at the beginning or the end of the working day are more convenient than the mid-day slots.

In our experiments, the booking system is fully automated (e.g., via a website), and the service requests may arrive at any hour during the day. The arrival process of the requests is assumed to be periodic (daily). The fraction of the arrivals during each hour of the day, denoted by d_t, is presented in Figure 2. Note that the arrival process reflects two peaks: one in the late morning and another in the early evening. Naturally, the arrival rate is considerably low during the nighttime.

19

Figure 2: Fraction of requests arriving at each hour of the day

In our problem instances, we consider three values of the expected number of requests per day, denoted by D, namely, 64, 72, and 80 requests. The service time τ related to each instance is determined such that the expected amount of total required service time is identical among all the instances and equals $2/3$ of the six technicians' total working time, i.e., 2,400 minutes, which represents 400 minutes for each of the technicians. For example, for instances with 80 customers, the service time per customer is $\tau = 2400/80 = 30$ minutes.

In all the problem instances, the customers' locations are scattered on a 100×100 square, and the depot is located at the center. We generate three types of spatial configurations, namely, urban, suburban, and rural. For the urban and suburban configurations, 8 points are randomly generated from a uniform distribution, i.e., $x_z \sim U(10,90), y_z \sim U(10,90), z = 1 \dots 8$. Each point constitutes the center of a zone. For the rural configuration, 72 points are generated to represent the centers of zones. In this case, $x_z \sim U(2,98), y_z \sim U(2,98), z = 1 \dots 72$. The locations of the settlement centers are kept constant throughout the experiment.

Upon the arrival of each request, its location is generated in two steps: first, we randomly pick a zone, z, for it, and then we generate its exact location using random polar coordinates. An angle θ and distance r are randomly generated using a uniform distribution $\theta \sim U(0,2\pi), r \sim U(0,R)$, where $R = 5, 10,$ and 2 for the urban, suburban, and rural configurations, respectively. Next, based on the polar coordinates, (θ, r), the x coordinate is obtained as $x = x_z + r\cos(\theta)$, and the y coordinate is obtained as $y = y_z + r\sin(\theta)$. The travel times between the locations are the Euclidean distances.

The urban and suburban configurations are created to reflect a setting with a small number of high-demand cities. The travel time within each city is short to moderate. The rural configuration represents a highly dispersed setting in which the expected demand at each settlement (village or small town) is low. Figure 3 demonstrates the locations of requests that arrive on a typical day for the three types of configurations.

20

Figure 3: Example of the locations of the requests in each spatial configuration

A class of instances is determined by the spatial configuration and the demand intensity D since all the other parameters are constant. A training stream and a testing stream of requests are generated for each instance. The training stream is used to fit the parameters of each instance. The same stream is used for all the parameter combinations explored by our search procedure. The testing stream is used to evaluate the performance of the booking system with the final parameters, (α, β, γ), fitted in the search process. The training and test streams consist of requests obtained over 600 and 6000 booking days, respectively. They include 60 warmup days that have been excluded from the estimations to ensure that the observed measures represent the steady-state. The lengths of the simulation and the warmup period are determined and verified by extensive statistical analysis using a batch means technique.

We generate three instances with different random zone centers and random streams of requests for each unique combination of the spatial configuration (urban, suburban, rural) and the value of D (64, 72, 80 requests per day). There are 27 benchmark instances in total. A Python script that reproduces this dataset and its detailed description are available in the electronic appendix, which can be downloaded from http://www.eng.tau.ac.il/~talraviv/Publications/OnlineKit_Booking_Mar21.zip.

The instances are coded according to the number of requests D, the number of zones Z, the radius of the settlements R, and the serial replication number S. For example, the code name D72Z8R10S2 refers to the second instance with 72 requests per day and eight different settlements with a radius of 10 each (suburban setting).

5.2 Performance evaluation of the proposed method

In this section, we compare the performance of the proposed method with that of several commonly used policies from the literature for generating assortments and show the superiority of our method. We aim to gain insights related to the following questions:

- How well does the proposed method perform compared to the previously proposed and used myopic and rule of thumb policies in terms of the acceptance rate? Maximizing the acceptance rate is the objective of the assortment optimization process. In addition, what are the merits of fitting the opportunity cost to a Cobb–Douglas function rather than the previously proposed linear one?

- How does the proposed method affect the proportion of unserved customers, namely, customers who are offered no slots or customers who choose not to book a time slot even after a nonempty assortment of possible slots is offered to them? The arriving customers are divided into three

types: accepted, rejected, and abandoning. The division between the second and third types can be seen as an indicator of customer satisfaction.

- Does the proposed method induce more efficient routing plans and schedules? While the traveling distance is not explicitly included in the objective function, it is still of interest to the service provider.

- What is the typical structure of the assortments offered to customers? To answer this question, we report the cardinality and average utility of an assortment. Both are indicators of the quality of the offers. The *utility of an assortment* is defined as the sum of the utilities of its slots.

- What is the quality of the actual service the customers receive? To answer this question, we report the following features: the average utility of the (actual) service slot and the proportion of customers served on each day of the service horizon.

The following policies are considered:

- A myopic baseline policy, denoted as MYOPIC. Under this policy, upon the arrival of a request, all feasible slots are offered. Once the customer selects a slot, his request is inserted to the technician and the position in the route that minimizes the additional travel time. It should be noted that under any reasonable choice model, the optimal assortment in a policy that ignores future requests includes all the feasible time slots since extending the assortment by additional slots cannot increase the probability of abandonment. In particular, in the MNL model used in our experiment, any additional slot in the assortment strictly decreases the chance that a customer will abandon the booking process. Therefore, the assortment offered by our baseline policy is optimal in the myopic sense.

- A rule-of-thumb policy, denoted as X-SLOTS. Under this policy, upon the arrival of a request, only a subset of at most X feasible slots most valuable to the customer is offered, i.e., the feasible slots with the highest expected utility. Once a slot is selected, it is inserted exactly as in MYOPIC. This rule of thumb is reasonable to follow when resources for intense optimization of the assortments are unavailable, but the service provider wants to withhold some slots for future customers. In our experiment, we tested the 3-SLOTS and 5-SLOTS policies.

Our proposed learned policies:

- The linear approximation opportunity cost function is denoted as LINEAR.
- The Cobb-Douglas approximation opportunity cost function is denoted as COBB-DOUGLAS.

All the policies are compared based on the same geographies and testing streams of the requests. The same customer choice model and routing oracle (see Section 4.2) are used. For the COBB-DOUGLAS and LINEAR policies, the search process begins with an initial solution: $\alpha^0 = 0, \beta^0 = 0, \gamma^0 = 0$. The initial Δ is set as 0.1 for COBB-DOUGLAS and as 0.001 for LINEAR. Δ is updated by a factor of $\omega=2$. The numerical tolerance of the Cobb–Douglas function (COBB-DOUGLAS policy) is $\varepsilon = 0.01$. These values were selected after some preliminary experimentations. LINEAR was shown to benefit from a much smaller value of Δ. The process is stopped after four consecutive iterations without improving the objective function.

In Table 3, we summarize some performance measures collected in our experiment. Detailed statistics regarding the performance of each of the 27 instances and the fitted parameters of

approximated cost functions can be found in the underline{electronic appendix}. Columns 1 and 2 of Table 3 specify the demand rate and the policy used to create the assortments. For each combination of demand rate and policy, we present average performance measures for nine demand scenarios of 6000-day (the test stream), excluding the warmup period. Columns 3 and 4 present the proportion of the rejected and abandoning customers, respectively. Rejected customers are those whose assortments are empty. Customers presented with a nonempty assortment that chose to end the process without booking are considered abandoning customers. Columns 5-7 display the average number of slots offered on each of the three days of the service horizon. These averages are calculated based only on customers whose offered assortments are nonempty, that is, excluding the rejected customers. The data presented in these columns is also visualized in Figure 4. Columns 8-9 display the proportion of customers served on the first and second days. Only the accepted customers in each policy are considered. Column 10 displays the average utility of the assortments for each demand rate and each tested policy. Again, the presented measures are calculated based only on customers whose offered assortments are nonempty. In our view, this is the best measure of the quality of service, as perceived by customers who are not rejected. The last column displays the average utility of the actual service slot. The data are calculated based on the customers that are served.

Table 3: Summary of results for various demand rates and each of the tested policies

| D | Policy | % rejected | % abandon. | # slots offered | | | % Accepted | | Assortment utility | Selected slot utility |
				1st day	2nd day	3rd day	1st day	2nd day		
64	MYOPIC	6.86%	6.36%	0.10	0.62	4.12	1.22%	8.27%	6.37	1.45
	5-SLOTS	6.47%	6.70%	0.09	0.46	3.68	1.52%	8.07%	5.74	1.46
	3-SLOTS	4.36%	8.38%	0.10	0.30	2.44	2.84%	8.84%	4.09	1.49
	LINEAR	0.33%	4.52%	0.90	2.58	3.21	13.20%	38.94%	9.81	1.67
	COBB-DOUGLAS	0.12%	4.47%	0.69	1.90	3.73	11.12%	29.77%	8.96	1.62
72	MYOPIC	8.49%	6.66%	0.06	0.45	4.09	0.78%	6.08%	6.03	1.45
	5-SLOTS	8.07%	6.98%	0.06	0.36	3.74	1.07%	6.14%	5.60	1.45
	3-SLOTS	5.56%	8.69%	0.08	0.25	2.48	2.31%	7.35%	4.01	1.48
	LINEAR	0.59%	4.70%	0.93	2.62	3.07	14.11%	40.38%	9.80	1.69
	COBB-DOUGLAS	0.07%	4.83%	0.65	1.81	3.55	11.29%	30.42%	8.48	1.62
80	MYOPIC	10.52%	7.03%	0.03	0.32	4.01	0.49%	4.38%	6.03	1.43
	5-SLOTS	10.20%	7.27%	0.04	0.26	3.76	0.65%	4.52%	5.41	1.43
	3-SLOTS	7.28%	9.06%	0.05	0.20	2.50	1.62%	5.92%	3.89	1.46
	LINEAR	0.90%	5.02%	0.75	2.38	3.14	11.42%	38.12%	9.12	1.64
	COBB-DOUGLAS	0.16%	5.15%	0.54	1.60	3.60	9.61%	27.26%	7.93	1.58

Columns 3-4 in Table 3 illustrate the superiority of COBB-DOUGLAS over the other policies. For all demand rates, the proportion of the unserved customers is the lowest under this policy. This result is significant (p-value < 0.02 in four paired t-tests, between COBB-DOUGLAS and each of the other methods). Moreover, the proportion of rejected customers under COBB-DOUGLAS, i.e., those with an empty assortment, is negligible. In terms of the proportion of served customers, LINEAR is second best. We also note that the rule-of-thumb policies are slightly better than MYOPIC in terms of served customers. In paired t-tests carried out using the results of the 27 instances, the acceptance rates of both

5-SLOTS and 3-SLOTS are greater than that of MYOPIC (p-value < 0.001). However, the rule of thumb policies lag far behind LINEAR and COBB-DOUGLAS.
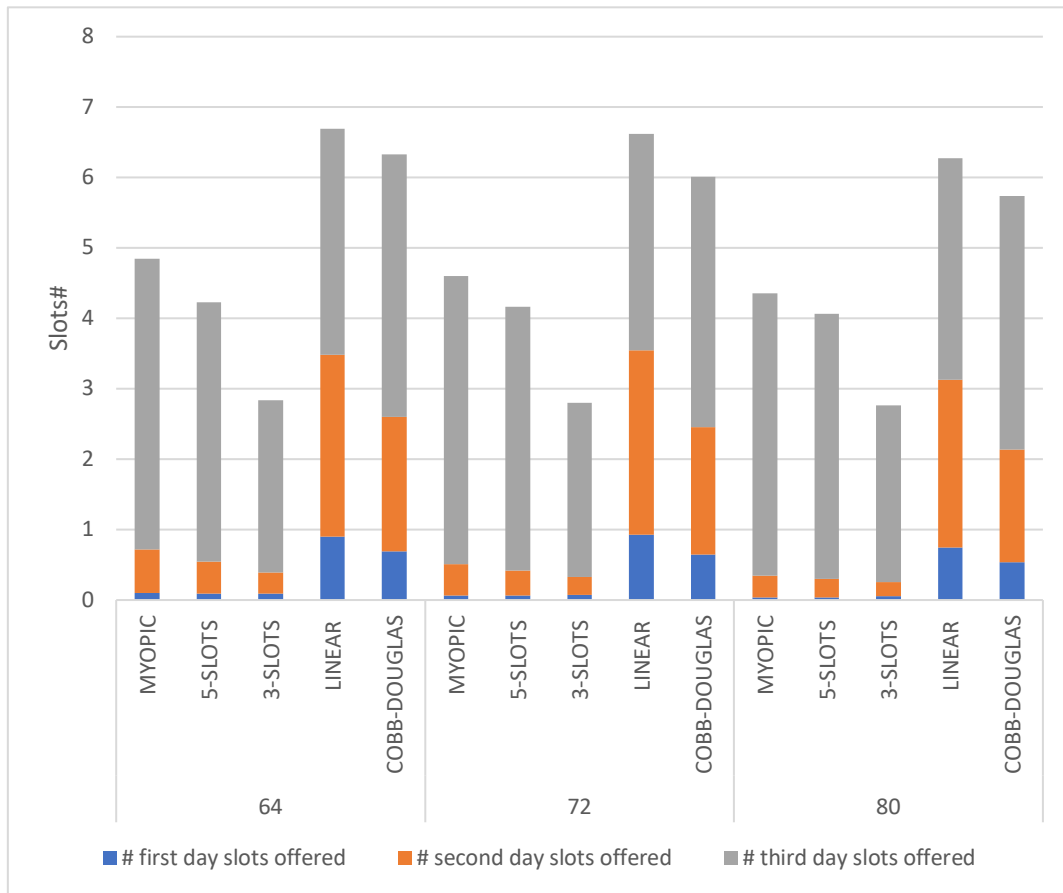
Figure 4: Number of slots offered on each working day for the different policies

It is apparent from Columns 5-7 in Table 3 and Figure 4 that COBB-DOUGLAS and LINEAR are superior to all other policies in terms of both the total number of slots offered and the number of slots offered for the first and second days. LINEAR offers the largest number of slots and more first and second-day slots than the other methods.

Note that the assortments offered by all the policies are subsets of the assortments that MYOPIC would offer in the same states. However, by withholding some feasible slots, the system may utilize its resources more efficiently and consequently may be able to offer more slots on average. Both COBB-DOUGLAS and LINEAR achieve this efficiency gain in a manner that enables enhancing the long-run availability of slots. Note that 3-SLOTS and 5-SLOTS offer assortments with fewer slots on average, especially on the second and third days. Moreover, it is apparent from Columns 8-9 that the proportions of customers served on the first and second days are the highest for LINEAR and second-highest for COBB-DOUGLAS.

From the measures presented in Columns 10-11 in Table 3, we observe that LINEAR induces the highest utility for the served customers. This claim holds when we consider the total utility of the assortments (Column 10) and the utility of the selected slots (Column 11). It seems that while COBB-

DOUGLAS is preferred over LINEAR when we aim to maximize the long-run acceptance ratio of the system, LINEAR can obtain a somewhat higher service level without a considerable efficiency loss. However, our objective function is to maximize the acceptance rate, and according to this measure, COBB-DOUGLAS dominates all the other policies.

The routing costs are not optimized directly by the proposed methods. However, we hypothesized that a high acceptance rate is correlated with efficient routing since it allows to better utilize the service provider's resources. Figure 5 displays the average travel time per served request for the three configurations (urban, suburban and rural) for the five tested policies.

Figure 5: Average travel time per served request for the three configurations

It is clear that the two policies that apply advanced assortment optimization methods, namely, COBB-DOUGLAS and LINEAR, induce considerably shorter travel times compared to those of the other policies and thus utilize the fleet much more efficiently. The effect is most significant for the rural instances. A decrease of almost 50% in the average travel time is noted when following COBB-DOUGLAS rather than MYOPIC or the two rule-of-thumb policies. This advantage is expected since when customers are sparser, assortment optimization can save more in terms of travel time. COBB-DOUGLAS yields slightly lower average travel times than LINEAR in the urban configuration. Its advantage over LINEAR is more apparent for the rural and suburban configurations.

We conclude Section 5.2 by summarizing our main findings:

1. COBB-DOUGLAS leads to significantly better routes and schedules, as reflected by the larger number of accepted customers and average travel time of the vehicles per request. LINEAR, which also applies an advanced assortment optimization scheme, is the second-best in terms of acceptance rate and routing costs per served customer.
2. While COBB-DOUGLAS maximizes the number of accepted customers for service, LINEAR is the best policy in terms of the provided service level.
3. COBB-DOUGLAS and LINEAR yield similar results in all the performance measures and strictly dominate all other policies. COBB-DOUGLAS and LINEAR are the closest policies to

an efficiency frontier between the number of served customers and the quality of service provided to the accepted customers.

5.3 The importance of the model components

In this section, we verify that all three arguments (RTS, RTR, and TT) of the Cobb-Douglas opportunity cost function f are essential by applying our method while eliminating each of the arguments, one at a time. These reduced models can be viewed as special cases of the full model, with α, β, or γ fixed at 0. Figure 6 presents the estimated acceptance ratio for these three modified models for each of the 27 instances. The results of the full model are displayed as well. The performances of the three partial models are reported in detail in the electronic appendix.
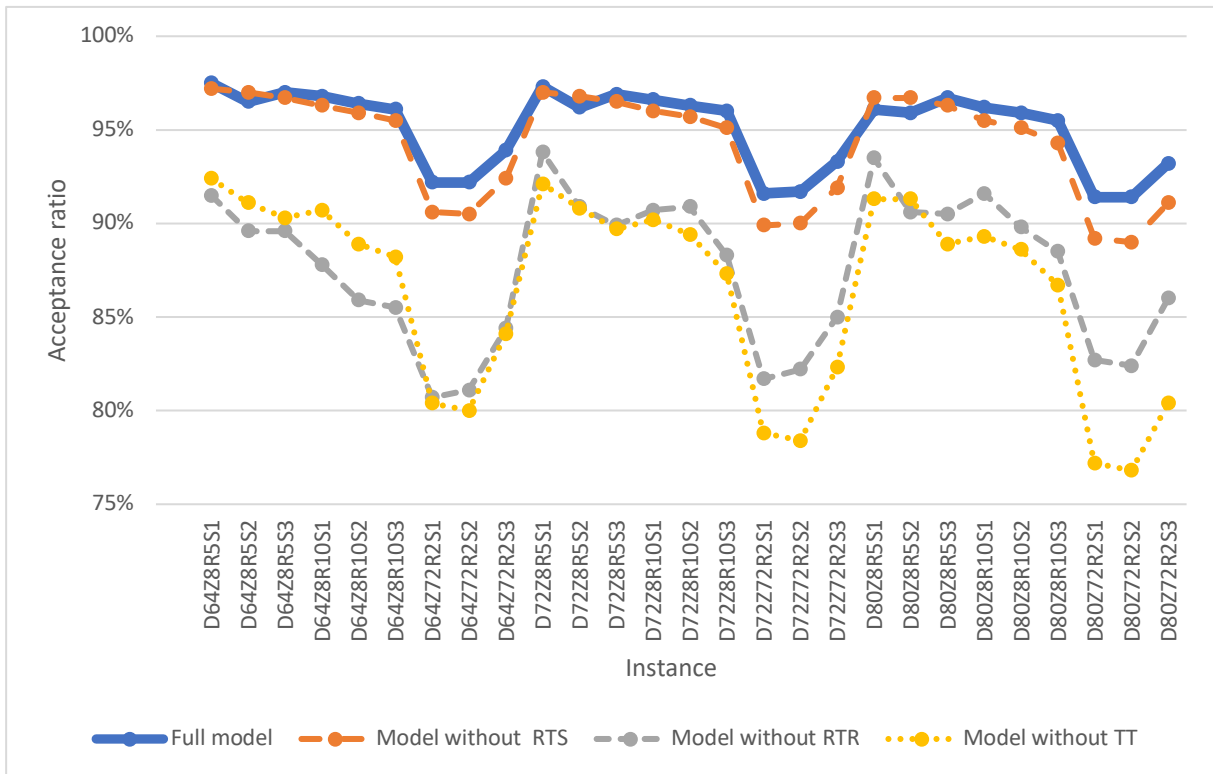

**Figure 6**: Acceptance ratio of the full model and its special cases

The full model obtains significantly better average results than each of the partial models (p-value < 0.001 in paired t-tests that compare the full model with each of the partial ones). The marginal effect of the RTS measure on the model's performance is relatively moderate. In contrast, the effect of the TT and RTR measures is considerably more notable, especially in rural settings.

5.4 The merits of a multiday service horizon

Recall that the proposed booking mechanism allows technician appointment scheduling several days into the future with multiple cutoff times. In contrast, previous literature focused on a service horizon that consists of a single day and overlooked the interactions between consecutive working days. In this section, we examine the merits of considering a multiday service horizon. To this end, we repeat our main experiment for a subset of 9 out of 27 instances, namely, the instances for which $S = 1$, with two

additional service horizon lengths: one day and five days. The detailed results of each instance are presented in the underlined electronic appendix.

Figure 7 displays the average proportion of abandoning and rejected customers for the five tested policies and the three lengths of the service horizon.
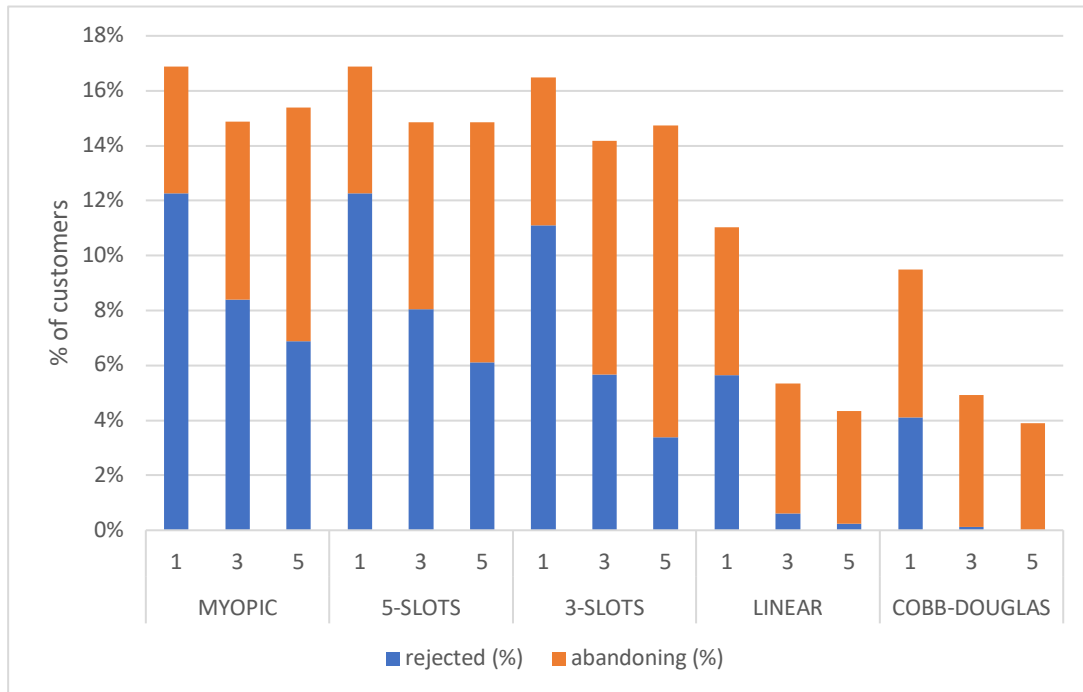

Figure 7: The proportion of rejected and abandoning customers for different service horizons

This experiment verifies that the advantage of COBB-DOUGLAS over the other policies in terms of the acceptance ratio is not sensitive to the length of the service horizon. For all policies, increasing the service horizon strictly decreases the proportion of customers with an empty assortment, i.e., the rejected customers. This finding is expected since increasing the service horizon increases the number of feasible slots.

The proportion of customers that abandon the process strictly increases for all policies that do not apply advanced assortment optimization mechanisms. Under these policies, the proportion of unserved customers, which is the sum of the abandoning and rejected ones, is minimized following a three-day service horizon. In other words, COBB-DOUGLAS and LINEAR are proven to be able to exploit the greater flexibility obtained by a longer service horizon. To understand this phenomenon, recall that in a steady-state, increasing the length of the service horizon does not increase the capacity of the resources available for the operator. In fact, limiting the number of days in the service horizon acts as an additional availability control measure and causes the operator to reject some of the customers that are hard to serve. These customers could be assigned a slot, if the service horizon would be longer, which could eliminate the possibility to serve a greater number of future customers. COBB-DOUGLAS and LINEAR can estimate the opportunity cost accurately enough to mitigate this effect.

The increase in the proportion of abandoning customers under the rule of thumb policies is caused by the structure of the presented assortments. Figure 8 displays the average cardinality of the assortments for the five policies and three lengths of the service horizon.

Figure 8: The number of slots in an average assortment

Figure 8 demonstrates that increasing the service horizon has little effect on the size of the assortments for simple policies that do not perform assortment optimization. For the 3-SLOTS policy, the effect is almost negligible, which follows from the fact that the number of offered slots is already limited to a very small value under this policy. The significant increase in the size of the assortments noted for COBB-DOUGLAS and LINEAR demonstrates that they are suited for controlling the demand in a multiday setting. Indeed, LINEAR can offer larger assortments than COBB-DOUGLAS, with a relatively small effect on the acceptance ratio in the long run.

Finally, Figure 9 displays the total utility of the offered assortments for all policies and lengths of the service horizon. We observe that the effect of the service horizon's length on the quality of the offered assortments depends on the policy. Greedier policies tend to "sell" all attractive slots early and, in the long run, can offer the poorest slots only. Hence, under these policies, the total utility decreases with the length of the service horizon. COBB-DOUGLAS and LINEAR exploit the additional flexibility obtained by a longer service horizon to improve the quality of the average offer made to customers. MYOPIC and 5-SLOTS generate offers with higher utility in a single-day setting, but recall that their acceptance rates are much lower.

**Figure 9:** Utility of assortments for all policies and service horizons

We conclude Section 5.4 with a summary of our findings. We observed that COBB-DOUGLAS and LINEAR are best suited to exploit the multiday setting. Under these policies, a longer service horizon allows a higher acceptance rate and better service to the customer than MYOPIC and X-SLOTS.
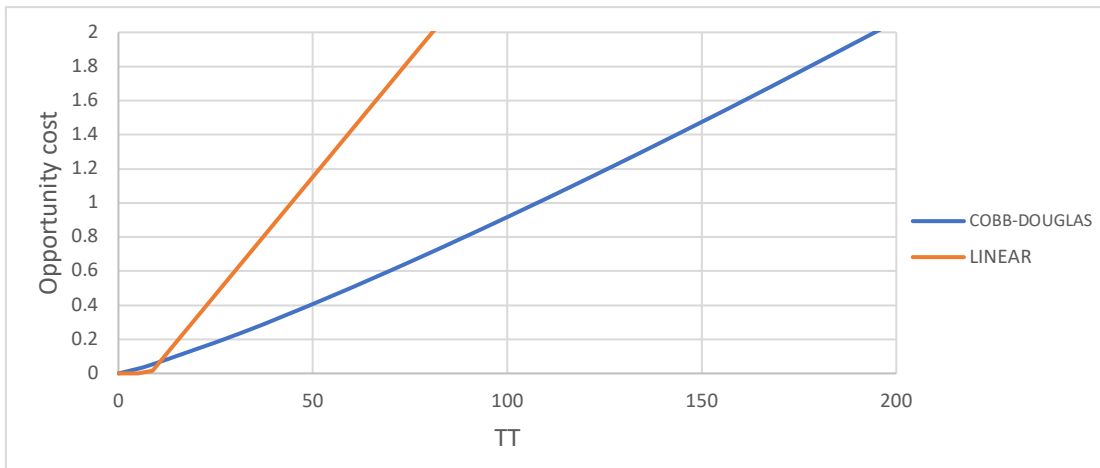
5.5 The opportunity cost functions

In this section, we provide examples of the estimated opportunity cost function using LINEAR and COBB-DOUGLAS policies. We focus on two instances with arrival rates of 80 requests per day: one in rural and one in urban settings. Each opportunity cost function is presented by three plots in Figures 10 and 11, respectively. In each of the three plots, one of the features (RTS, RTR, or TT) is the independent variable, and the other two are fixed at their median value. The dependent variable on the vertical axis is the estimated opportunity cost. The two opportunity cost functions (based on linear and Cobb-Douglas approximation) are plotted.
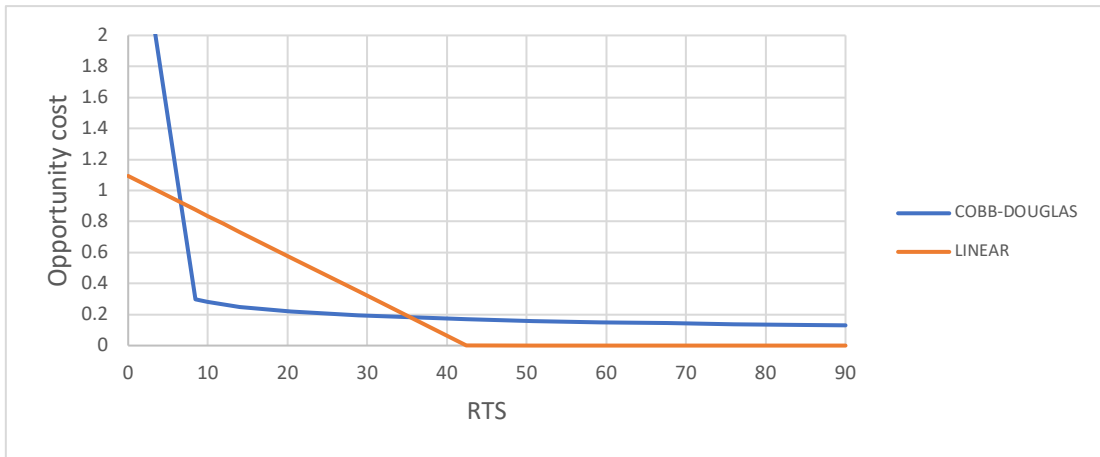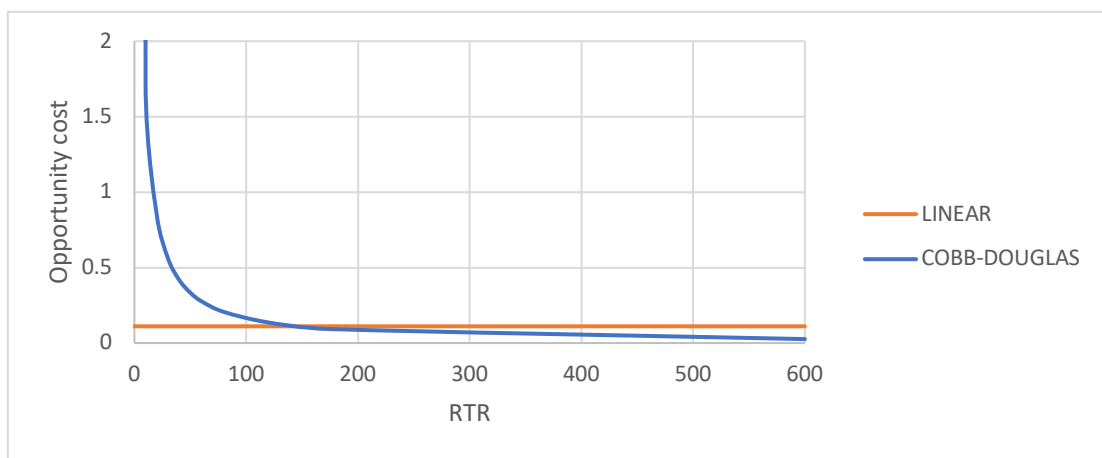
(a)
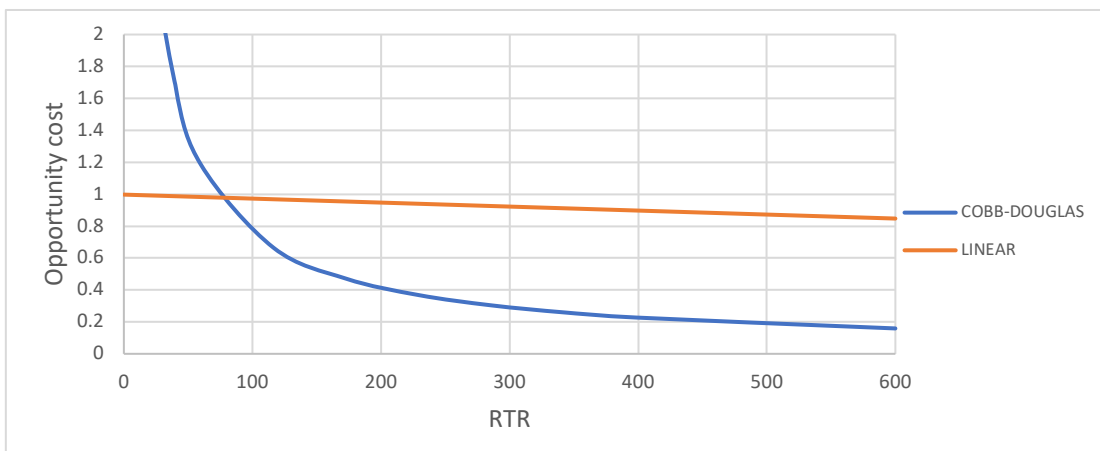
(b)

(c)

Figure 10: The opportunity cost function of an instance in rural setting

(a)

(b)

(c)

Figure 11: The opportunity cost function of an instance in urban **setting**

In both the rural and urban setting and with both the linear and Cobb-Douglas value function approximations, the opportunity cost is decreasing in the remaining time in the slot (RTS), non-increasing and the remaining time in the route (RTR) and increasing in the added travel time due to the insertion of the request (TT). These are expected results since it is worthy to accept requests at routes and slots with large remaining time and when the additional travel time incurred by the requests is small .

Interestingly, while LINEAR policy is more likely than COBB-DOUGLAS to reject a request that requires long additional travel time [See Figure 10(c) and 11(c)], the routes created by the latter are more efficient, as demonstrated in Figure 5. This may be because, in the long run, COBB-DOUGLAS accepts more requests and succeeds in reducing the average travel time per request.

The opportunity cost function as estimated with Cobb-Douglas is nearly constant in most of the range of the RTS feature, which implies that the information gained from this feature is not very useful. A fact that we already demonstrated in Figure 6. A similar phenomenon is observed for the RTR feature with the linear value function approximation .

We observe that the difference between the linear and Cobb-Douglas opportunity cost functions is sharper in the rural setting (Figure 10) than in the urban setting (Figure 11). We note that spatial considerations are more critical when the customers are geographically dispersed. Recall that we already demonstrated above that the advantage of COBB-DOUGLAS over LINEAR policy is more significant in the rural setting.

While the estimated value functions differ from each other, we conjecture that the similar performance of LINEAR and COBB-DOUGLAS steams from a similar estimation of the opportunity cost in the relevant argument domain. To test this conjecture, we calculated the estimated opportunity cost using the two functions based on 10000 actual triplets of the measured arguments (RTS, RTR, TT) sampled from our simulation, for each of the 27 instances.

It should be noted that the assortment optimization process eliminates any slot with an opportunity cost greater than one. Therefore, we conducted the comparison separately for values greater and smaller than one. In about 92% (resp., 93%) of the cases where the opportunity cost was estimated to be smaller than one by the Cobb-Douglas (resp., linear) function, it was also estimated to be smaller than one by the linear (resp. Cobb-Douglas) function. That is, if a slot is a candidate for the assortment by one of the methods, it is also likely to be a candidate when using the other method. In about 83% of the triplets, both functions estimated opportunity cost to be smaller than one. The average absolute difference between the two estimated opportunity costs when it was smaller than one (with both functions) was 0.13. The observed similarities between the two functions explain the similar assortment and therefore similar performances obtained by the two methods.

6 Summary and conclusions

In the past, the booking of field service personnel and attended home delivery was usually performed through a sequential interaction between customers and service representatives in call centers. The current practice is increasingly based on using online services (such as web or mobile apps), which can conveniently present customers with all the offered alternatives at once. In this case, the process of determining the alternatives to offer requires an algorithmic approach.

In this study, we present an effective framework that addresses this challenge in realistic settings when the service horizon spans several days with different cutoff times and when the price of the service is exogenous and independent of the service time slot. Previous studies focused on the single cutoff case and overlooked the interactions between consecutive days. When the service provider can accept requests several days in advance, the decisions made on different days are interrelated. In this case, the problem is no longer separatable by working days. Therefore, the well-studied objective of maximizing

the number of accepted requests (during a single planning period) should be replaced to maximize the steady-state acceptance rate.

The idea of the proposed method is to estimate the opportunity cost of accepting a request in each feasible slot based on a compact approximated representation of the state of the system. Using these costs, the provider optimizes the offered assortment of slots for each service request, assuming a discrete choice model representing the customers' preferences. The successful implementation of the proposed methods is highly dependent on a reliable forecast of future demand and sufficiently accurate modeling of the user choice model.

An important insight gained by this study is that allowing the customers to book their service requests over a multiday horizon may result in a significant improvement in both the acceptance rate and the utility of the customers. This insight holds only when the assortments of offered slots are constructed using anticipatory approaches and optimized intelligently. When the assortments are created myopically based on feasibility alone, a multiday horizon may improve the acceptance rate, but this always comes at the expense of the quality of the offered slots.

We demonstrate that optimizing the offered assortments allows increasing the acceptance rate of service requests significantly even without a pricing mechanism, which is popular in the literature but not suitable for all the business scenarios. Finally, we show that it is possible to approximate and fit the opportunity cost to the Cobb-Douglas function. This approach is somewhat more effective than the more commonly used linear approximation.

A closely related line of research considers booking home-attended deliveries. In this case, the booking system should also consider the capacity constraint of the vehicles. Introducing this constraint into the proposed method requires including the remaining capacity in the state representation used for estimating the opportunity cost. It also straightforwardly affects the routing module.

The literature regarding the booking of field service personnel in multiday settings is still sparse, and we believe that the following extensions of our model require further investigation:

(1) Consider several classes of customers that differ in their time preferences, value to the service provider, required technician skills, and expected service duration. Each class may have a different optimal booking policy.

(2) If the provider has no capacity to serve a request during the service horizon, it may be extended to make additional service slots available. An interesting direction for future research is to introduce decisions regarding the length of the service horizon into the model and allowing its dynamic extension.

(3) Our solution approach makes a simplifying assumption that the choice model of the customers is fixed for each day of the week. In practice, the preferences of customers can be affected by seasonality and weekly cycles. For example, customers may prefer weekend time slots over weekdays. Introducing seasonality and day-of-the-week effects would require additional parameters for the function that estimates the opportunity cost. While this is possible using the method proposed in this study, much more computational effort will be required to fit the parameters correctly. A future study to meet this challenge will have to devise a much more efficient method to solve the parameter fitting problem.

(4) Concurrency issues should be resolved. When a large service provider performs the booking process, the system is likely to interact simultaneously with several customers. The slots offered to a particular customer may or may not be available to others while the interaction commences. The optimal booking policy may be affected by these concurrency considerations.

References

1. Niels Agatz, Ann Campbell, Moritz Fleischmann, Martin Savelsbergh (2011) Time Slot Management in Attended Home Delivery. Transportation Science 45(3):435-449.
2. Maria Albareda-Sambola, Elena Fernández, Gilbert Laporte (2014) The dynamic multiperiod vehicle routing problem with probabilistic information. Computers & Operations Research, 48, 31–39.
3. Enrico Angelelli, Nicola Bianchessi, Renata Mansini, M. Grazia Speranza (2009) Short term strategies for a dynamic multi-period routing problem. Transportation Research Part C: Emerging Technologies, 17(2): 106–119.
4. Kursad Asdemir, Varghese S. Jacob, Ramayya Krishnan (2009) Dynamic pricing of multiple home delivery options. European Journal of Operational Research 196: 246-257.
5. Edison Avraham, Tal Raviv (2020) The data-driven time-dependent traveling salesperson problem. Transportation Research Part B: Methodological 134: 25-40.
6. Russell W. Bent, Pascal Van Hentenryck (2004) Scenario-Based Planning for Partially Dynamic Vehicle Routing with Stochastic Customers. Operations Research 52(6):977-987.
7. Kris Braekers, Katrien Ramaekers, Inneke Van Nieuwenhuyse (2016) The vehicle routing problem: State of the art classification and review. Computers & Industrial Engineering 99: 300-313.
8. Dominic Bühler, Robert Klein, Michael Neugebauer (2016) Model-based delivery cost approximation in attended home services. Computers & Industrial Engineering 98: 78-90.
9. Ann Melissa Campbell, Martin W. P. Savelsbergh (2005) Decision Support for Consumer Direct Grocery Initiatives. Transportation Science 39(3):313-327.
10. Ann Melissa Campbell, Martin Savelsbergh (2006) Incentive Schemes for Attended Home Delivery Services. Transportation Science 40(3):327-341.
11. Xi Chen, Barrett W. Thomas, Mike Hewitt (2017) Multi-period technician scheduling with experience-based service times and stochastic customers. Computers & Operations Research, 82, 1–14.
12. Catherine Cleophas, Jan Fabian Ehmke (2014) When Are Deliveries Profitable? Business & Information Systems Engineering 6(3):153-163.
13. Jan Fabian Ehmke, Ann Melissa Campbell (2014) Customer acceptance mechanisms for home deliveries in metropolitan areas. European Journal of Operational Research 233: 193-207.

14. Guillermo Gallego, Richard Ratliff, Sergey Shebalov (2015) A General Attraction Model and Sales-Based Linear Program for Network Revenue Management Under Customer Choice. Operations Research 63(1):212-232.

15. Gianpaolo Ghiani, Emanuele Manni, Antonella Quaranta, Chefi Triki (2009) Anticipatory algorithms for same-day courier dispatching. Transportation Research Part E: Logistics and Transportation Review 45(1): 96-106.

16. Gianpaolo Ghiani, Emanuele Manni, Barrett W. Thomas (2012) A Comparison of Anticipatory Algorithms for the Dynamic and Stochastic Traveling Salesman Problem. Transportation Science 46(3):374-387.

17. Florent Hernandez, Michel Gendreau, Jean-Yves Potvin (2017) Heuristics for tactical time slot management: a periodic vehicle routing problem view. International Transactions In Operational Research. 24(6):1233–1252.

18. Robert Klein, Sebastian Koch, Claudius Steinhardt, Arne K. Strauss (2020) A review of revenue management: Recent generalizations and advances in industry applications. European Journal of Operational Research 284: 397-412.

19. Robert Klein, Jochen Mackert, Michael Neugebauer, Claudius Steinhardt (2018) A model-based approximation of opportunity cost for dynamic pricing in attended home delivery. OR Spectrum 40(4): 969-996.

20. Robert Klein, Michael Neugebauer, Dimitri Ratkovitch, Claudius Steinhardt (2019) Differentiated Time Slot Pricing Under Routing Considerations in Attended Home Delivery. Transportation Science 53(1):236-255.

21. Sebastian Koch, Robert Klein (2020) Route-based approximate dynamic programming for dynamic pricing in attended home delivery. European Journal of Operational Research 287: 633-652.

22. Charlotte Köhler, Jan Fabian Ehmke, Ann Melissa Campbell (2020) Flexible time window management for attended home deliveries. Omega 91:102023

23. Averill L. Law (2013), Simulation modeling and analysis (Fifth Edition), McGraw-Hill Education.

24. Jochen Mackert (2019) Choice-based dynamic time slot management in attended home delivery. Computers & Industrial Engineering 129: 333-345.

25. Arturo E. Pérez Rivera, Martijn R.K. Mes (2017) Anticipatory freight selection in inter- modal long-haul round-trips. Transportation Research Part E: Logistics and Transportation Review, 105, 176–194.

26. Victor Pillac, Michel Gendreau, Christelle Guéret, Andrés L. Medaglia (2013) A review of dynamic vehicle routing problems. European Journal of Operational Research 225 (1): 1-11.

27. Warren B. Powell (2019) A unified framework for stochastic optimization. European Journal of Operational Research 275 (3): 795-821.

28. Harilaos N. Psaraftis, Min Wen, Christos A. Kontovas (2016) Dynamic vehicle routing problems: Three decades and counting. Networks **67**(1) 3-31.

29. María I. Restrepo, Frédéric Semet, Thomas Pocreau (2019) Integrated Shift Scheduling and Load Assignment Optimization for Attended Home Delivery. Transportation Science 53(4):1150-1174.

30. Sheldon Ross (2013), Simulation (Fifth Edition), Academic Press.

31. Ninja Soeffker, Marlin Ulmer, Dirk Mattfeld (2017) On Fairness Aspects of Customer Acceptance Mechanisms in Dynamic Vehicle Routing. Proceedings of Logistikmanagement 2017: 17-24.

32. Marius M. Solomon (1987) Algorithms for the vehicle routing and scheduling problem with time window constraints. Operations Research 35(2): 54-65.

33. Arne Strauss, Nalan Gülpınar, Yijun Zheng (2020) Dynamic pricing of flexible time slots for attended home delivery. European Journal of Operational Research https://doi.org/10.1016/j.ejor.2020.03.007.

34. Arne K. Strauss, Robert Klein, Claudius Steinhardt (2018) A review of choice-based revenue management: Theory and methods. European Journal of Operational Research 271: 375-387.

35. Marlin W. Ulmer, Dirk C. Mattfeld, Felix Köster (2018a) Budgeting Time for Dynamic Vehicle Routing with Stochastic Customer Requests. Transportation Science 52(1):20-37.

36. Marlin W. Ulmer, Ninja Soeffker, Dirk C. Mattfeld (2018b) Value function approximation for dynamic multi-period vehicle routing. European Journal of Operational Research 269: 883-899.

37. Marlin W. Ulmer (2019) Anticipation versus reactive reoptimization for dynamic vehicle routing with stochastic requests. *Networks* 73:277–291.

38. Marlin W. Ulmer, Barrett W. Thomas (2019) Meso-parametric value function approximation for dynamic customer acceptances in delivery routing. European Journal of Operational Research 285(1): 183-195.

39. Marlin W. Ulmer (2020) Dynamic Pricing and Routing for Same-Day Delivery. Transportation Science 54(4):1016-1033

40. Stacy A. Voccia, Ann Melissa Campbell, Barrett W. Thomas (2019) The Same-Day Delivery Problem for Online Purchases. Transportation Science 53(1):167-184.

41. Min Wen, Jean-François Cordeau, Gilbert Laporte, Jesper Larsen (2010) The dynamic multi-period vehicle routing problem. Computers & Operations Research, 37(9): 1615–1623.

42. Xinan Yang, Arne K. Strauss, Christine S. M. Currie, Richard Eglese (2016) Choice-Based Demand Management and Vehicle Routing in E-Fulfillment. Transportation Science 50(2):473-488.

43. Xinan Yang, Arne K. Strauss (2017) An approximate dynamic programming approach to attended home delivery management. European Journal of Operational Research 263: 935-945.