Transportation Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Customer-Centric Dynamic Pricing for Free-Floating Vehicle Sharing Systems

Christian Müller, Jochen Gönsch, Matthias Soppert, Claudius Steinhardt

To cite this article:
Christian Müller, Jochen Gönsch, Matthias Soppert, Claudius Steinhardt (2023) Customer-Centric Dynamic Pricing for Free-Floating Vehicle Sharing Systems. Transportation Science 57(6):1406-1432. https://doi.org/10.1287/trsc.2021.0524

informs
https://pubsonline.informs.org/journal/trsc

TRANSPORTATION SCIENCE
Vol. 57, No. 6, November–December 2023, pp. 1406–1432
ISSN 0041-1655 (print), ISSN 1526-5447 (online)

Customer-Centric Dynamic Pricing for Free-Floating Vehicle Sharing Systems

Christian Müller,[a] Jochen Gönsch,[a,*] Matthias Soppert,[b] Claudius Steinhardt[b]

[a] University of Duisburg-Essen, 47057 Duisburg, Germany; [b] University of the Bundeswehr Munich, 85577 Neubiberg, Germany

*Corresponding author

Contact: christian.mueller.9@uni-due.de, https://orcid.org/0000-0002-2950-446X (CM); jochen.goensch@uni-due.de, https://orcid.org/0000-0002-5699-1320 (JG); matthias.soppert@unibw.de, https://orcid.org/0000-0003-3399-0525 (MS); claudius.steinhardt@unibw.de, https://orcid.org/0000-0003-4263-6608 (CS)

Received: December 15, 2021
Revised: August 31, 2022; March 10, 2023; June 19, 2023
Accepted: August 8, 2023
Published Online in Articles in Advance: November 6, 2023

https://doi.org/10.1287/trsc.2021.0524

Abstract. Free-floating vehicle sharing systems such as car or bike sharing systems offer customers the flexibility to pick up and drop off vehicles at any location within the business area and, thus, have become a popular type of urban mobility. However, this flexibility has the drawback that vehicles tend to accumulate at locations with low demand. To counter these imbalances, pricing has proven to be an effective and cost-efficient means. The fact that customers use mobile applications, combined with the fact that providers know the exact location of each vehicle in real-time, provides new opportunities for dynamic pricing. In this context of modern vehicle sharing systems, we develop a profit-maximizing dynamic pricing approach that is built on adopting the concept of customer-centricity. Customer-centric dynamic pricing here means that, whenever a customer opens the provider's mobile application to rent a vehicle, the price optimization incorporates the customer's location as well as disaggregated choice behavior to precisely capture the effect of price and walking distance to the available vehicles on the customer's probability for choosing a vehicle. Two other features characterize the approach. It is origin-based, that is, prices are differentiated by location and time of rental start, which reflects the real-world situation where the rental destination is usually unknown. Further, the approach is anticipative, using a stochastic dynamic program to foresee the effect of current decisions on future vehicle locations, rentals, and profits. We propose an approximate dynamic programming-based solution approach with nonparametric value function approximation. It allows direct application in practice, because historical data can readily be used and main parameters can be precomputed such that the online pricing problem becomes tractable. Extensive numerical studies, including a case study based on Share Now data, demonstrate that our approach increases profits by up to 8% compared with existing approaches from the literature.

History: This paper has been accepted for the *Transportation Science* Special Issue on 2021 TSL Workshop: Supply and Demand Interplay in Transport and Logistics.

Supplemental Material: The e-companion is available at https://doi.org/10.1287/trsc.2021.0524.

Keywords: free-floating vehicle sharing system • customer-centric dynamic pricing • data-driven nonparametric value function approximation

1. Introduction

Vehicle sharing systems (VSSs), such as car sharing, bike sharing, or scooter sharing, are specific shared mobility systems (Mourad, Puchinger, and Chu 2019) that allow users to flexibly and spontaneously rent vehicles for a short period of time (Ataç, Obrenović, and Bierlaire 2021). In contrast to other popular shared mobility systems, VSSs enable individual trips (conducted by the user), whereas, for example, ride-hailing (e.g., Wang and Yang 2019) requires connecting passengers with drivers in a two-sided market, and ride-pooling (e.g., Ke et al. 2020) as well as dial-a-ride (e.g., Qiu, Li, Zhao 2018) concepts strive for shared trips.

There are three fundamental types of VSSs which, from the customers' view, decisively differ with regard to the degree of flexibility they offer. In *two-way station-based* systems, customers have to return vehicles to the pick-up station, whereas in *one-way station-based* systems, customers can pick-up and drop-off the vehicle at any station. *Free-floating* is the most flexible variant, as it allows customers to pick-up and drop-off vehicles at any public parking spot in the business area of the VSS provider (e.g., Chow and Yu 2015). For this reason, free-floating VSSs have become a very popular type in urban areas (e.g., Statista 2022). However, higher degrees of flexibility come with an important drawback: Due to unbalanced demand patterns and the oscillation of the demand intensity over the course of the day, vehicles accumulate at certain locations (usually the outskirts) over time, while other areas lack vehicles (usually downtown). This so-called "tide phenomenon" (spatiotemporal demand asymmetries (Jorge and Correia 2013, Côme 2014)) is even more pronounced for free-floating VSSs than for station-based one-way VSSs. The reason is that while demand and supply in a free-floating VSSs spread across the entire business area, they concentrate to rather few locations in station-based VSSs (Wagner et al. 2015).

Pricing is an established tool to counter these imbalances and to improve the system's profit. The idea is to nudge some customers to slightly adapt their travel plans, for example, to pick up a sharing vehicle at a low demand location instead of a high demand location. Thus, adequate pricing can achieve a vehicle availability that assures an appropriate service level.

Existing pricing approaches from the related literature on VSSs determine vehicle-specific prices, meaning that all customers are getting displayed the same (minute) price for a particular vehicle when checking prices at the same time. These approaches do not leverage on the detailed disaggregate information which is available in modern free-floating VSSs. Opposed to these approaches, a *customer-centric* pricing approach determines prices under consideration of the situation-specific information of each customer. As depicted in Figure 1(a), this particularly comprises a customer's location as well as the information about the available vehicles within the customer's walking distance. Thus, in contrast to vehicle-specific pricing, customer-centric pricing may result in different prices for the same vehicle when customers are at different locations (Figure 1(b)). It can be considered as an application of choice-based revenue management (Strauss, Klein, and Steinhardt 2018) with the particularity in free-floating VSSs that the consideration set of available vehicles is determined by the customer's location. The core idea of customer-centric pricing has successfully been applied to shared mobility systems other than VSSs, that is, to mobility on demand services (Qiu, Li, Zhao 2018). However, these mobility systems differ fundamentally from VSSs in that the provider operates a platform that receives requests for the customers' desired trips and then matches demand to supply. This (moving) supply of vehicles may either be under full control of the provider (Ke et al. 2020) or it may be influenced by pricing as well, resulting in a two-sided market (Wang and Ma 2019). Thus, these approaches are not directly applicable to our context, and in this work we propose such customer-centric dynamic pricing which is specifically tailored to free-floating VSSs.

In this work, we consider a profit-maximizing free-floating VSS provider's dynamic online pricing problem with a strong focus on applicability in practice, meaning that problem definition and solution approach design are based on the circumstances and requirements in practice. More precisely, whenever a customer considers prices in the mobile application, prices need to be determined based on the currently available information. For this problem, we develop a new pricing approach that is characterized by the following three distinguishing features:

- First, and most importantly, the pricing is *customer-centric*. As briefly described above, adopting the concept of customer-centricity means that prices are situation-specific for each customer. As a consequence, the online price optimization can leverage on detailed disaggregate information like the customer's location. This, in turn, allows to exploit that the location not only determines the vehicles within the customer's walking distance, but

Figure 1. (Color online) Illustration of the First Distinguishing Feature of the Developed Pricing Approach: *Customer-Centricity*

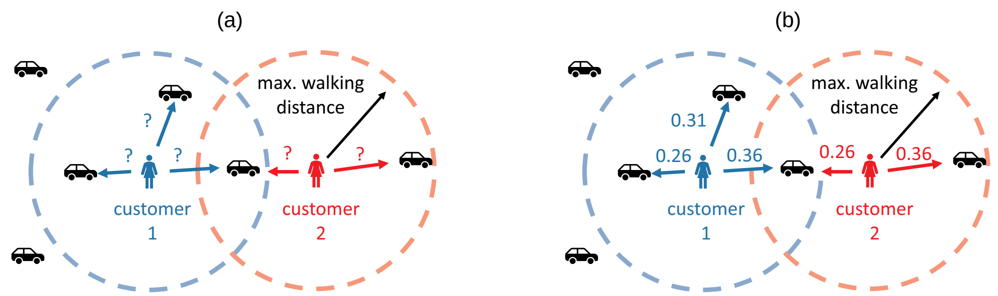

Notes. (a) Provider's pricing problem: The "?" indicate prices to be optimized online. (b) Resulting customer's choice situation: Customers might see *different* prices for a vehicle.

that the distance to a vehicle also impacts the customer's utility and, thus, the probability of choosing it. Due to the customer-centricity, the pricing approach can incorporate the customer choice behavior through an appropriate choice model. As stated above, this pricing can result in one specific vehicle having different prices for different customers, but this does not mean that pricing is *personalized*. More specifically, we do *not* use socio-demographic characteristics such as age or income to potentially exploit individual willingness-to-walk or individual price sensitivity. Only the location of the customer's device when she looks for vehicles is used to account for the impact of distance to different vehicles on customer utility. Thus, prices are identical for every customer who faces the same situation. Note, however, that—as for all pricing approaches in which locations of customers or vehicles are decisive—prices may indirectly be dependent on socio-demographic characteristics, for example, because of location-dependent income levels.

• Second, within our pricing approach, prices can be varied (solely) based on location and time of a rental's start, denoted as *origin-based pricing*, in line with the business decision of Share Now. In particular, information on a rental's destination cannot be used, because it is not available in reality: Asking customers for their destination beforehand contradicts the spontaneous selling proposition of free-floating VSSs (Soppert et al. 2022). The alternative, that is, displaying prices for all potential origin-destination (and -time) combinations in advance of a rental, is impracticable in general, given that free-floating VSS providers often discretize their business area into up to hundred zones. Note, however, that if zones (or stations in a station-based system) are aggregated to fewer categories, such as "incentivized rental (return) station" and "neutral station" which allow the customer to precalculate incurring rental prices based upon these zone categories (Chung, Freund, and Shmoys 2018), this becomes practical.

• Third, for the provider, it is important *how* prices are determined. In this respect, our pricing approach is *anticipative* as it considers future profits based on dynamic programming. The papers using mathematical models largely rely on myopic optimization models. In addition, as we will discuss in-depth in Section 2, they cannot be applied to the problem we consider for various other reasons. The ways we design the anticipation allows to use historical data that is readily available in practice.

This pricing approach takes into account the typical characteristics of free-floating VSS (pick-up and drop-off possible at any location within the business area). However, given that station-based VSSs can be considered as a special case of free-floating systems, the approach is applicable to both. The pricing problem and our approach's practical relevance is ensured by, among other things, close cooperation with Share Now, Europe's largest car sharing provider operating in eight countries and 16 cities (Share Now 2021).

The contributions of our work are the following:

• We present a dynamic pricing approach for modern (free-floating) VSSs like car sharing and bike sharing, which is characterized by the three distinguished features mentioned above, that is, it is *customer-centric*, *origin-based*, and *anticipative*.

• We formulate the pricing problem underlying our approach as a dynamic program which considers stochasticity of the VSS. We show that regarding the action space at each stage of the dynamic program, only vehicles within walking distance need to be considered, such that online pricing becomes tractable. Based on the dynamic programming formulation, we develop an approximate dynamic programming solution method for the online pricing problem. The approach incorporates a nonparametric regression which allows the approximation of future profits based on historical data. This enables the precalculation of state-values such that the numerical operations of the online pricing problem can be reduced to a minimum.

• We conduct several computational studies, including sensitivity analyses as well as a case study based on Share Now data from the city of Vienna. We consider a discrete set of five (see Sections 5.4–5.6) or three price points (see Section 6) as VSS providers aim for a transparent and easy-to-communicate pricing mechanism. These studies show that our new dynamic pricing approach dominates all of the considered benchmarks in terms of realized profit, including state-of-the-art approaches from the literature. Further, these results are shown to be robust across the various considered settings and parameter variations, such as different VSSs sizes, overall demand levels, and customer preferences.

• We derive a number of relevant managerial insights from the computational studies. In particular, we show that our pricing approach is particularly effective when there is spatial variation in demand and that sophisticated anticipation of future states and profits is key. Another finding is that our pricing approach realizes higher profits compared with the benchmarks while maintaining the overall level of rentals, which is beneficial for service-oriented metrics of a VSS provider.

The remainder of the paper is organized as follows. In Section 2, we review the relevant literature. Section 3 formalizes the problem. Based on this, we develop the new dynamic pricing approach in Section 4. Section 5 contains the computational study with a discrete set of five price points. Section 6 presents the Share Now case study with the set of the three original price points. Section 7 concludes the paper and gives an outlook on future research. The appendix contains additional numerical results and a list of notation.

2. Related Literature

The literature on VSS optimization is broad, covering various types of systems, optimization problems, control approaches, and methodologies. General overviews on VSS optimization problems have been presented in survey papers on bike sharing (e.g., DeMaio 2009; Fishman, Washington, and Haworth 2013; Ricci 2015), car sharing (e.g., Jorge and Correia 2013; Ferrero et al. 2015a, b; Illgen and Höck 2019; Golalikhani et al. 2021a, b), and VSSs in general (e.g., Laporte, Meunier, and Wolfler Calvo 2015, 2018).

To define the scope of the following discussion of the related literature, please first note that VSSs are one specific type of shared mobility systems (Mourad, Puchinger, and Chu 2019), with the latter also comprising other modern mobility concepts such as ride-hailing (e.g., Wang and Yang 2019), ride-pooling (e.g., Ke et al. 2020), or dial-a-ride (e.g., Qiu, Li, Zhao 2018). VSSs substantially differ from these other mobility concepts in terms of the service offered as well as in terms of the provider's operating tasks. For example, in contrast to ride-hailing where the provider operates a platform that connects passengers and drivers in a two-sided market for individual trips, providers of VSSs operate in a one-sided market with a fixed supply (fleet) of (nonautonomous) vehicles. Further, for example, in dial-a-ride concepts, the provider is in charge of executing trips and, thus, strives for shared trips among passengers, whereas users in VSSs rent vehicles for individual trips and then drive these vehicles themselves. Due to these substantial differences in terms of mobility services and provider tasks, we exclude the other mobility concepts from our literature discussion and exclusively concentrate on VSSs. As a side remark, please note that the term "shared mobility system" is sometimes also used in a narrower sense as a synonym for VSS (e.g., Laporte, Meunier, and Wolfler Calvo 2015, 2018).

Within the literature on VSS pricing, we further focus on the closest related works that address *dynamic* pricing in VSSs. Dynamic—in contrast to *differentiated* (or static) pricing for VSSs like in Waserhole and Jost (2012) or Ren et al. (2019)—means that the pricing is performed on-line and depends on the system's current state (e.g., vehicle locations) (Agatz et al. 2013).

Our presentation of the closest related literature is organized in two groups according to the methodology proposed. First, in Section 2.1, we summarize approaches in which prices are determined through *business rules*. Second, in Section 2.2, we present those which are based on *optimization*. For both groups, we distinguish further into approaches that are *anticipative* in contrast to the *myopic* ones. Please note that some of these works have another main focus than pricing, but we discuss them specifically with regard to their pricing mechanism. In Section 2.3, we briefly refer to other literature streams that are related only in a broader sense. Finally, in Section 2.4, we position our approach in the closest related literature on dynamic pricing for VSSs.

2.1. Rule-Based Approaches

Rule-based approaches propose business rules to derive prices, for example by comparing endogenously given thresholds to the current state of the system. Among them, a group of anticipative approaches incorporates expected future states of the VSS into the pricing decision. Threshold values are usually compared with the ratio of future supply and demand at individual locations, which is derived from historical data and the system's current state (Brendel, Brauer, and Hildebrandt 2016; Dötterl et al. 2017). Wagner et al. (2015) in contrast consider exogenously given rules based on expected idle times. These three works have in common that they propose an overall framework for dynamic pricing, of which the rule-based pricing approach is one component.

Other works that use business rules propose myopic approaches. Bianchessi, Formentin, and Savaresi (2013) compare the number of vehicles at a location with the mean value across all locations to determine prices. Zhang, Meng, and David (2019) define prices by comparing the current number of vehicles with demand and propose a negative price that is linear in the undersupply of a rental's destination location. If there is no undersupply, the regular positive price applies. Barth, Todd, and Xue (2004) propose a system that, once it recognizes an imbalance, provides incentives for joint rides of independent customers in one car or splitting a party of customers into multiple cars. Mareček, Shorten, and Yu (2016) derive drop-off charges for vehicles depending on the intended destination location's distance to the nearest vehicle. Angelopoulos et al. (2016, 2018) propose two algorithms for promoting trips based on the priorities of vehicle relocations between locations. Chung, Freund, and Shmoys (2018) study the impact of an incentive program of New York City's bike sharing system and propose two myopic dynamic pricing approaches which are based on a performance metric that incorporates the estimated reduction of future out-of-stock events as well as the costs incurred by the incentives. Whereas the first considers each individual trip, the second determines decisions less frequently, that is, for several periods. Neijmeijer et al. (2020) is difficult to classify regarding methodology. They formulate a MIP that minimizes deviations of idle times from a desired value plus the costs of incentives, but in the empirical evaluations of a scooter sharing system, they test the effect of two possible discounts on vehicles' idle times.

2.2. Optimization-Based Approaches

Optimization-based dynamic pricing approaches are characterized by having a formal objective metric that is maximized (profit, revenue, rentals, service, etc.) or minimized (costs, unsatisfied demand, etc.). Methodologically, these approaches comprise those built on mathematical optimization like mixed-integer programming as well as learning methods that iteratively improve the objective, like reinforcement learning.

A couple of these optimization-based approaches use anticipative models. Singla et al. (2015) design a complete architecture of an operational incentive system which comprises, among other components, an" Incentive Deployment Schema" that decides whether to offer an alternative station with incentives or not with the objective to align future demand and supply. They evaluate using a real world survey as well as simulations. Pfrommer et al. (2014) propose an approach that uses quadratic programming and combines user-based and operator-based relocation. Prices are recalculated each period in a rolling horizon fashion. Ruch, Warrington, and Morari (2014) build on Pfrommer et al. (2014) and investigate simplified variants that can be used to benchmark more complex approaches. Di Febbraro, Sacco, and Saeednia (2012) aim at balancing supply and demand at all locations. They suggest alternative drop-off locations with a discount to customers. Assuming a given acceptance probability for these suggestions, a simulation evaluates the benefit for vehicle availability. Di Febbraro, Sacco, and Saeednia (2019) follow up on their earlier paper and formulate and test corresponding optimization models. Kamatani, Nakata, and Arai (2019) optimize thresholds by simulation-based optimization (Q-learning), while Clemente et al. (2017) propose a decision support system that uses a simulation-based heuristic (particle swarm optimization).

The remaining papers use myopic optimization. While they overall focus on user-based relocation, in one subsection Chemla et al. (2013) determine myopic prices period by period. They aim at a service-maximizing fleet distribution in bike-sharing systems through user-based relocation, where customer satisfaction is measured by successful and unsuccessful customer actions (available or non-available bike, empty or full rack). They use a linear program to determine the number of customers who change their travel plans because of the price incentive to reach the given target inventory of vehicles for each location. Two papers do not directly solve a mathematical model, but use it as a basis to develop a heuristic. Haider et al. (2018) formulate a bilevel program, where the upper level determines prices and minimizes vehicle imbalance, while the lower level represents the cost-minimizing route choice of customers. The problem is transformed into a single-level problem and a heuristic is proposed that iteratively adjusts prices (and, in contrast to the bilevel program, contains some anticipation). Wang and Ma (2019) consider the objective of keeping inventory within a certain range for a period. For this purpose, they define lower and upper thresholds for each location. The number of rentals from or to a location can be affected by pick-up and drop-off fees. They formulate a simple quadratic program to determine optimal dynamic pick-up and drop-off fees and solve it with a genetic algorithm. Kanoria and Qian (2019) define a myopic and location-based dynamic pricing algorithm for a trip-based VSS with time-variant arrival rates, without knowledge of arrival rates and with a finite time horizon. The algorithm needs the information about the queue length of origin and destination and a partitioning of the business area into zones.

2.3. Further Literature

There are several further literature streams which have some similarities with the considered problem and the applied methods, but which we do not discuss in detail. In particular, this concerns the determination of relocation prices with an auction process for VSSs (Ghosh and Varakantham 2017). Furthermore, we do not consider papers that do not describe the price setting process in detail. For example, Fricker and Gast (2016) show that user-based relocation is worthwhile, but they do not elaborate on how the prices are calculated.

Another stream of the literature investigates the steady state of stationary settings in mobility sharing concepts (including VSSs as considered in this work and ride-hailing applications), using techniques from closed-queuing networks. Waserhole and Jost (2016) maximize the number of trips taken, assuming time-invariant demand and zero travel time. Banerjee, Freund, and Lykouris (2022) extend this work by considering additional performance metrics and allowing nonzero travel time. Besbes, Castro, and Lobel (2021) focus on the single-location case of this setting. Benjaafar and Shen (2023) present an alternative solution approach with better performance bounds for several of the particular problems addressed in the above. Since all of these works consider stylized steady-state settings with time-invariant arrival-rates, exact knowledge of arrival-rate and an infinite time horizon. The derived pricing policies are differentiated (static) (see Benjaafar and Shen 2023), and, thus, not within the scope of our work. Further, they consider the systems on an aggregated level in which customer choices on a disaggregated level cannot be easily incorporated. Note that there are papers from this stream that differ from the aforementioned papers and use dynamic pricing. We have already mentioned an example (Kanoria and Qian 2019) of this in Section 2.2.

2.4. Positioning in the Literature

Our approach is positioned as follows in the literature on dynamic pricing for VSSs. With regard to methodology, the pricing approach we propose is anticipative and optimization-based. Thus, it belongs to the first group of works named in Section 2.2. It substantially differs from the above-named closest related works in particular by incorporating the concept of customer-centricity. As explained in Section 1 and illustrated by Figure 1, this means that prices are situation-specific for each customer, considering in particular a customer's location, the available vehicles within walking distance and other situational characteristics that influence the customer choice behavior, like walking distances to these vehicle. Using this situation-specific information for each customer has also been applied in shared mobility systems other than VSSs, for example, for mobility on demand services (Qiu, Li, Zhao 2018). To adopt this idea to the particularities of free-floating VSSs, we design our approach using disaggregate demand modeling which is capable of capturing the uncertain choice behavior. More specifically, our approach is designed to incorporate discrete choice models (Train 2009) in which the available vehicles within walking distance represent the discrete alternatives of the customer's consideration set that have individual choice probabilities. In the context of dynamic pricing in VSSs, we are not aware of any other comparable pricing approach.

As a consequence, our work also differs in terms of the mathematical modeling. Almost all of the closest related works which use mathematical models for determining prices use deterministic models (e.g., Chemla et al. 2013, Pfrommer et al. 2014, Wang and Ma 2019). In contrast, our approach is built upon a stochastic dynamic model and uses approximate dynamic programming with a nonparametric value function approximation to become tractable and scalable.

With regard to the type of the VSS, there are a few papers that also consider free-floating systems, albeit with rule-based approaches. Di Febbraro, Sacco, and Saeednia (2012) and Di Febbraro, Sacco, and Saeednia (2019) are the only two optimization-based approaches. Their vehicle-specific pricing is based on aggregated demand modeling.

Further, our approach differs from most of the literature in that prices are origin-based, meaning that they only depend on time and location of a rental's start. Such origin-based pricing, as applied by Share Now, is the most practicable variant for free-floating VSSs (see Section 1 or Soppert et al. 2022) for differentiated pricing, and thus a popular business practice. Other variants are destination-based prices (e.g., target-specific discounts as in Singla et al. 2015) and trip-based prices (or other incentives as in Chung, Freund, and Shmoys 2018), that depend on both origin and destination. Only Neijmeijer et al. (2020) in a rule-based approach consider origin-based pricing.

3. Model

In this section, we formally model the free-floating VSS provider's decision problem as a dynamic program. To do so, in Section 3.1, we give an overview on the sequence of events. In Sections 3.2–3.5, we explain how we choose the standard ingredients of dynamic programs: states, decisions, state transitions, and cost/revenue function. Based on this, we formally state the provider's optimization problem in Section 3.6.

3.1. Sequence of Events

We consider a free-floating VSS provider who operates a fleet of vehicles $\mathcal{C} = \{1, \ldots, C\}$ which is distributed spatially across a continuous business area. At any given point in time, a vehicle $i \in \mathcal{C}$ is either *idle* (standing available) or *in use* (currently rented). The provider seeks to maximize his profit over a finite planning horizon (e.g., one day) by pricing.

We follow the standard approach in the literature on pricing and revenue management by which this planning horizon is discretized into micro periods $t \in \{0, \ldots, T\} = \mathcal{T}$ and we have Δ micro periods per minute. These micro periods are w.l.o.g. so short that at most one customer request arrives and/or one rental terminates per period.

For every customer who opens the mobile application and requests prices, the provider has the ability to optimize and display prices. Hence, in this *online pricing problem*, the four *steps* within a micro period t are the following: (I) A customer may arrive, (II) if so, prices are determined by the provider, and (III) the customer chooses among the available vehicles under consideration of the offered prices. Finally, (IV) another moving vehicle that was previously rented (before period t) by another customer may return. One micro period of this process is illustrated in Figure 2, where decision nodes are represented as squares and stochastic nodes as circles.

Please note that in the following description, t denotes the current period, t' an earlier period, and t'' some future period, that is, $t' < t < t''$. Further, for random variables, following conventional notation, the indices t', t, and t'' reflect when these random variables *realize* and, thus, become known to the provider.

Step I: At the beginning of period t, the system is in state S_t which contains information about idle and driving vehicles. Now, at the stochastic node (circle) in Step I, with probability λ_t a *customer* k_t *arrives*, that is, she opens the mobile application and looks for available vehicles. The coordinates of the requesting customer's specific location in the business area are random variables (X_t^O, Y_t^O) which follow a given, time-dependent origin probability distribution $O(t)$. Realizations of these random variables, meaning the coordinates where a customer opens the mobile application, are denoted with (x_t^O, y_t^O).

Step II: The *provider optimizes the* $C \times 1$ *price vector* $\vec{p}_t$, visualized by the decision node (square).

Figure 2. Illustration of Dynamic Pricing Problem

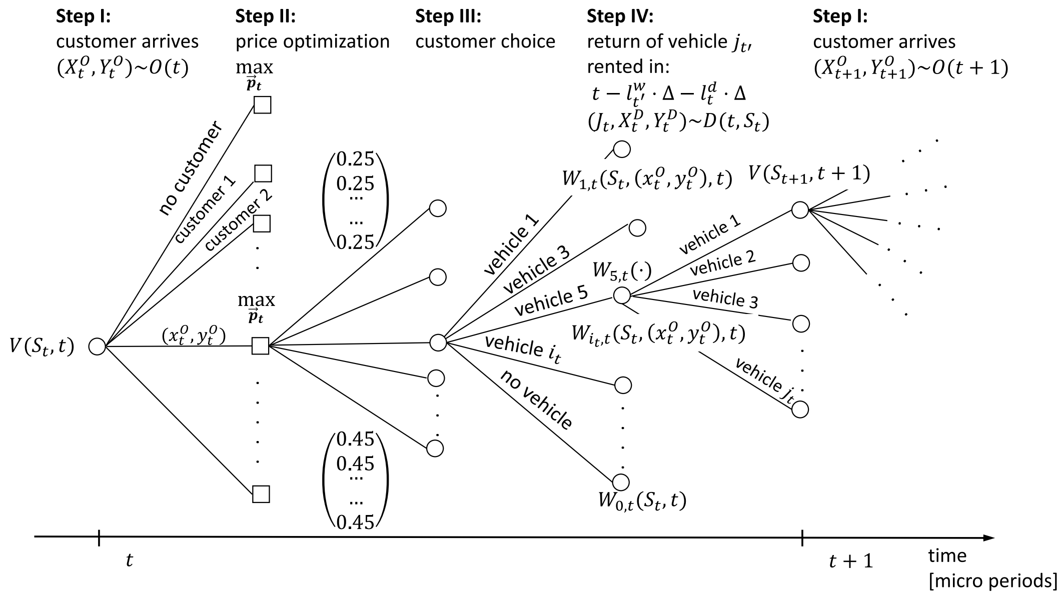

Step III: Based on these prices, the *customer* k_t *decides* whether and which vehicle to rent. The vehicle chosen is denoted by the random variable I_t with realizations i_t. The *customer choice behavior* is formalized as follows: Customers have a (fixed) maximum willingness to walk $\overline{d}$ (assumed to be known), meaning that a customer only considers idle vehicles i_t for which the walking distance $d_{i_t,t}$ between the customer's current location (x_t^O, y_t^O) and the idle vehicle is smaller than this radius, that is, the consideration set is $\mathcal{C}_{t,(x_t^O,y_t^O)} = \{i_t \in \mathcal{C} | d_{i_t,t} \le \overline{d} \wedge \tau_{i_t,t}^v = 0\}$ ($\tau_{i_t,t}^v$ is explained below in Section 3.2, it contains the information whether a vehicle is idle or in use). This is a well-known behavior of customers in VSSs and has been reported in multiple studies (e.g., Niels and Bogenberger 2017). More technically, we assume that the customer's choice probability $q_{i_t,t}$ for vehicle $i_t \in \mathcal{C}_{t,(x_t^O,y_t^O)}$ follows a known choice model and depends on the prices and the distances of the vehicles within reach, that is, $\mathcal{C}_{t,(x_t^O,y_t^O)}$ and $\vec{\boldsymbol{p}}_t$. The probability of not choosing any of the available vehicles is denoted by $q_{0,t}$. Note that our problem formulation is generic in this regard, meaning that *arbitrary* choice models providing these probabilities can be used. In the numerical studies, we apply a multinomial logit model (e.g., Train 2009). If she chooses a vehicle, the customer needs to walk there. This walking time l_t^w in minutes is a realization of L_t^W and depends on her distance to the vehicle. We assume a constant walking speed. As the chosen vehicle is stochastic, L_t^w is a random variable.

Step IV: Finally, in each micro period, w.l.o.g., at most one *rental may terminate*, which started in some period $t' < t$ *before* the current period t. More specifically, a customer who arrived at $t' = t - l_{t'}^w \cdot \Delta - l_t^d \cdot \Delta$ terminates her rental in t (see also Figure 3). Similar to the customer origin probability distribution $O(t)$, when and where a rental terminates is random. More technically, (J_t, X_t^D, Y_t^D) is a random variable which denotes that vehicle J_t (zero if none) is returned at location (X_t^D, Y_t^D). It follows a given *destination probability distribution* $D(t, S_t)$ that depends on the *state* S_t at the beginning of period t. The state definition is explained below. In particular, to capture typical traffic flow patterns, $D(t, S_t)$ may depend on where and when the currently driving customers have originated. Realizations of these random variables are denoted with j_t and (x_t^D, y_t^D). The driving/rental time l_t^d in minutes is a realization of the random variable L_t^d, which follows a distribution $\rho(S_{t'})$ and depends on the stochastic travel speed and the travel distance from pick-up to drop-off location (X_t^D, Y_t^D) of the vehicle, all unknown to the provider before the rental ends at micro period t.

3.2. State

The VSS's state $S_t = (\vec{\boldsymbol{x}}_t^c, \vec{\boldsymbol{y}}_t^c, \vec{\boldsymbol{\tau}}_t^c, \vec{\boldsymbol{x}}_t^v, \vec{\boldsymbol{y}}_t^v, \vec{\boldsymbol{\tau}}_t^v)$ at the beginning of period t consists of six vectors of dimension $C \times 1$, where the i_t-th element in each vector describes a property of the i_t-th vehicle of the fleet. The vectors $\vec{\boldsymbol{x}}_t^v$ and $\vec{\boldsymbol{y}}_t^v$ contain the coordinates of all vehicles of the fleet, that is, $x_{i_t,t}^v$ and $y_{i_t,t}^v$ $\forall i_t \in \mathcal{C}, t \in \mathcal{T}$, respectively. More specifically, for an idle vehicle they contain the coordinates of its location at the current micro period t. For a rented vehicle, they contain where the currently driving customer picked up the vehicle. The vector $\vec{\boldsymbol{\tau}}_t^v$ contains when the vehicle was picked up, with the value 0 indicating a vehicle standing idle. The vectors $\vec{\boldsymbol{x}}_t^c$, $\vec{\boldsymbol{y}}_t^c$, and $\vec{\boldsymbol{\tau}}_t^c$ describe when and where the respective

Figure 3. Sequence of Events

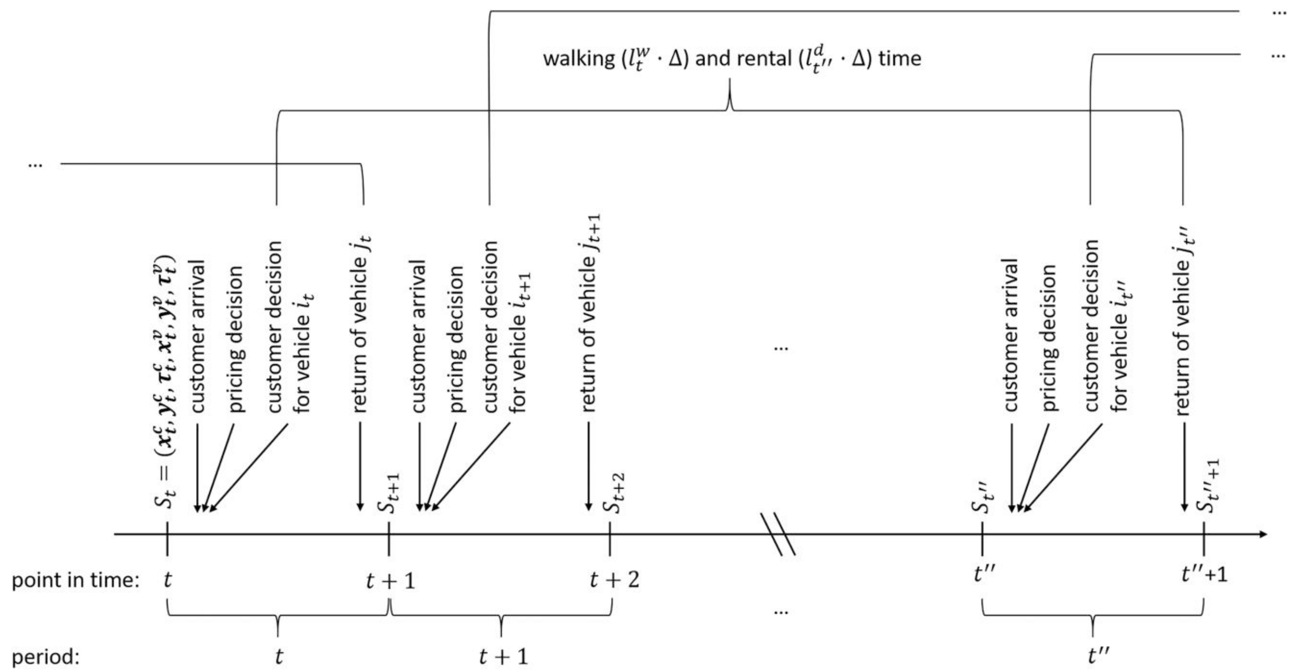

customers have requested the rental (i.e., for driving vehicles when and where the *customer* initially opened the mobile application).

3.3. Actions

Regarding the VSS provider's *pricing* decisions, we assume that the provider seeks to maximize profits by means of dynamic pricing. More precisely, when a customer opens the mobile application to look for available vehicles in micro period t, the VSS provider needs to optimize prices. As explained in Section 1, prices are *origin-based minute prices* and they are chosen from a discrete finite price set $\mathcal{M}$.

3.4. State Transitions

The transition function describes the evolution of the system from state S_t at the beginning of period t to state S_{t+1} at the beginning of period $t+1$. It depends on the current state S_t and the following *realizations* of random variables: the arriving customer's location (x_t^O, y_t^O), the chosen vehicle i_t (0 indicates the customer decides against renting a vehicle), and the returned vehicle j_t together with its return location (x_t^D, y_t^D) if (another) vehicle j_t rented before is returned (likewise, $j_t = 0$ indicates no vehicle is returned), that is,

$$S_{t+1} = S_{t+1}\big(S_t, (x_t^O, y_t^O), i_t, (x_t^D, y_t^D), j_t\big). \tag{1}$$

Please note that the probability distribution of the chosen vehicle I_t and therewith specific choices i_t depend on $\vec{p}_t$. Technically speaking, S_{t+1} is *probabilistically dependent* (Powell 2011, Chapter 3) on the pricing decision $\vec{p}_t$. By contrast, j_t does not depend on $\vec{p}_t$, but solely on which vehicles are currently in use and when and where these rentals started (stored in $\vec{x}_t^c, \vec{y}_t^c, \vec{\tau}_t^c, \vec{x}_t^v, \vec{y}_t^v, \vec{\tau}_t^v$).

The transitions of the state vectors are as follows. When a customer selects a vehicle i_t, the respective entries of the vectors $\vec{x}_t^c$ and $\vec{y}_t^c$ are filled with the customer's origin location (x_t^O, y_t^O), the *origin* time is updated to $\tau_{i_t,t}^c = t$ and the rental start time is set to $\tau_{i_t,t}^v = t + l_t^w \cdot \Delta$. When vehicle i_t is returned, $x_{i_t,t}^v$ and $y_{i_t,t}^v$ change to the destination location (x_t^D, y_t^D) and the corresponding $\tau_{i_t,t}^v$ and $\tau_{i_t,t}^c$ change back to 0.

3.5. Cost/Revenue Structure

For each minute that a vehicle i_t is rented, the provider collects the corresponding revenue $p_{i_t,t}$. During the rental, when the customer drives, the provider incurs variable costs per minute of c (e.g., for fuel). Following common practice, we assume that the short time frame the customer needs to walk to the vehicle is free of charge. Thus, the profit resulting from renting out vehicle i_t to the customer who arrived in period t and will return it in period t'' is given by $(p_{i_t,t} - c) \cdot l_{t''}^d$, where $l_{t''}^d$ is the realization of $L_{t''}^d \sim \rho(S_t)$.

3.6. Dynamic Programming Formulation

In this subsection, we model the problem described previously as a Markov decision process and state the corresponding Bellman equation:

$$V(S_t,t)=\lambda_t\cdot\mathop{\mathbb{E}}_{\substack{(X_t^O,Y_t^O)\\\sim O(t)}}\Bigg[\max_{\vec{p}_t}\overbrace{\sum_{i_t\in\mathcal{C}_{t,(x_t^O,y_t^O)}}q_{i_t,t}(\vec{p}_t)\cdot\Big((p_{i_t,t}-c)\cdot\mathop{\mathbb{E}}_{L^d_{t''}\sim\rho(S_t)}[L^d_{t''}]+\mathop{\mathbb{E}}_{\substack{(J_t,X_t^D,Y_t^D)\\\sim D(t,S_t)}}\big[V(S_{t+1}(S_t,(X_t^O,Y_t^O),i_t,(X_t^D,Y_t^D),J_t),t+1)\big]\Big)}^{\text{customer arrives and chooses a vehicle}}$$
$$+\overbrace{q_{0,t}(\vec{p}_t)\cdot\mathop{\mathbb{E}}_{\substack{(J_t,X_t^D,Y_t^D)\\\sim D(t,S_t)}}\big[V(S_{t+1}(S_t,0,0,(X_t^D,Y_t^D),J_t),t+1)\big]}^{\text{customer arrives and chooses no vehicle}}\Bigg]$$
$$+\overbrace{(1-\lambda_t)\cdot\mathop{\mathbb{E}}_{\substack{(J_t,X_t^D,Y_t^D)\\\sim D(t,S_t)}}\big[V(S_{t+1}(S_t,0,0,(X_t^D,Y_t^D),J_t),t+1)\big]}^{\text{no customer arrives}} \tag{2}$$

with the boundary condition $V(S_T,T)=0\ \forall S_T$. The Bellman equation recursively calculates the optimal expected profit from future rentals $V(S_t,t)$ for being in state S_t at the beginning of period t. Each micro period t corresponds to a stage in this dynamic program. In the following, we explain how the four steps (I–IV) within each stage (= micro period) are represented in (2).

In the first and the second line of (2), a customer arrives (Step I) with probability λ_t at a location (x_t^O,y_t^O) and in this case, the optimal price vector $\vec{p}_t$ for all available vehicles is determined (Step II). The following customer choice process has different potential outcomes (Step III): With probability $q_{i_t,t}(\vec{p}_t)$ (first line), vehicle i_t is chosen, and, in expectation a profit of $(p_{i_t,t}-c)\cdot\mathbb{E}_{L^d_{t''}\sim\rho(S_t)}[L^d_{t''}]$ is obtained. Another vehicle j_t that was rented before micro period t may be returned at location (x_t^D,y_t^D) (Step IV) and the system evolves to the next state in micro period $t+1$ where expected future profit is $V(S_{t+1}(S_t,(x_t^O,y_t^O),i_t,(x_t^D,y_t^D),j_t),t+1)$.

With probability $q_{0,t}(\vec{p}_t)$ (second line), no vehicle is chosen. Nonetheless a vehicle j_t may be returned at location (x_t^D,y_t^D), and the system evolves into the state in micro period $t+1$ with expected future profit $V(S_{t+1}(S_t,0,0,(x_t^D,y_t^D),j_t),t+1)$. The third line of the Bellman equation considers the case—occurring with probability $(1-\lambda_t)$—in which no customer arrives, so again, only a vehicle may be returned at location (x_t^D,y_t^D). Hence, we have the same expected future profit as in the second line of the equation.

If a customer arrives in micro period t, the optimal prices need to be calculated. When doing so, the provider is aware of the customer's coordinates (x_t^O,y_t^O), which are thus deterministic for the problem of pricing. Obviously, profit-maximizing prices in (2) are given by

$$\vec{p}_t^{\,*}=\mathop{\arg\max}_{\vec{p}_t}\sum_{i_t\in\mathcal{C}_{t,(x_t^O,y_t^O)}}q_{i_t,t}(\vec{p}_t)\cdot\Big((p_{i_t,t}-c)\cdot\mathbb{E}_{L^d_{t''}\sim\rho(S_t)}[L^d_{t''}]+W_{i_t}(S_t,(x_t^O,y_t^O),t)\Big)+q_{0,t}(\vec{p}_t)\cdot W_0(S_t,t) \tag{3}$$

with

$$W_{i_t,t}(S_t,(x_t^O,y_t^O),t)=\mathop{\mathbb{E}}_{\substack{(J_t,X_t^D,Y_t^D)\\\sim D(t,S_t)}}[V(S_{t+1}(S_t,x_t^O,y_t^O),i_t,(X_t^D,Y_t^D),J_t),t+1)]\quad\forall i_t\in\mathcal{C}_{t,(x_t^O,y_t^O)}, \tag{4}$$

$$W_{0,t}(S_t,t)=\mathop{\mathbb{E}}_{\substack{(J_t,X_t^D,Y_t^D)\\\sim D(t,S_t)}}[V(S_{t+1}(S_t,0,0,(X_t^D,Y_t^D),J_t),t+1)]. \tag{5}$$

Since the provider knows the customer's coordinates, he also knows her consideration set $\mathcal{C}_{t,(x_t^O,y_t^O)}$ and only the prices for the idle vehicles $i_t\in\mathcal{C}_{t,(x_t^O,y_t^O)}$ within reach of the current customer at location (x_t^O,y_t^O) need to be optimized, as the choice probabilities only depend on them (see Section 3.3 and Figure 1(a)). Thus, instead of the $C\times 1$ vector $\vec{p}_t$, a smaller price vector $\vec{p}_{t,(x_t^O,y_t^O)}$ with only $|\mathcal{C}_{t,(x_t^O,y_t^O)}|\times 1$ entries (a subset of the entries of the original price vector) needs to be optimized. More specifically, this new $\vec{p}_{t,(x_t^O,y_t^O)}$ contains the entries i of $\vec{p}_{t,(x_t^O,y_t^O)}$, for which $i\in\mathcal{C}_{t,(x_t^O,y_t^O)}$. Thus, the action space reduces from pricing all idle vehicles of the fleet to a handful and the online pricing problem becomes

$$\vec{p}^{\,*}_{t,(x_t^O,y_t^O)}=\mathop{\arg\max}_{\vec{p}_{t,(x_t^O,y_t^O)}}\sum_{i\in\mathcal{C}_{t,(x_t^O,y_t^O)}}q_{i_t,t}(\vec{p}_{t,(x_t^O,y_t^O)})\cdot\Big((p_{i_t,t}-c)\cdot\mathop{\mathbb{E}}_{L^d_{t''}\sim\rho(S_t)}[L^d_{t''}]+W_{i_t,t}(S_t,(x_t^O,y_t^O),t)\Big)+q_{0,t}(\vec{p}_{t,(x_t^O,y_t^O)})\cdot W_{0,t}(S_t,t). \tag{6}$$

The dynamic program considers all steps (Step I-Step IV) in a micro period, that is, at the beginning of the micro period, it is not yet clear if and where a customer arrives. Therefore, the outer expectation in (2) is over the location $(X_t^O,Y_t^O)\sim O(t)$ where the customer arrives. However, the supplier determines the prices of the reachable vehicles when the customer has already arrived (Step II). Thus, when deciding on prices, the provider knows the location (x_t^O,y_t^O) of the customer and the expectation over $(X_t^O,Y_t^O)\sim O(t)$ can be dropped in (6).

4. Solution Method

In this section, we describe the solution method we propose for the considered problem. First, in Section 4.1, we derive a convexity result that allows us to obtain an efficient linear reformulation of the pricing problem. Then, in Section 4.2, we develop our approximate dynamic programming solution method. Section 4.3 contains the proposed nonparametric value function approximation, including a description how historical data are used.

4.1. Convexity of the Pricing Problem and Linear Reformulation

When solving the online pricing problem (6) by complete enumeration, the number of calculations that must be performed for each micro period t is of the order $\mathcal{O}(|\mathcal{M}|^{|\mathcal{C}_{t,(x_t^O,y_t^O)}|})$, implying that the calculation time in one single stage of the dynamic program increases exponentially with the number of reachable vehicles $|\mathcal{C}_{t,(x_t^O,y_t^O)}|$ and polynomially with the number of price levels $|\mathcal{M}|$. Further, given that the state space grows exponentially across the multiple stages, an alternative solution method is required to efficiently solve larger instances.

Under the assumption that demand follows a multinomial logit model (see Section 3.1) as it does in our numerical studies (see Section 5.4.2), the pricing problem at stage t (i.e., the one-stage optimization problem) can be formalized by the following fractional program (see, e.g., Davis, Gallego, and Topaloglu 2013), using decision variables $z_{i_t,m}$ that equal 1 if the provider selects price level m for vehicle i_t:

$$\max \frac{1}{B_t + \sum_{i_t \in \mathcal{C}_{t,(x_t^O,y_t^O)}} \sum_{m\in\mathcal{M}} A_{i_t,t,m}\cdot z_{i_t,t,m}} \cdot \left[\sum_{i_t \in \mathcal{C}_{t,(x_t^O,y_t^O)}} \sum_{m\in\mathcal{M}} \left((p_{i_t,m} - c)\cdot \mathop{\mathbb{E}}_{L^d_{t''}\sim\rho(S_t)}[L^d_{t''}] + W_{i,t}(S_t,(x_t^O,y_t^O),t) \right) \cdot A_{i_t,t,m}\cdot z_{i_t,t,m} + B_t \cdot W_{0,t}(S_t,t) \right] \tag{7}$$

s.t.

$$\sum_{m\in\mathcal{M}} z_{i_t,t,m} = 1 \qquad \forall i_t \in \mathcal{C}_{t,(x_t^O,y_t^O)} \tag{8}$$

$$z_{i_t,t,m} \in \{0,1\} \qquad \forall i_t \in \mathcal{C}_{t,(x_t^O,y_t^O)}, m\in\mathcal{M} \tag{9}$$

In the objective function (7), $A_{i_t,t,m}$ denote given attraction values for each vehicle i_t and price level m, capturing the attractiveness of the corresponding alternative to the customer. The attraction values correspond to $e^{(u_{i_t,t})}$, with $u_{i_t,t}$ referring to the utility of alternative i_t in the underlying multinomial logit demand model. The attraction value for the no-choice alternative is B_t and corresponds to $e^{(u_{0,t})}$. Constraints (8) ensure that exactly one price level m is selected for each reachable vehicle i_t, and Constraints (9) define the variables $z_{i_t,t,m}$ as binary.

The structure of the fractional program (7)–(9) is identical to that of the product line and price selection problem analyzed by Chen and Hausman (2000). They show that its objective function is strictly quasi-convex in $z_{i_t,m}$. Further, the coefficient matrix of the model's Constraints (8) is totally unimodular (also see Lemma 2 in Chen and Hausman 2000). Consequently, the binary Constraints (9) can be LP-relaxed (i.e., $0 \le z_{i_t,t,m} \le 1\ \forall i_t \in \mathcal{C}_{t,(x_t^O,y_t^O)}, m\in\mathcal{M}$) and, due to the convexity of the resulting model, every local maximum is also a global maximum, allowing to solve the problem by standard nonlinear programming codes.

It is even possible to use linear programming to solve the problem, as we can linearize the objective of the relaxed fractional program by applying a Charnes-Cooper transformation (see, e.g., Stancu-Minasian 1997). The idea is to substitute the reciprocal of the denominator, that is, $\left(B_t + \sum_{m\in\mathcal{M}}\sum_{i_t\in\mathcal{C}_{t,(x_t^O,y_t^O)}} A_{i_t,t,m}\cdot z_{i_t,m}\right)^{-1}$ in (7) by a new variable $v_t > 0$ and ensure correctness by additionally imposing the constraint:

$$B_t\cdot v_t + \sum_{i_t\in\mathcal{C}_{t,(x_t^O,y_t^O)}} \sum_{m\in\mathcal{M}} A_{i_t,t,m}\cdot \hat{z}_{i_t,t,m} = 1, \tag{10}$$

where the variable $\hat{z}_{i_t,t,m}$ substitutes $v_t \cdot z_{i_t,t,m}$.

After performing the latter substitution also in the existing constraints of the relaxed fractional program and removing redundant constraints, the following equivalent, linear reformulation for period t is obtained:

$$\max \sum_{i_t\in\mathcal{C}_{t,(x_t^O,y_t^O)}} \sum_{m\in\mathcal{M}} \left((p_{i_t,m} - c)\cdot \mathop{\mathbb{E}}_{L^d_{t''}\sim\rho(S_t)}(L^d_{t''}) + W_{i_t}(S_t,(x_t^O,y_t^O),t) \right)\cdot A_{i_t,t,m}\cdot \hat{z}_{i_t,t,m} + (B_t\cdot v_t)\cdot W_{0,t}(S_t,t) \tag{11}$$

s.t.

$$B_t\cdot v_t + \sum_{i_t\in\mathcal{C}_{t,(x_t^O,y_t^O)}} \sum_{m\in\mathcal{M}} A_{i_t,t,m}\cdot \hat{z}_{i_t,t,m} = 1 \tag{12}$$

$$\sum_{m\in\mathcal{M}} \hat{z}_{i_t,t,m} = v_t \qquad \forall i_t \in \mathcal{C}_{t,(x_t^O,y_t^O)} \tag{13}$$

$$\hat{z}_{i_t,t,m} \ge 0 \qquad \forall i_t \in \mathcal{C}_{t,(x_t^O,y_t^O)}, m\in\mathcal{M} \tag{14}$$

$$v_t \ge 0 \tag{15}$$

4.2. Approximate Dynamic Programming Solution Method

The state space as well as the outcome space of the dynamic program (2) depend on the coordinates of the arriving customer. Since the customer can arrive anywhere within the continuous business area, these coordinates are continuous. The same holds for the return location of a vehicle. Thus, state and outcome space are of infinite size. As a consequence, the dynamic program cannot be solved exactly (curse of dimensionality, see, e.g., Powell 2011, Chapter 1.2).

We use approximate dynamic programming to obtain a tractable solution method and exploit the fact that we are only interested in the price decisions $\vec{\boldsymbol{p}}^*_{\boldsymbol{t},(x^O_t,y^O_t)}$, that is, the solution of (6).

In particular, we approximate the values $W_{i_t,t}, W_{0,t}$ of the stochastic nodes immediately *after a customer's decision* (Step III) and before a potential return of a vehicle becomes known (see Figure 2). This allows to reduce the size of the online pricing problem tremendously by only optimizing *one* period explicitly while still taking into account the customer choice behavior. Graphically, this corresponds to "trimming" the decision tree in Figure 2 after Step IV and capture the parts that are cut away by $W_{0,t}$ and $W_{i_t,t}$. The challenge, however, is to find accurate approximations $\tilde{W}_{i,t}$, $\tilde{W}_{0,t}$ for $W_{i_t,t}$, $W_{0,t}$, respectively. Our approximation is based on the key simplification that V and, thus, W is *additive* in the values of all vehicles. The intuition here is that, since the overall revenue obtained is composed of the revenues realized by the individual vehicles, adding up the revenue-to-come from a certain point in time on for all vehicles—the vehicle values—yields the state value. Clearly, this is a simplification because the actual spatial vehicle distribution never perfectly matches the one observed in history such that potential interdependencies between vehicle values for the current vehicle distribution are neglected. However, this additivity assumption has a very favorable property with regard to runtime and, thus, implementation in practice: In order to calculate the optimal prices, we no longer need to consider *all* available vehicles of the fleet, but only the *reachable* ones. This is because the vehicles out of the customer's reach remain idle for any choice outcome and, thus—technically speaking—form a constant term in the online pricing problem (6) that can be neglected. Our results show that this value function approximation works well and that this policy provides in a small instance, the same results (= prices) as the theoretically optimal policy (TINY, see Sections 5.2 and 5.3).

For the same reason, we neglect the currently moving vehicles. Clearly, this also is an approximation, because, in reality, unreachable and moving vehicles obviously impact future revenues and may also alter a decision—for example, if a vehicle is part of a large agglomeration of vehicles.

Clearly, a vehicle's value (= expected future profit until the end of the time horizon) depends on whether it remains standing *idle* at its current location or whether it *departs* to another location through a rental. Hence, for a certain vehicle i_t, we denote these approximate vehicle-specific values as $\tilde{w}^{idle}_{i_t}$ and $\tilde{w}^{depart}_{i_t}$, respectively. With this assumption, the approximated values $\tilde{W}_{i_t,t}$ and $\tilde{W}_{0,t}$, thus, can be obtained by

$$W_{i_t,t} \approx \tilde{W}_{i_t,t} = \sum_{j_t \in \mathcal{C}_{t,(x^O_t,y^O_t)} \setminus \{i_t\}} \tilde{w}^{idle}_{j_t,t} + \tilde{w}^{depart}_{i_t,t} \quad \forall i_t \in \mathcal{C}_{t,(x^O_t,y^O_t)}, \tag{16}$$

$$W_{0,t} \approx \tilde{W}_{0,t} = \sum_{j_t \in \mathcal{C}_{t,(x^O_t,y^O_t)}} \tilde{w}^{idle}_{j_t,t}. \tag{17}$$

The idea in (16) is that the value of the state after vehicle i_t has been chosen ($W_{i_t,t}$) is approximately the sum of the values of the remaining idling vehicles from the consideration set $\mathcal{C}_{t,(x^O_t,y^O_t)}$, plus the value of the departing (= chosen) vehicle i_t. Accordingly in (17), the state value when no vehicle was chosen ($W_{0,t}$) is approximately the sum of all idling vehicles from $\mathcal{C}_{t,(x^O_t,y^O_t)}$.

Hence, the online pricing problem (6) solved in Step II becomes

$$\begin{aligned}\vec{\boldsymbol{p}}^*_{\boldsymbol{t},(x^O_t,y^O_t)} \approx \underset{\vec{\boldsymbol{p}}_{\boldsymbol{t},(x^O_t,y^O_t)}}{\arg\max} & \sum_{i_t \in \mathcal{C}_{t,(x^O_t,y^O_t)}} q_{i_t,t}(\vec{\boldsymbol{p}}_{\boldsymbol{t},(x^O_t,y^O_t)}) \cdot \Big((p_{i_t,t} - c) \\ & \cdot \mathbb{E}_{L^d_{t''} \sim \rho(S_t)}[L^d_{t''}] + \tilde{W}_{i_t,t}\Big) \\ & + q_{0,t}(\vec{\boldsymbol{p}}_{\boldsymbol{t},(x^O_t,y^O_t)}) \cdot \tilde{W}_{0,t}.\end{aligned} \tag{18}$$

4.3. Nonparametric Value Function Approximation

In this subsection, we describe the specific approach for obtaining the values $\tilde{w}^{idle}_{i_t,t}$ and $\tilde{w}^{depart}_{i_t,t}$. We first give an overview of our approach in Section 4.3.1. Then, we present the details of data selection and the kernel function used in Section 4.3.2.

4.3.1. General Idea. Based on the current time t and the location, a vehicle i_t is evaluated. This evaluation is performed for two different cases. First, for the case that vehicle i_t is chosen by the customer, departs and is dropped off at the destination after the rental time. Second, when the customer does not choose the vehicle and, thus, the vehicle remains idle at the current location. For the first case ($\tilde{w}^{depart}_{i_t}$), the value indicates the expected value of vehicle i_t after it has been chosen and parked by the customer. For the second case ($\tilde{w}^{idle}_{i_t}$), the value indicates the expected value of the vehicle after the customer has not chosen the vehicle and it is still idle at the same location.

More technically, the approximate vehicle values $\tilde{w}^{idle}_{i_t,t}$ and $\tilde{w}^{depart}_{i_t,t}$ for a vehicle i_t are determined by a nonparametric value function approximation (see Powell 2011, Chapter 8.4 for an introduction to this technique). Building on this, our approach is as follows. The values $\tilde{w}^{idle}_{i_t,t}$

and $\tilde{w}^{depart}_{i_t,t}$ are calculated as weighted averages across corresponding data points k from historical and/or simulated data that reflects current system behavior. That is, for an idle vehicle, $\tilde{w}^{idle}_{i_t,t}$ is a weighted average of corresponding idle vehicle values $\hat{w}^{idle}_k$ in the data and $\tilde{w}^{depart}_{i_t,t}$ is a weighted average of corresponding departing vehicle values $\hat{w}^{depart}_k$ in the data. The values $\tilde{w}^{idle}_{i_t,t}$ and $\tilde{w}^{depart}_{i_t,t}$ depend on the location of the vehicle i_t and the time at which the vehicle i_t to be valuated is located. These approximate vehicle values are location and time dependent, that is, they depend on a subset of the state. More specifically,

$$\tilde{w}^s_{i_t,t} = \sum_{k \in \mathcal{K}^s_{i_t}} \kappa^s_{k,i_t} \cdot \hat{w}^s_k \quad \forall i_t \in \mathcal{C}_{t,(x^O_t,y^O_t)}, s \in \{idle,\ depart\}, t \in \mathcal{T} \tag{19}$$

where κ^{idle}_{k,i_t} and κ^{depart}_{k,i_t} are the weights that capture how "similar" a specific data point k is to vehicle i_t (see next subsection for details). The sets $\mathcal{K}^{idle}_{i_t,t}$ and $\mathcal{K}^{depart}_{i_t,t}$ represent the sets of observations relevant to approximate the value of vehicle i_t (see next subsection for details).

To explain the process of obtaining these values $\hat{w}^{idle}_k$ and $\hat{w}^{depart}_k$ from data, we assume for the following illustration w.l.o.g. that the problem's time horizon is one day and that we dispose of data that only comprises one specific date. For each vehicle, we know over the day when and where it was standing idle, when it departed, and how much profit the corresponding rental generated, as well as when and where each rental terminated. Figure 4 illustrates such "paths" in the historical data, consisting of idle times (thick blue/red lines) and rentals (thin blue/red arrows) exemplarily for two vehicles (red and blue). For now, consider only the temporal dimension on the horizontal axis. The remainder of this figure (with the spatial dimension on the vertical axis) is explained in the next subsection. Thus, for any given point in time, we can determine the current status of each vehicle from this data, and the required values $\hat{w}^{idle}_k$ and $\hat{w}^{depart}_k$ capture the—loosely speaking—profit the vehicle generates from this point in time onwards until the end of the day.

Obviously, robustness improves with increased amount of data available, and, thus, one would combine data from multiple comparable historical/simulated dates, for example from multiple identical days of the week. Then, regarding a data's timestamps, only the time (and not the date) is relevant and observations from different dates are considered as different vehicles. Further, an implicit assumption to note here is that spatial vehicle distributions, demand patterns and, thus, vehicle values are similar for comparable days and, even more important, to the situation when the pricing approach is applied. This, of course, is the case only to a certain extent, because the above-mentioned are endogenous and an important point of pricing is to influence the vehicle distribution. To address this issue, the underlying database may be iteratively updated (see also the numerical results in Sections 5.3 and 5.7.2)

The described nonparametric value function approximation has two decisive benefits for practice. First, historical data can readily be used. Second, the approximate vehicle values can continuously be precomputed such that they do not need to be determined in the moment the pricing problem (18) needs to be solved.

4.3.2. Data Collection, Data Selection and Kernel Function.

The remaining part to fully specify our approach is the determination of the sets $\mathcal{K}^{idle}_{i_t,t}$ and $\mathcal{K}^{depart}_{i_t,t}$ relevant for the evaluation of vehicle i_t from the sets of all data points $\mathcal{K}^{idle}$ and $\mathcal{K}^{depart}$, as well as the weights $\kappa^{idle}_{i_t,k}$ and $\kappa^{depart}_{i_t,k}$.

Figure 4. (Color online) Illustration of Historical Data Considered for Evaluation of Vehicle i_t

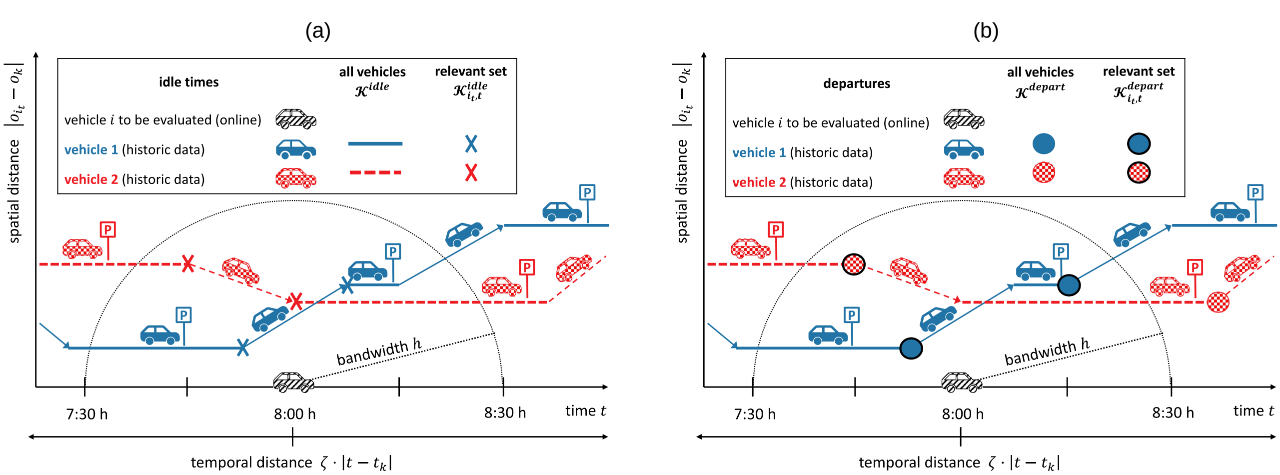

Notes. (a) Idle vehicles. (b) Departing vehicles.

Regarding *departing vehicles*, the set of all data points $\mathcal{K}^{depart} = \{(\hat{w}_k^{depart}, o_k, t_k)\}$ consists of collected data points k with location o_k and time t_k of the departure event. The value $\hat{w}_k^{depart}$ is the profit earned by this vehicle *after* the rental that started at t_k (this is necessary for consistency with (18)) until the considered time frame (one day in our case). Thus, a VSS provider can simply reconstruct this data from the observed historical vehicle movements, which are collected anyway for invoicing.

As mentioned above, one central idea is to approximate values for departing vehicles based on "similar" data points. Since all events in the free-floating VSS are characterized by a certain location and time, it is reasonable to integrate the spatial as well as the temporal dimension in the metric that measures "similarity". This implies that the vehicle values $\tilde{w}_{i_t,t}^{idle}$ and $\tilde{w}_{i_t,t}^{depart}$ are time- and location-dependent. To determine $\mathcal{K}_{i_t}^{depart}$ for a vehicle i_t whose value is to be approximated (with location $o_{i_t,t} = (x_{i_t,t}^v, y_{i_t,t}^v)$ and at time t), we define the following filter:

$$\mathcal{K}_{i_t,t}^{depart} = \left\{(\hat{w}_k^{depart}, o_k, t_k) \in \mathcal{K}^{depart} \,\middle|\, \zeta \cdot |t - t_k| + |o_{i_t,t} - o_k| \le h\right\}. \tag{20}$$

where $|t - t_k|$ is some *temporal distance*, $|o_{i_t,t} - o_k|$ is some *spatial distance*, ζ is a scaling parameter, and h is a bandwidth. This idea of a spatiotemporal "similarity" and a bandwidth h can be thought of as a (stretched) circle. It is illustrated in Figure 4(b). The black (diagonally striped) vehicle at a certain location at 8:00 h is to be evaluated. The departure event data points are the red and blue circles. According to the filter, only data points (red and blue circles within the semicircle) within radius h (black dotted) are to be considered and marked by a black circle.

For the *idle vehicles*, this step is slightly more complex, because collected data points on idle vehicles $\mathcal{K}^{idle} = \{(\hat{w}_k^{idle}, o_k, \bar{t}_k)\}$ refer to the time *intervals* $\bar{t}_k$ when the vehicles stood idle (the horizontal thick lines in Figure 4(a)). For an interval $\bar{t}_k$, data point k has the future value $\hat{w}_k^{idle}$ that equals the profit earned by this vehicle after the interval until the end of the horizon (there is obviously no profit during the interval). To determine distance in time, we need to compare these intervals with the point in time t of the vehicle to evaluate. To do so, from each interval, we consider the point in time closest to t. More formally, the set of relevant observations to evaluate an idle vehicle i_t (depicted as red and blue crosses in the figure) is

$$\mathcal{K}_{i_t,t}^{idle} = \left\{(\hat{w}_k^{idle}, o_k, t_k) \,\middle|\, \exists (\hat{w}^{idle}, o_k, \bar{t}_k) \in \mathcal{K}^{idle} \wedge t_k = \arg\min_{\bar{t}_k''' \in \bar{t}_k} |t_k''' - t| \wedge \zeta \cdot |t - t_k| + |o_{i_t,t} - o_k| \le h\right\}. \tag{21}$$

Next, the weights κ_{k,i_t}^s for every historical/simulated data point $k \in \mathcal{K}_{i_t,t}^s$ $\forall s \in \{idle, depart\}$ are determined with a *kernel function* K. As described above, a scaling ensures that the weights sum to one. In particular, we use

$$\kappa_{i_t,k}^s = \frac{K_{i_t,k}^s}{\sum_{l=1}^{|\mathcal{K}_{i_t,t}^s|} K_{i_t,l}^s} \quad \forall k \in \mathcal{K}_{i_t,t}^s, s \in \{idle, depart\}. \tag{22}$$

The unscaled weights $K_{i_t,k}^s$ $\forall s \in \{idle, depart\}$ can be determined using various kernel weighting functions, including the Gaussian, Uniform, Epanechnikov, or Bi-Weight kernel weighting functions. As kernel weighting function, we use the following Epanechnikov kernel function (Powell 2011, Chapter 8.4.2)

$$K_{i_t,k}^s = \frac{3}{4} \cdot \left(1 - \left(\frac{d_{i_t,k}}{h}\right)^2\right) \quad \forall k \in \mathcal{K}_{i_t,k}^s, s \in \{idle, depart\} \tag{23}$$

with

$$d_{i_t,k} = \sqrt{(\zeta \cdot (t - t_k))^2 + (|o_{i_t,t} - o_k|)^2} \quad \forall k \in \mathcal{K}_{i_t,k}^s, s \in \{idle, depart\}. \tag{24}$$

5. Computational Studies

In this section, we evaluate the developed dynamic pricing approach in comparison with different benchmark approaches. To that end, we consider two different groups of settings. The first group only consists of one artificial, small setting (denoted TINY) that allows analytically solving the dynamic program by backward recursion to obtain the optimal prices (= policy) as a benchmark. The second group with its three settings SMALL, MEDIUM, and LARGE allows to evaluate the approaches in more realistic instances.

The new pricing approach and the benchmark approaches are described in Section 5.1. Regarding TINY, Section 5.2 describes the setup and Section 5.3 describes the evaluation procedure and the results. The setup of the second group (i.e., SMALL, MEDIUM, LARGE) is described in Section 5.4, followed by the main results in Section 5.5 and additional sensitivity analyses in Section 5.6. In Section 5.7, we briefly analyze variations of the developed solution approach.

5.1. Pricing Approaches

We compare our developed pricing approach with nine benchmark approaches (see Table 1). First we describe the *customer-centric* pricing approaches:

- C-ANT: Our *customer-centric and anticipative* pricing approach determines dynamic prices for each customer by considering current and future (approximate) state values (see Section 4).
- OPT: Calculation of the optimal price for each arriving customer by backward recursion. As usual in

Table 1. Overview of Pricing Approaches

		Focus of dynamic pricing		
		Customer-centric	Location-based	Uniform pricing
Foresight	Anticipative	C-ANT, C-HEUR, OPT	L-ANT	BASE, LOW, HIGH
	Myopic	C-MYOP	L-MYOP, RUBA	

dynamic programming, this pricing approach is only feasible for very small instances.

- C-MYOP: *Myopic* version of C-ANT without anticipation: $\tilde{w}^{idle}_{i_t,t} = \tilde{w}^{depart}_{i_t,t} = 0$ for all $i_t \in \mathcal{C}_{t,(x^O_t, y^O_t)}$, resulting in $\tilde{W}_{i_t,t} = \tilde{W}_{0,t} = 0$ for all $i_t \in \mathcal{C}_{t,(x^O_t, y^O_t)}$.
- C-HEUR: *Heuristic* improvement of C-MYOP. Instead of $\tilde{w}^{idle}_{i_t,t} = 0$, $\tilde{w}^{idle}_{i_t,t}$ equals the average profit per expected rental duration across all vehicles for all $i_t \in \mathcal{C}_{t,(x^O_t, y^O_t)}$. More specifically, no distinction is made in the valuation of the idle vehicles: $\tilde{w}^{idle}_{i_t,t} = \tilde{w}^{idle}\ \forall i_t$. Thus, in the price optimization (18), a rental is no longer "for free" (no opportunity cost) as in C-MYOP, but $\tilde{w}^{idle}_{i_t,t}$ now reflects that if vehicle i_t is not rented now, it obtains in expectation a profit of $\tilde{w}^{idle}_{i_t,t}$ during the expected rental time $\mathbb{E}_{L^d_{t''} \sim \rho(S_t, S_{t,d})}[L^d_{t''}]$ because it may be chosen by another customer already in the next period.

In addition, we consider some pricing approaches that are *location-based*. Here we differentiate between approaches that are based on partitioning the business area into tiles and approaches that are based on business rules.

- RUBA: *Rule-based* pricing approach, in which the business area is partitioned into 1 km × 1 km *tiles* that can be thought of as stations, as it is common in the literature. To obtain prices for the vehicles in each tile, we follow the approach of Bianchessi, Formentin, and Savaresi (2013) who compare the number of vehicles in each tile to the average number of available vehicles in the entire business area. If the number of vehicles in a tile falls below the average number of available vehicles, the price of the vehicles in the tile is increased and the magnitude of the increase depends on the severity of the imbalance. Vice versa, if the number of idle vehicles rises above the average number of available vehicles, the price is decreased. Whereas in the original approach continuous prices are used, we require discrete prices for the considered problem. Thus, in a further step, the calculated continuous prices are discretized by rounding to the nearest price point.
- L-ANT: In a first step, the business area is also partitioned into 1 km × 1 km tiles. This location-based pricing approach determines dynamic prices for every tile and 1h-period (e.g., 1 a.m.–2 a.m., 2 a.m.–3 a.m., etc.). At the beginning of each period, a Faure sequence is used to generate multiple realizations of artificial customer arrivals for each tile. C-ANT is then applied to determine the prices for the available vehicles for each artificial customer. The prices for the vehicles in a tile are then averaged and rounded to the next price point. All vehicles located in the tile obtain this price. This benchmark is anticipative but is not customer-centric, since it does not use the situation-specific information of the customer's location as in C-ANT. Therefore, by comparing the results of L-ANT to C-ANT's, the value of considering the location of the customer in anticipative pricing is isolated, that is, the importance of customer-centric pricing can be quantified.
- L-MYOP: *Myopic* version of L-ANT without anticipation: This means that instead of C-ANT we use C-MYOP to determine the prices for the vehicles in the different tiles. This benchmark is neither anticipative nor customer-centric. Therefore, compared with C-MYOP, this benchmark can be used to measure the impact of considering the location of the customer in myopic pricing. Compared with the benchmark above, we see the value of anticipation in non-customer-centric (= location-based) pricing.

Regarding the relation between C-ANT and L-ANT, note that C-ANT dominates L-ANT in the (theoretical) case that the state value approximation is exact. In this case, C-ANT is the optimal policy. However, since state value approximations are not exact in general, both approaches are heuristics.

Last, we consider the following *uniform* pricing approaches. As they use only one price, they do not require a pricing decision from the provider.

- BASE: Constant uniform pricing, where $p_{i,t}$ is the median price from the set of price points (following our industry partner, we also call it *base* price) for all vehicles. Due to its wide adoption over all VSS types, this pricing approach can be considered as the de facto standard in practice.
- LOW: Constant uniform pricing, where $p_{i,t}$ is the *lowest* price for all $i \in \mathcal{C}$, $t \in \mathcal{T}$.
- HIGH: Constant uniform pricing, where $p_{i,t}$ is the *highest* price for all $i \in \mathcal{C}$, $t \in \mathcal{T}$.

Each pricing approach is evaluated in $n = 500$ simulation runs with common random numbers and we report average values.

5.2. Setup Artificial Setting

This artificial setting (denoted TINY) is set up such that it is still possible to optimally solve the dynamic program

Figure 5. (Color online) Locations and Stations in Artificial Setting in the First Micro Period (TINY)

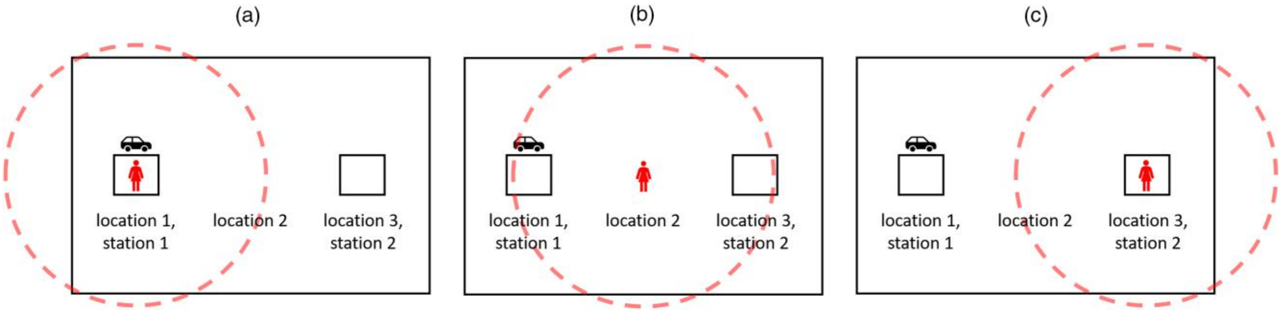

Notes. (a) Customer arrives at location 1. (b) Customer arrives at location 2. (c) Customer arrives at location 3.

by backward recursion. To that end, a spatial discretization of the business area is required such that the state space becomes finite. Thus, the setting can be considered as a station-based VSS. It consists of two *stations* where vehicles can be rented and returned, two micro periods, and one vehicle. The customer can arrive (i.e., open the app) at three possible *locations*, of which two (location 1 and 3) coincide with the two stations (see Figure 5). If the customer arrives at one station, the other is too far to walk to. Both stations are within reach of the customer if she arrives at location 2. The customer arrival probability for both micro periods is 80% for location 1, 10% for location 2, and 5% for location 3. With 5%, no customer arrives. The provider can set two different prices, that is, the high price (0.46 €/min) or the low price (0.18 €/min).

The customer choice probability depends on the distance between customer and vehicle (positive correlation) and the price (negative correlation): If the vehicle is out of reach, the choice probability is 0. If the customer arrives at the station where the vehicle is located, she chooses the vehicle with 80% for the low price and with 30% for the high price. If the customer arrives at location 2 and the vehicle is located at one of the two stations, then she chooses the vehicle at the low price with 60% and at the high price with 20%.

The rental destination probabilities are set as follows: If the customer chooses the vehicle at station 1 (station 2), she terminates the rental at station 1 with probability 5% (50%) and at station 2 with 95% (50%). For simplicity, the walking time is neglected in this setting ($l_1^w = l_2^w = 0$) and the rental time is always one micro period ($l_1^d = l_2^d = 1/\Delta$). At the beginning of the first micro period, the vehicle is at station 1. Figure 1 in Online Appendix A.1 shows the corresponding decision tree. The optimal policy always sets high prices in the first micro period and low prices in the second micro period.

5.3. Simulation Procedure and Results Artificial Setting

In this subsection, we compare the results of LOW, C-ANT, and C-MYOP with OPT. For the evaluation of each approach, $n = 10{,}000$ runs are simulated.

Regarding C-ANT, the required historical data are generated by simulating 10,000 runs with LOW. As the only vehicle is always at station 1 at the beginning of period 1, we obtain two vehicle values. In case the vehicle departs from station 1, it has an expected future profit (i.e., in period 2) of $\tilde{w}_1^{depart} = 0.0242$ €. In case it remains idle at station 1, we have $\tilde{w}_1^{idle} = 0.125$ €. As the vehicle is never at station 2 in period 1, no values are available or needed. The values for period 2 are $\tilde{w}_1^{idle} = \tilde{w}_2^{idle} = \tilde{w}_1^{depart} = \tilde{w}_2^{depart} = 0$, as no rentals can occur beyond this period.

Figure 6(a) depicts the profits as relative improvements over LOW. It shows that C-MYOP generates the same profit as LOW and that C-ANT generates the same profit as OPT, that is, over 20% more profit than

Figure 6. (Color online) Performance Indices and Vehicle Value for Station 1 and Period 1

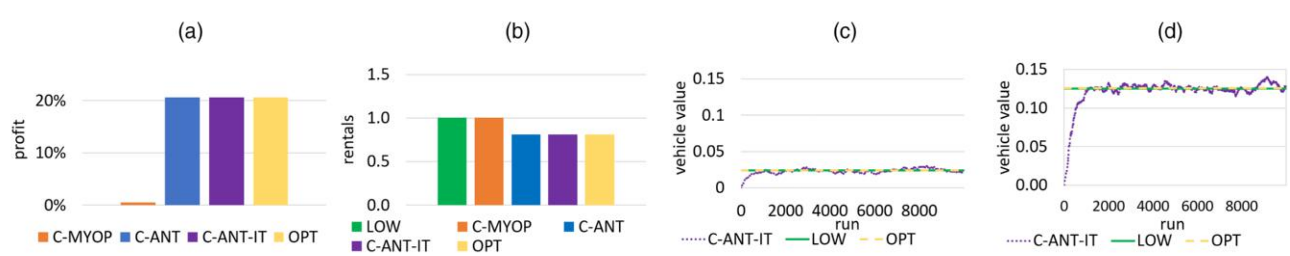

Notes. (a) Profit impr. over LOW. (b) Rentals. (c) $\tilde{w}_1^{depart}$ generated by. (d) $\tilde{w}_1^{idle}$ generated by.

LOW. Further analyses (not shown here) reveal that both LOW and C-MYOP only set the low price. C-ANT and OPT also yield the same policy with a frequency of 44% for the low price and 56% for the high price. This is also reflected in the number of rentals, which are about 20% lower for C-ANT and OPT (Figure 6(b)). Overall, this shows that C-ANT can indeed yield the theoretically optimal policy. The optimality achieved by C-ANT in this small example cannot be generalized. However, it shows that in principle C-ANT can achieve optimality.

For C-ANT, the vehicle values were determined based on LOW runs. However, as discussed in Section 4.3.1, vehicle values reflect customer behaviour and are, thus, influenced by prices. If past pricing considerably deviates from current practice, iteratively updating the vehicle values based on new data are an alternative. To demonstrate this, we initialize the values $\tilde{w}^{idle}$ and $\tilde{w}^{depart}$ with 0, perform a simulation run, and then adjust the vehicle values. In doing so, we used a constant stepsize and weighted the previous vehicle value with 0.995 and the new vehicle value from the recent simulation run with 0.05. The vehicle values for the departing and idle vehicle at station 1 in period 1 in comparison with the vehicle values obtained by runs with LOW (as before) and OPT are shown in Figure 6, (c) and (d).

Although the initialization is quite bad, we observe the values to quickly converge to the value obtained with the LOW or the OPT pricing. After 10,000 iterations we obtain the values $\tilde{w}^{idle} = 0.1263$ € and $\tilde{w}^{depart} = 0.025$ €, which—in this example—yield the same policy as described above. We will further investigate the value of iterative updating in a more realistic setting (Section 5.7.2).

Thus, we could show two things by comparing C-ANT with OPT: First, C-ANT can generate the same results as OPT if the vehicle values are well estimated. Second, we demonstrated how to iteratively update the vehicle values based on new data.

5.4. Setup Realistic Settings

5.4.1. Settings and Parameters. We consider three *settings* that differ mainly in the size of the business area and the number of vehicles (SMALL, MEDIUM, and LARGE). The area of the SMALL setting has a size of 9 km^2 and is equipped with 36 vehicles (MEDIUM 16 km^2 and 64 vehicles, LARGE 25 km^2 and 100 vehicles, all areas are square). The planning horizon is one day and at the beginning, all vehicles are randomly uniformly distributed across the business area. The demand patterns we use replicate what is observed in practice. Demand intensity varies over the course of the day with two peaks (Figure 2 in Online Appendix A.2, see, e.g., Reiss and Bogenberger 2016). Furthermore, in line with practice, there is also a spatial variation of demand, for example, between the city center and peripheral areas. This is modeled by the density (pdf) of the origin probability distribution $O(t)$ (see Section 3.1), which is exemplarily shown for all settings and two different times (8:00 h, 16:00 h) in Figure 7. Density here means that the demand for the respective time period is spread over the size of a certain area, for example, the city center (the very center of Figure 7(a)). In other words, integrating the demand densities displayed in Figure 7(a) over the whole business area yields a value of 1 (note that density values are rounded) which corresponds to the normalized demand

Figure 7. (Color online) Exemplary Density (pdf) of Customer Arrivals (Demand) Over Business Area

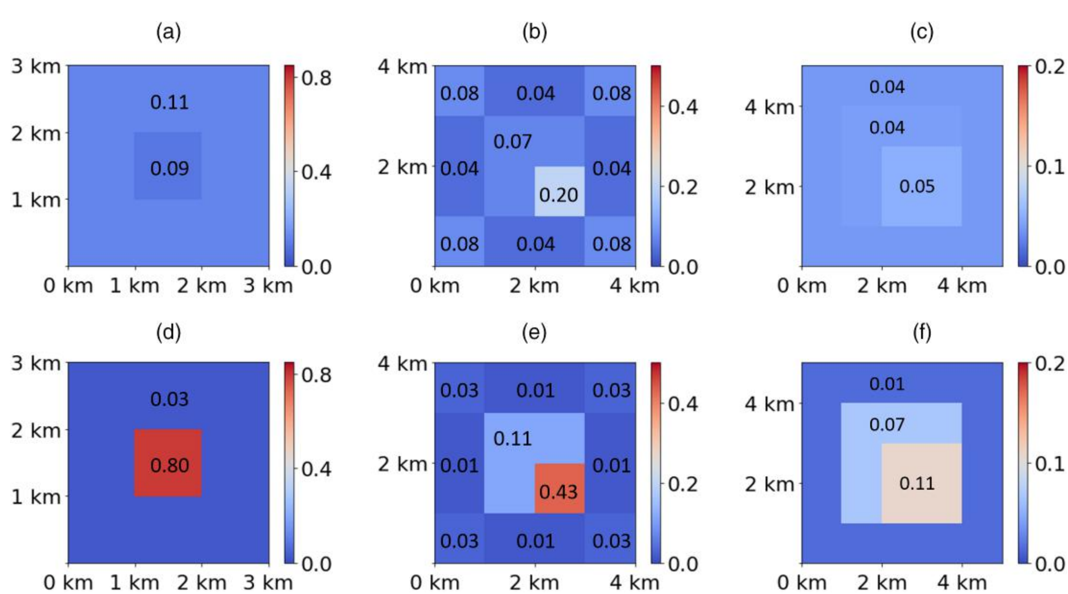

Notes. (a) 8:00 h, SMALL. (b) 8:00 h, MEDIUM. (c) 8:00 h, LARGE. (d) 16:00 h, SMALL. (e) 16:00 h, MEDIUM. (f) 16:00 h, LARGE.

value from Figure 7 for a certain period. The destination probability distribution for a customer who departed in the center is exemplarily shown for all settings and at two different times in Figure 6 in Online Appendix A.4.

Each of the three settings is examined for three different overall demand levels, which differ in the *demand-supply-ratio* (DSR). The DSR is the maximum period demand (second peak) divided by the fleet size and we consider the values $\in \left\{\frac{1}{3}, \frac{2}{3}, 1\right\}$ by scaling demand appropriately.

The other parameters are constant throughout all three settings: We choose a presumably small number of possible prices ($M = 5$ price points). The reason is related to the practitioners' important pursuit for a transparent and easy to communicate pricing mechanism. In particular, a pricing mechanism with a small number of prices in comparison with infinite possibilities in a continuous range is much more transparent and easier to communicate to customers. These price points (prices for short) $p^m \in \mathcal{M}$ are predefined with regard to typical prices in practice: We chose a *base price* per minute of $p^3 = 0.35$ €/min and a price difference of 0.10 €/min to the so-called *low* and *high* prices, so that $p^1 = 0.25$ €/min and $p^5 = 0.45$ €/min. The other prices are $p^2 = 0.30$ €/min and $p^4 = 0.40$ €/min. Variable costs are $c = 0.07$ €/min. We calculate the travel time of a rental by drawing the speed from a realistic distribution for urban traffic. We then get the rental/driving time l_t^d as the product of the driving speed and the distance $d_{o_{i_t},(X_t^D, Y_t^D)}$. Further, we assume a willingness to walk of $\overline{d} = 500$ m.

The parameter $\overline{d}$ is assumed to be known and its estimation was performed based on two analyzes. In the first analysis, we conducted a literature search with regard to the maximum walking distance. For example, Singla et al. (2015) and Herrmann, Schulte, and Voß (2014) show that below 20% and at most 20.43%, respectively, of the respondents were willing to walk more than 500 m. In the second analysis, we examined the data on customer choices from our practice partner which, amongst others, contains information on the distances to reachable vehicles. In particular, we analyzed two aspects. First, we analyzed the share of customers walking more than 500 m to a vehicle. This share is below 2%. Second, we examined the choice situations in which all of the available vehicles were located more than 500 m away. In this situation, the share of customers who chose one of these vehicles was close to 0. Thus, by combining the insights of literature and data, we decided to set the maximum walking distance to $\overline{d} = 500$ m.

Furthermore, we only consider the ten closest vehicles, as we observed from looking at the Share Now data that on average there are 4.3 vehicles within walking distance and the customer in average chooses the 2.1 nearest vehicle. Furthermore, an analysis of the distribution of the number of available vehicles shows that 90% of customers have seven or fewer vehicles available upon arrival. Only about 4% have ten or more vehicles available (see Figure 5).

5.4.2. Customer Choice Model.

As described in Section 3.1, a customer at position (x_t^O, y_t^O) chooses among the vehicles $i \in \mathcal{C}_{t,(x_t^O, y_t^O)}$ within reach and may also decide not to rent (no choice option), which is denoted by $i_t = 0$. In the numerical study, customer choice behavior follows a multinomial logit model (see e.g., Train 2009, Chapter 3). Accordingly, the choice probabilities $q_{i_t,t}$ depend on the alternatives' deterministic utilities $u_{i_t,t}$ for the customer (see Figure 4(b) in Online Appendix A.3):

$$q_{i_t,t} = \frac{e^{u_{i_t,t}}}{\sum_{n \in \mathcal{C}_{t,(x_t^O, y_t^O)} \cup \{0\}} e^{u_{n,t}}}. \tag{25}$$

The deterministic utility $u_{i_t,t}$ of a vehicle i_t depends on its price $p_{i_t,t}$ and its distance to the customer $d_{i_t,t}$ (see Figure 4(a) in Online Appendix A.3):

$$u_{i_t,t} = \beta^{price} \cdot p_{i_t,t} + \beta^{distance} \cdot d_{i_t,t}. \tag{26}$$

The no-choice option has utility $u_{0,t} = ASC^{NoChoice}$ where $ASC^{NoChoice}$ stands for the alternative-specific constant for the no-choice option. These assumptions imply homogeneous customers and that customers decide solely based on current circumstances (myopic behavior). In particular, they do not act strategically (see, e.g., Gallego and Van Ryzin 1997; Talluri and Van Ryzin 2004, chapter 5.1.4; Gönsch et al. 2013, for discussions of strategic or forward looking customers).

The parameters for the one choice model which is fit across all locations are estimated with a maximum-likelihood estimation based on 200,000 observations of mobile application openings. Technically, we used the Python package PandasBiogeme 3.2.10 (Bierlaire 2020). Among others, we used the likelihood ratio test and the Akaike information criterion to compare different model specifications. We tested attributes for the nearest vehicle, the time of the day, the different vehicle types used in practice, and displayed discount badges. However, their impact was minimal, so, in the end, we then implemented the above utility function with the two attributes with the biggest impacts on utility, that is, price and distance, as one might intuitively expect. We can state that 60 meters of walking distance reduction approximately correspond to a price reduction of 0.10 €/min.

5.5. Main Results Realistic Settings

5.5.1. Profit.

We first discuss profit, whose maximization is the objective of the optimization problem and obviously the most important metric from the provider's perspective. The results for all three settings and DSRs are summarized in Figure 8. All profits are presented as relative profit improvements over the BASE pricing approach.

Figure 8. (Color online) Profit Improvement Over BASE

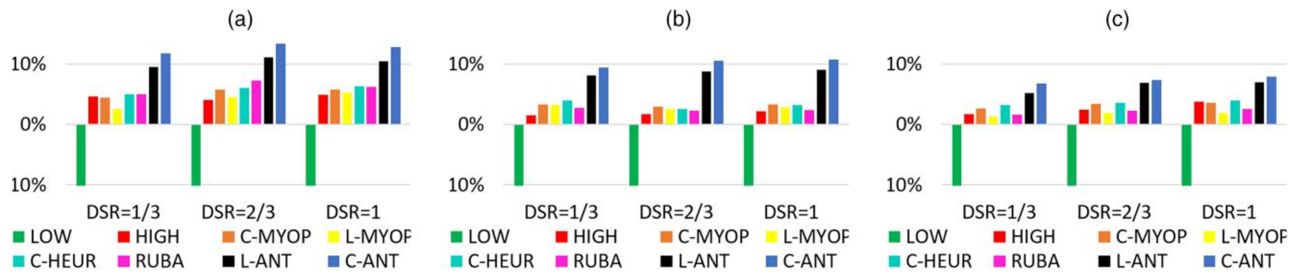

Notes. (a) SMALL. (b) MEDIUM. (c) LARGE.

We observe that C-ANT clearly provides the highest profit for all settings and DSRs. Compared with BASE, C-ANT shows profit improvements of up to 13.4%. The improvement over LOW is 22.9 to 32.7 percentage points, over HIGH 4.2 to 9.3, over C-MYOP 4 to 7.6, over C-HEUR 3.6 to 8, over RUBA 5.1 to 8.4, over L-MYOP 5.5 to 9.3, over L-ANT 0.5 to 2.3 percentage points. By contrast, LOW performs much worse than BASE. L-ANT always performs second best across all settings and DSRs with an improvement of 5.1–10.5 percentage over BASE. For the benchmarks HIGH, C-MYOP, L-MYOP, C-HEUR, and RUBA, there is no clear order.

The fact that C-ANT generates up to 7.6 percentage points higher profits than C-MYOP shows that including anticipation has substantial value. However, the comparison of C-ANT, L-ANT and C-HEUR shows that it is important *how* anticipation is done. A simple constant valuation for $\tilde{w}_i^{idle}$ as done in C-HEUR is not effective, since in some cases, for example, in the SMALL setting with DSR = 2/3, C-MYOP performs better than C-HEUR. The comparison between C-ANT and L-ANT shows the additional profit that C-ANT gains, because it is customer-centric, that is, it has the advantage to take situation-specific customer information (location of customer and distance to each vehicle within walking distance) into account. C-ANT is up to 2.3 percentage points better than L-ANT.

We conclude that C-ANT dominates all other pricing approaches with regard to profit and that its anticipative and customer-centric design is key for the performance—with about three quarters of C-ANT's improvement over simple heuristics such as L-MYOP coming from anticipation and one quarter from being customer-centric.

5.5.2. Prices.

Now, we compare the prices resulting from the different pricing approaches. To that end, we consider results from the SMALL setting with all three DSRs. Figure 9 illustrates the average price across all areas during the day (we left out LOW, BASE, and HIGH that set constant prices), and Figure 7 in Online Appendix A.6 shows the average price for different parts of the business area from C-ANT and C-MYOP. Figure 10 shows the relative frequency of prices for all approaches. The results for MEDIUM and LARGE are depicted in Online Appendix A.7.

Regarding the average price curves (Figure 9), we observe two different groups of approaches. For both groups there is a similar pattern in the average prices, but the prices of the nonanticipative (myopic) pricing approaches fluctuate less than the anticipative pricing approaches. For example in Figure 9(a), the average price of C-MYOP fluctuates between 0.35 and 0.40 €/min, whereas the average price of C-ANT fluctuates between 0.34 and 0.42 €/min. These results can be explained as follows: The anticipatives approaches (C-ANT, L-ANT, C-HEUR) attempt to incentivize the use of the vehicles in certain parts of the business area during the morning such that they become available in other parts with high demand later during the day. This explains the comparably low average prices of

Figure 9. (Color online) Average Prices Over the Course of the Day (SMALL)

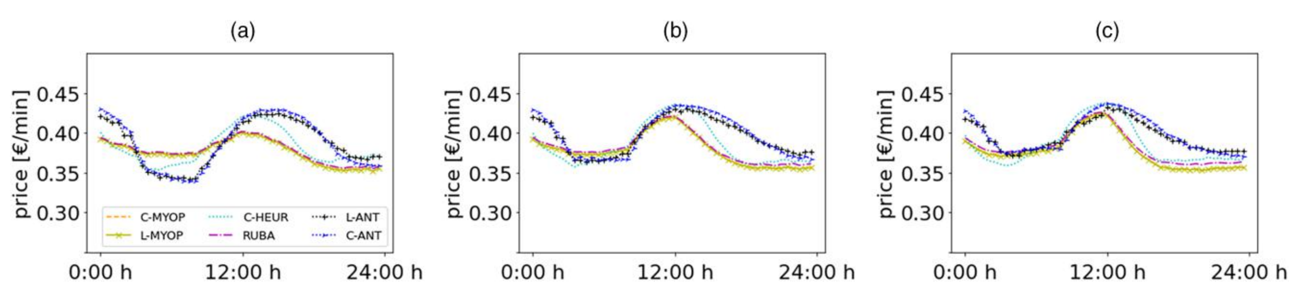

Notes. (a) DSR = 1/3. (b) DSR = 2/3. (c) DSR = 1.

Figure 10. (Color online) Relative Price Frequency (SMALL)

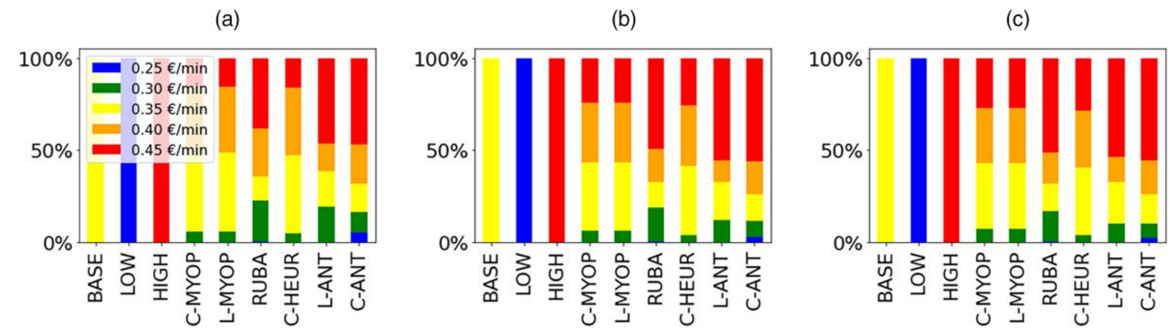

Notes. (a) DSR = 1/3. (b) DSR = 2/3. (c) DSR = 1.

C-HEUR and especially C-ANT and L-ANT during the morning. On the other hand, the myopic approaches do not consider futures states and profits and, thus, set higher average prices during the morning hours which are more profitable in the short term but less profitable in the long term, as the profit results above show.

The difference in terms of pricing between anticipative and myopic approaches becomes even more apparent when considering the temporal *and* spatial differences of prices by C-ANT and C-MYOP in Figure 7 in Online Appendix A.6. C-MYOP sets relatively high average prices in all parts of the business area throughout the entire day. In contrast, C-ANT varies prices in time and space. For example, in all peripheral parts, relatively low prices are set in the morning, while prices in the center at the same time are comparably high. Again, the purpose of this is to incentivize customers to drive vehicles from the outer areas to the center. In the center there is always high demand, so the price here is always quite high.

The discussed differences in price patterns between the pricing approaches can also be seen at the aggregate level by comparing the frequency of prices in Figure 10. While C-MYOP and L-MYOP do not set the lowest price, C-ANT also sets lowest prices. L-ANT, on the other hand, does also not use the lowest prices, but the share of second lowest prices is higher than for L-MYOP. Thus, these lowest prices (C-ANT) and the higher share of lower prices (C-ANT and L-ANT) cannot be motivated by myopic considerations, but only by regard to future profits. The lower profits by the higher share of lower prices in the morning are overcompensated by profits from later rentals. This also works in the opposite direction: C-ANT and L-ANT also choose highest prices more often then C-MYOP and L-MYOP.

We conclude that lowest prices and the higher share of lower prices, especially during morning hours, can be used as an incentive for customers and allow to generate higher profits at highest prices later during the day when the vehicle distribution is better aligned with the demand. This only works when future profits are taken into account like it is done by C-ANT or L-ANT, because vehicle values are approximated more accurately—in particular their dependence on both location and time is considered. Taking the customer's location into account allows C-ANT to better tailor incentives to the customer.

5.5.3. Rentals. For the analysis of the rentals, we consider Figure 11, which shows the average hourly rentals for the different pricing approaches over the course of the day for different DSRs in the SMALL setting. The respective results for MEDIUM and LARGE are depicted in Online Appendix A.7.

The rental curves resemble the demand curve (Figure 2 in Appendix A.2) in that there is a minimum of rentals in the early morning and a maximum in the afternoon. As expected, the number of rentals increases in the DSR

Figure 11. (Color online) Rentals Over the Course of the Day (SMALL)

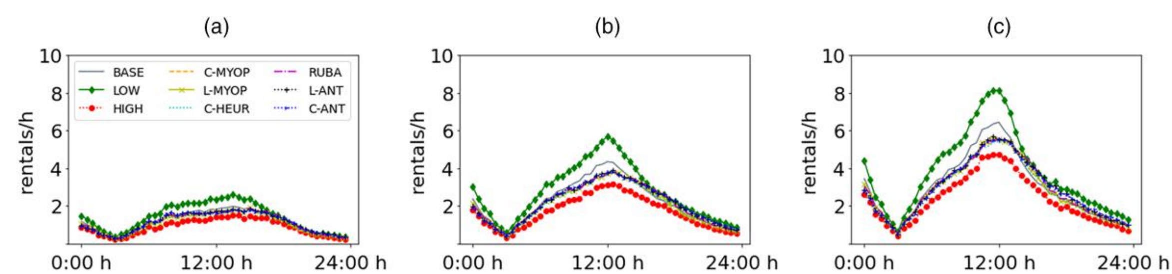

Notes. (a) DSR = 1/3. (b) DSR = 2/3. (c) DSR = 1.

Figure 12. (Color online) Vehicle Distribution for Different Pricing Approaches (SMALL, DSR = 2/3)

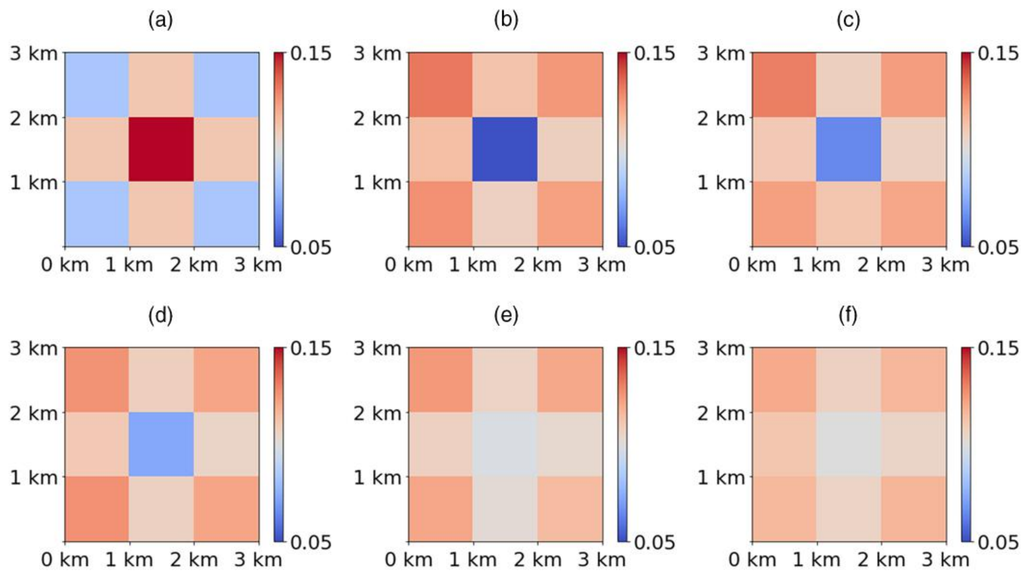

Notes. (a) Initial, all approaches. (b) 17:30, BASE. (c) 17:30, C-MYOP. (d) 17:30, RUBA. (e) 17:30, L-ANT. (f) 17:30, C-ANT.

and the number of rentals is lowest (highest) for HIGH (LOW).

The rental curves for C-MYOP and L-MYOP are very similar to the rental curves for C-ANT and L-ANT. Thus, although they all have very similar aggregated rentals, C-ANT and L-ANT manage to obtain considerably higher profits. It is also interesting that the rental curve for C-ANT is very similar to the rental curve of L-ANT, but C-ANT generates a considerably higher profit.

The rental curves of BASE lie above the ones for C-ANT, L-ANT, C-MYOP, and L-MYOP for all DSRs. Three important insights can be drawn thereupon. First, myopic pricing (C-MYOP, L-MYOP) leads to a significant decrease in the number of rentals compared with BASE, but an improvement in profit. Second, including anticipation, as in C-ANT and L-ANT compared with the C-MYOP and L-MYOP, leads to similar number of rentals and *at the same time* to an increase in profit. Thus, besides the increased profit, C-ANT and L-ANT arguably provides the same service to customers with a higher profit. Third, including customer's location as in C-ANT compared with L-ANT leads to a similar number of rentals, but a considerably higher profit.

5.5.4. Service-Oriented Metrics. In this section, we consider three metrics: the share of customers who have (1) *no available vehicle within reach*, the (2) *walking effort* by considering the share of customers who choose the nearest, second nearest, etc. vehicle, and the comparison of the (3) *spatial distribution of vehicles* at the beginning of the day and immediately before the afternoon peak. All results are obtained in the SMALL setting.

Regarding (1), Figure 13(a) shows that—as expected—the share of customers who cannot find an reachable vehicle increases for each pricing approach, as the overall

Figure 13. (Color online) Different Results

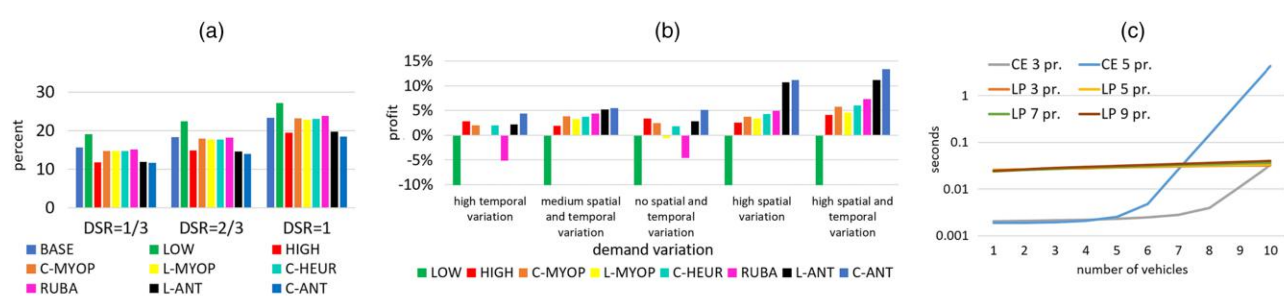

Notes. (a) Share of customers with no vehicle within walking distance (SMALL, DSR = 2/3). (b) Profit improvement over BASE (SMALL, DSR = 2/3). (c) Computational time of linear program (LP) and complete enumeration (CE) (SMALL, DSR = 2/3) for the reachable vehicles within walking distance.

demand level (DSR) increases. For each demand level, this share is highest (lowest) for LOW (HIGH). It is also noticeable that for the two pricing approaches C-ANT and L-ANT, this share is considerably lower than for BASE, C-MYOP, and L-MYOP. Overall, the results suggest that the service level improves with anticipation (C-ANT and L-ANT), because it leads to a lower share of customers who do not have a reachable vehicle within the walking distance.

Regarding (2), Figure 8 in Online Appendix A.6 shows that most customers (>50%) take the nearest vehicle and then, to a decreasing extent, the second nearest, the third nearest, and then the others. Since these shares for the nearest (second nearest, etc.) vehicle hardly differ between the different pricing approaches, we conclude that the applied pricing approach has only a minor influence on the customers' walking distance. This implies that total customer experience does not suffer.

Regarding (3), Figure 12 shows the spatial vehicle distributions over the business area at the beginning of the day as well as at 17:30 (immediately before the afternoon peak) for BASE, C-MYOP, RUBA, L-ANT, and C-ANT in the SMALL setting with DSR = 2/3. For illustrative purposes, we partitioned the 3 km × 3 km business area into 9 tiles and calculated the average share of idle vehicles in each tile. While at the start of the day 22% of the vehicles are in the center of the business area (Figure 12(a)) where demand is strongest in the afternoon, at 17:30 h with BASE that number declined to 5%. By contrast, C-ANT manages to have 10% of vehicles in the center which explains the higher availability observed above.

5.6. Sensitivity Analysis Realistic Settings

We consider two aspects in more detail. We examine the robustness of the above results regarding changes in demand preferences which, for example, vary across cities and countries. To that end, first, we examine whether the dominance of C-ANT discussed above holds if the spatial and temporal variation of the demand intensity is less pronounced (Section 5.6.1). Second, we analyze the impact of customer preference variation regarding price sensitivity and disutility from walking in Section 5.6.2.

A common *standard demand pattern* serves as a basis for parameter variations throughout the sensitivity analysis. We use the demand pattern of the SMALL setting for DSR = 2/3, depicted in the top right of Figure 3 in Online Appendix A.2.

5.6.1. Variation of Spatial and Temporal Demand Intensity.

5.6.1.1. Parameter Variations. In addition to the standard demand pattern, we define four additional *demand patterns* which range from spatial and temporally homogeneous demand intensity to spatially and temporally heterogeneous demand intensity (the standard demand pattern), as illustrated in Figure 3 in Online Appendix A.2. In the most homogeneous demand pattern, there is no spatial and no temporal variation at all (bottom left in Figure 3 in Online Appendix A.2). In the most heterogeneous demand pattern, there is a high spatial and temporal variation, as observed in practice (top right in Figure 3 in Online Appendix A.2, *standard demand pattern*). Moreover, we also consider patterns with only spatial or temporal and intermediate variation.

5.6.1.2. Results. Regarding profit, there is a clear impact of spatial and temporal demand variation (Figure 13(b)). The superiority of C-ANT (and L-ANT) over the other benchmark pricing approaches is more pronounced the more spatial variation there is. When there is no spatial variation the difference between the anticipative and myopic approaches is considerably smaller. Thus, the spatial variation is the main driver of C-ANT's (and L-ANT's) advantage over C-MYOP (and L-MYOP): C-ANT (L-ANT) performs around 7 percentage points better than C-MYOP (L-MYOP) when there is only spatial variation or spatial and temporal variation but the approaches performs only around 2 percentage points better when there is only temporal variation. However, as the results for high spatiotemporal demand variation show, C-ANT (and L-ANT) leverages most on its anticipation when there is both high spatial *and* high temporal demand variation, as it is observed in practice. Overall, the dominance of C-ANT as discussed in Section 5.5 can be confirmed and C-ANT proves to be robust against changes in spatial and temporal demand variation. The main insight here is that more sophisticated pricing approaches are of particular value when there is more demand variation—especially spatial demand variation. For the analysis of rentals and prices we refer to Online Appendix A.8.1. The most interesting insight there is that C-ANT only uses the low price with high spatial variation. This shows that it indeed sacrifices revenue to nudge customers to drive to "better" areas.

5.6.2. Variation of Customer Preferences.

5.6.2.1. Parameter Variations. In this section, we again use the standard demand pattern. We define three *choice patterns* in which we alter the parameters $\beta^{distance}$ and β^{price} of the multinomial logit model which describes the customer choice behavior (see Section 5.4.2). As we cannot disclose the choice parameters estimated on Share Now data, we now use three new choice patterns (Table 2).

Table 2. Parameter Variations

Choice pattern	$\beta^{distance}$	β^{price}	$ASC^{NoChoice}$
Walking distance sensitive	−10	−7.5	−5
Price sensitive	−7.5	−10	−5
Walking distance and price sensitive	−10	−10	−5

Figure 14. (Color online) Profit Impr. Over BASE (SMALL, DSR = 2/3)

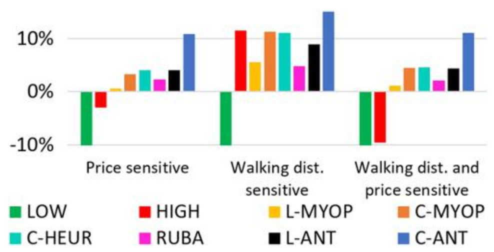

The first choice pattern (*walking distance sensitive*) is similar to the real values we estimated on Share Now data. Here, a walking distance change of 1 km has a higher impact on the customer's utility than a price change of 1 €/min. In the second choice pattern (*price sensitive*), the price is more important for the customer than the walking distance. In the last parameter variation, the customer is both *walking distance and price sensitive*. Please note that also customers always care about distance and price, for simplicity, we name the patterns according to the more pronounced sensitivity. For each choice pattern, we vary the DSR as in Section 5.5.

5.6.2.2. Results. Regarding profit, we consider Figure 14. C-ANT clearly outperforms all other pricing approaches across all choice patterns and all DSRs. Compared with C-MYOP, C-ANT yields a profit increase of up to 7.6 percentage points. However, there are substantial differences in the results between the three choice patterns. For example, with price sensitive customers, the improvements of all approaches over BASE are slightly lower (C-ANT's is the highest at 10.8%). With walking distance sensitive customers, improvements reach up to 16.2%. Regarding an analysis of rentals and prices we again refer to Online Appendix A.8.2.

In conclusion, we recommend C-ANT independent of customer preferences. It considerably improves profits and consistently provides the best result (significant at the 95% confidence level).

5.7. Variations of C-ANT

In the following, we briefly investigate straightforward variations of two aspects of C-ANT. First, in Subsection 5.7.1, we compare using the LP (11)–(15) from Section 4.1 and complete enumeration of all possible price combinations for the reachable vehicles to solve the pricing problem for each customer. Second, in Subsection 5.7.2, we look at the database used to approximate the vehicle values. As historical/simulated data depends on the pricing regiment active during that time, we consider iteratively updating the database based on new data, as already investigated in the TINY example (see Section 5.3). Throughout this section, as in the sections above, we use the standard demand pattern (SMALL, DSR = 2/3)

5.7.1. Comparison of Linear Program and Complete Enumeration.

5.7.1.1. Experiment. We define four cases in which we alter the number of possible prices. For each case, we calculate the optimal prices in C-ANT using the linear program (LP, Section 4.1) and complete enumeration (CE).

5.7.1.2. Results. Figure 13(c) shows computational times with 3 to 9 price points for up to 10 vehicles within the walking distance.

With three prices, there is no considerably difference in computational time between CE and the LP. For five or more prices, however, it is clearly visible that computational time with CE increases in the number of vehicles within reach and explodes above seven vehicles. For eight or more vehicles, CE's computational time clearly exceeds the LP's runtime, which remains consistently below 0.1 seconds. In conclusion, we recommend to use the LP to determine the optimal prices, independent of the number of possible prices $|\mathcal{M}|$.

5.7.2. Iterative Updating of Historical Data.

5.7.2.1. Experiment. In this analysis, we evaluate the impact of updating the historical data $\mathcal{K}^{idle}$ and $\mathcal{K}^{depart}$ (see Section 4.3.2) *iteratively* in batches based on new data. We denote these variants of C-ANT and L-ANT as C-ANT-IT and L-ANT-IT, respectively. This process is interesting for practical application, for example, a provider may update the data after several weeks of applying the most recent parameterization of C-ANT-IT (L-ANT-IT). More specifically, of the 5,000 simulation runs in this analysis, a batch consists of 1,000 runs each, such that the data are updated five times. For the first 1,000 runs, C-ANT-IT and L-ANT-IT use the same historical data as C-ANT and L-ANT (from 1,000 runs with BASE). For the second 1,000 runs, the *entire* batch of the historical data (1,000 runs BASE) is replaced by the data collected from these first 1,000 runs and so on.

5.7.2.2. Results. The results show that the iterative update based on new data improves the performance of C-ANT-IT and L-ANT-IT compared with C-ANT and L-ANT by 3.7 and 2.6 percentage points, respectively (see Figure 15(a)). With regard to the amount of rentals realized, there are no differences between these four approaches. However, with regard to pricing, Figure 15(b) shows that the updating leads to a higher proportion of lower prices being set. Overall, this analysis shows that there is additional potential for the

Figure 15. (Color online) Comparison of Iterative Updating Based on New Data and Static Database (SMALL, DSR = 2/3)

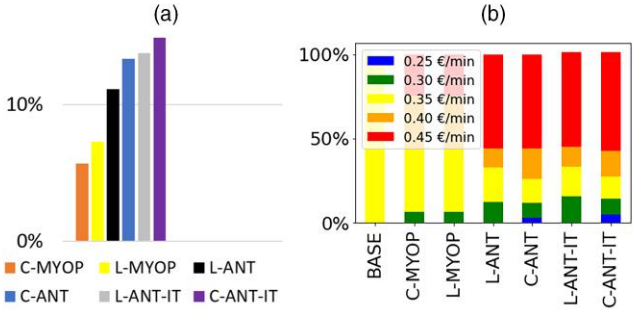

Notes. (a) Profit improvement over BASE. (b) Relative price frequency.

proposed solution approach when updating the historical data.

6. Case Study—Share Now in Vienna, Austria

In this section, we consider a real-world setting that reflects the origin-based dynamic pricing optimization of Share Now for a weekday in Vienna, Austria. This case study allows to conclude results and managerial insights from a real-world instance, as all parameters are based on real historical data which was collected over several months at Share Now. We introduce the scenario in Section 6.1 and discuss the results in Section 6.2.

6.1. Setting and Parameters

To respect the nondisclosure agreement, we do not share the exact origin and destination probability distributions $O(t)$ and $D(t, S_t)$, respectively. Instead, we present the course of the aggregate demand across the entire business area normalized to the maximum period demand (at base price) in Figure 16(b). Demand parameters are obtained from data that Share Now recorded during six months in 2018. We unconstrained the constrained demand, that is, the observed rentals, with the help of location- and time-specific app opening data which served as a proxy for the unconstrained demand. Such unconstraining is a standard issue in revenue management (see, e.g., Talluri and Van Ryzin 2004, Chapter 9.4).

The demand curve (Figure 16(b)) shows two peaks at the rush hour times, in the morning at 8:30 h and in the evening at 18:30 h, with the lowest level during the night at 3:00 h. This pattern is typical for weekdays in cities in central and northern Europe. The demand-supply-ratio is approximately DSR = 0.84, which is similar to the scenario with DSR = 2/3 above. As is in reality at the time of data recording, we use three price points. All other parameters (walking speed, stochastic rental times, etc.) are as in the computational studies (Section 5.4.1). Regarding the customer choice modeling, we applied the same multinomial logit model including the utility function as described in Section 5.4.2, which has been estimated on real-life data. Due to the very good performance of the C-ANT pricing approach in the sensitivity analysis, only this pricing approach and some benchmarks (BASE, LOW, HIGH, C-MYOP) are used for the case study.

6.2. Results

We first consider the profit of the different approaches (Figure 17(a)). Again, LOW leads to a reduction in profit compared with BASE. The approaches HIGH and C-MYOP deliver almost identical profits. As in the numerical study, C-ANT obtains the best result. Compared with C-MYOP (6.1% better than BASE), C-ANT's solution is more than 2 percentage points better in profit.

Overall, the rental curves (Figure 17(b)) follow the general course of the demand curve, with less pronounced peaks. During the night, the difference between demand and rentals is smaller than during the day. This can be explained by the higher availability of vehicles

Figure 16. (Color online) Share Now in Vienna, Austria

Notes. (a) Business area. (b) Normalized demand.

Figure 17. (Color online) Results for Case Study

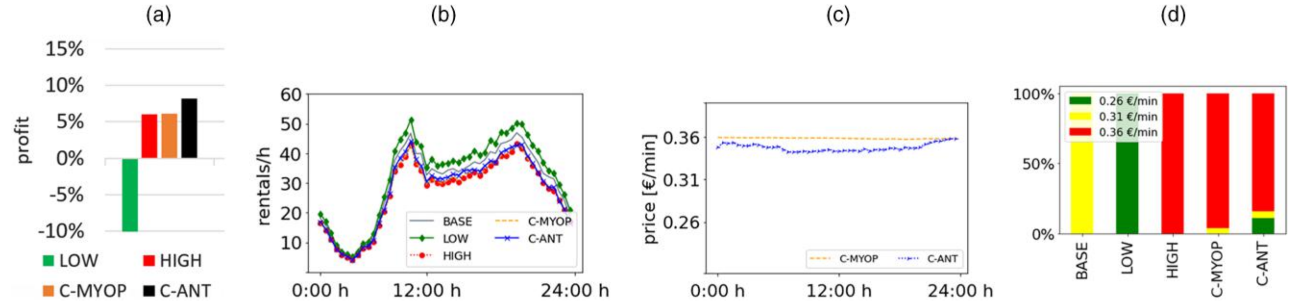

Notes. (a) Profit impr. over BASE. (b) Rentals over the course of the day. (c) Average prices over the course of the day. (d) Relative price frequency.

during the night, implying that potential customers almost always find an available vehicle. During the day, in particular during peak times, the probability that demand results in a rental is lower due to the relatively high number of vehicles in use. Regarding the pricing approaches, Figure 17(b) shows that LOW leads to the most rentals. Just below this is the curve of BASE.

The average price per period (Figure 17(c)) of C-MYOP is always above the average price of C-ANT and very close to the high price. In addition, differences in the average price over all periods (not depicted here) of the *reachable* vehicles and the *chosen* vehicles are also apparent. While the average price of C-MYOP is higher for the chosen vehicles than for the reachable vehicles, it is vice versa for C-ANT. This indicates that C-ANT sets low prices for reachable vehicles due to anticipation, so that the probability of choosing these vehicles increases. The rental curves (Figure 17(b)) of HIGH and C-MYOP are almost identical and the rental curve of C-ANT is above them. This can also be seen in the frequency of the prices (Figure 17(d)). Thus, most (96%) of C-MYOP's prices are the high price. Comparing C-ANT and C-MYOP, the frequency of the base price is larger (5% C-ANT, 4% C-MYOP). Furthermore, low prices are also more frequent (10% C-ANT, 0% C-MYOP). Therefore, the case study confirms that C-ANT is a viable pricing approach that can handle real-world problem instances.

7. Conclusion

In modern free-floating VSSs, providers have access to disaggregate real-time data regarding the locations of vehicles as well as of customers who open the mobile application to look for available vehicles. In this work, we demonstrate that this information can be leveraged in dynamic pricing to increase profitability. The *anticipative customer-centric* dynamic pricing approach in VSSs takes customers' location as well as their behavior regarding walking distances and prices explicitly into consideration in the online price optimization. Thus, vehicles can have different prices for customers who are requesting the price information at the same time but from different locations. Further, the specific pricing approach that we consider relies on *origin-based* minute prices. This origin-based feature is decisive for practice, because the information of a customer's intended destination is usually not available and its enquiry would contradict the spontaneous nature of free-floating VSSs. The third distinguishing feature of the developed pricing approach is that it is *anticipative*, that is, that future expected profits resulting from different spatial vehicle fleet distributions are taken into consideration.

We formally define the provider's online pricing problem as a Markov decision process and formulate the corresponding dynamic program by stating the corresponding Bellman equation. We show that in our approach, with regard to the action space of the pricing problem, only the vehicles within a customer's maximum walking distance have to be considered. Nevertheless, the dynamic program cannot be solved to optimality by classical backward induction due to the curse of dimensionality which, in our case, is (above all) caused by the state space containing the location of every vehicle in the business area.

To solve the online pricing problem, we develop a solution method based on approximate dynamic programming. We approximate state values representing expected future profits that occur *after* the current customer's decision, such that the current customer's choice behavior can still be considered explicitly with a disaggregated choice model in the optimization—in our case by a multinomial logit model. We take the assumption that state values are additive in the vehicle values which represent the profits that individual vehicles are expected to realize until the end of the considered time horizon. As a consequence of this assumption, vehicles which are not part of the current customer's consideration set can be neglected for the calculation of the state values, as they do not change their state for any possible choice and, thus, do not influence the online pricing optimization. To approximate the vehicle values, we propose a nonparametric value function approximation.

This type of approximation has two main benefits for implementation in practice. First, historical data can easily be used for the approximation and, second, approximate vehicle values can be precomputed such that the numerical operations of the online pricing problem can be reduced to a minimum.

In an extensive computational study with varying size of business area and fleet as well as varying demand patterns and overall demand levels, we demonstrate the advantages of our dynamic pricing approach compared with various benchmarks, including one from the literature and a myopic variant of customer-centric dynamic pricing. The new pricing approach outperforms all benchmarks significantly and considerably. It improves profits by up to 13.4% compared with the de facto standard in practice of constant uniform prices, as well as up to 7.6 percentage points compared with myopic dynamic pricing. From the latter, we conclude that the accurate approximation of our pricing approach is decisive for its performance. Compared with the benchmark from the literature, our approach obtains up to 8.4 percentage points more profit. Further, the numerical study demonstrates that the theoretical advantage of integrating the concept of customer-centricity in dynamic pricing compared with a location-based approach—in the case that state value approximations are exact—also applies when using the approximation that we propose. That is, considering situation-specific customer information like the position of the customer and distance to the vehicles within walking distance yields up to 2.3 percentage points more profit. The numerical results of a real-life case study based on Share Now data from Vienna also confirm the benefit of customer-centric and anticipative pricing and demonstrate the scalability of our approach. With regard to service level, we observe that anticipation in the pricing leads to an improvement, because there is a lower share of customers who do not have a reachable vehicle within the walking distance.

With a sensitivity analysis, we show that our results are robust regarding the decisive parameters of the customer choice behavior and we derive valuable managerial insights. We vary the influence of price and distance on the customers' utility of a vehicle and show that our pricing approach still always performs best in terms of profit. A detailed analysis indicates that this is because the new pricing approach leads to a higher variation of prices over different parts of the business area compared with a myopic pricing. The reason is the consideration of future vehicle locations and rentals. Thus, for example, our approach already raises prices in an area in the early morning if it anticipates a shortage of vehicles around noon. It would be very tedious to comprehensively mimic this anticipation with, for example, simple pricing rules. An analysis of spatial and temporal variations in demand shows that spatial variation, in contrast to temporal variation, has a stronger effect on the importance of anticipation. For a VSS provider this means that if there is no spatial demand variation, it is not necessary to anticipate the future in the pricing and rather straightforward approaches are sufficient—even a uniform pricing may be appropriate. If, however, there are already small spatial differences, it is worthwhile to anticipate the future. Another important insight for VSS providers is that our dynamic pricing approach manages to increase profits while maintaining the overall number of rentals that realize. This is important, since many service-related metrics that strive for customer satisfaction are related to a high number of rentals.

The comparison of the linear program formulation (based on Charnes-Cooper transformations) to solve the pricing problem with complete enumeration shows that the former is substantially more efficient for five or more prices. We further demonstrate that an iterative update of the historical data based on new data improves the performance of the developed solution approach.

To summarize, our new customer-centric, origin-based, and anticipative dynamic pricing approach for free-floating VSSs performs considerably well in comparison with existing approaches in terms of the relevant performance metrics. The nonparametric value function approximation solution method provides a scalable means to successfully account for the future evolution of the VSS based on current decisions, and allows to integrate disaggregated historical and real-time data which is readily available in practice for modern free-floating VSSs. Currently, to prepare implementation in practice, the approach is tested by our practice partner Share Now in an agent-based simulation (digital twin) which samples historical events on a disaggregated level, meaning with individual customer and vehicle events.

There are several reasonable paths for future research to extend this work. For example, in approximate dynamic programming, updating/stepsize rules often have an important impact on solution quality. Thus, there certainly is potential for additional improvements, given that there are several more sophisticated updating procedures than the applied batch update. Incorporating additional features such as idle-times in the vehicle value approximation may improve results and/or substitute vehicle values by another intuitive, often already available data source. Regarding the scope of the problem, a combined optimization of pricing and operator-based vehicle relocation seems natural.

Acknowledgments

The authors thank Share Now's Business Intelligence & Data Analytics department, in particular its head Clemens Kraus and Moritz Schattka, Team Lead Pricing, for their support and valuable input related to their pricing project. The authors also thank the three anonymous reviewers and the associate editor for their valuable feedback, which included, for example, the concept of the L-ANT benchmark.

References

Agatz N, Campbell AM, Fleischmann M, Van Nunen J, Savelsbergh M (2013) Revenue management opportunities for Internet retailers. *J. Revenue Pricing Management* 12(2):128–138.

Angelopoulos A, Gavalas D, Konstantopoulos C, Kypriadis D, Pantziou G (2016) Incentivization schemes for vehicle allocation in one-way vehicle sharing systems. *Proc. IEEE Internat. Smart Cities Conf.* (IEEE, Piscataway, NJ), 1–7.

Angelopoulos A, Gavalas D, Konstantopoulos C, Kypriadis D, Pantziou G (2018) Incentivized vehicle relocation in vehicle sharing systems. *Transportation Res. Part C Emerg. Technol.* 97:175–193.

Ataç S, Obrenović N, Bierlaire M (2021) Vehicle sharing systems: A review and a holistic management framework. *Eur. J. Transportation Logist.* 10:100033.

Banerjee S, Freund D, Lykouris T (2022) Pricing and optimization in shared vehicle systems: An approximation framework. *Oper. Res.* 70(3):1783–1805.

Barth M, Todd M, Xue L (2004) User-based vehicle relocation techniques for multiple-station shared-use vehicle systems. Working paper, UC Riverside, Riverside, CA.

Benjaafar S, Shen X (2023) Pricing in on-demand (and one-way) vehicle sharing networks. *Oper. Res.* 71(5):1596–1609.

Besbes O, Castro F, Lobel I (2021) Surge pricing and its spatial supply response. *Management Sci.* 67(3):1350–1367.

Bianchessi AG, Formentin S, Savaresi SM (2013) Active fleet balancing in vehicle sharing systems via feedback dynamic pricing. *Proc. IEEE Internat. Conf. Intelligence Transportation* (IEEE, Piscataway, NJ), 1619–1624.

Bierlaire MA (2020) A short introduction to PandasBiogeme. Technical Report, Transport and Mobility Laboratory, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland.

Brendel AB, Brauer B, Hildebrandt B (2016) Toward user-based relocation information systems in station-based one-way car sharing. *AMCIS Proc.* (Association for Information Systems, Atlanta), 1–10.

Chemla D, Meunier F, Pradeau T, Wolfler Calvo R, Yahiaoui H (2013) Self-service bike sharing systems: Simulation, repositioning, pricing. Working paper, CERMICS, École des Ponts ParisTech, Champs-sur-Marne, France.

Chen KD, Hausman WH (2000) Mathematical properties of the optimal product line selection problem using choice-based conjoint analysis. *Management Sci.* 46(2):327–332.

Chow Y, Yu JY (2015) Real-time bidding based vehicle sharing. *Internat. Foundation Autonomous Agents Multiagent Systems*, 1829–1830.

Chung H, Freund D, Shmoys DB (2018) Bike Angels: An analysis of Citi Bike's incentive program. *Proc. 1st ACM SIGCAS Conf. Comp. Sust. Soc.* (Association for Computing Machinery, New York), 1–9.

Clemente M, Fanti MP, Iacobellis G, Nolich M, Ukovich W (2017) A decision support system for user-based vehicle relocation in car sharing systems. *IEEE Trans. Syst. Man Cybern. Syst.* 48(8):1283–1296.

Côme E (2014) Model-based count series clustering for bike sharing systems usage mining: A case study with the Vélib' system of Paris. *ACM Transact. Intell. Syst. Technol.* Association for Computing Machinery, New York), 1–27.

Davis J, Gallego G, Topaloglu H (2013) Assortment planning under the multinomial logit model with totally unimodular constraint structures. Working paper, The University of Illinois at Urbana Champaign, Urbana and Champaign, IL.

DeMaio P (2009) Bike-sharing: History, impacts, models of provision, and future. *J. Public Transportation* 12(4):3.

Dötterl J, Bruns R, Dunkel J, Ossowski S (2017) Toward dynamic rebalancing of bike sharing systems: An event-driven agents approach. *EPIA Conf. Artificial Intell.* (Springer, Cham, Germany), 309–320.

Di Febbraro A, Sacco N, Saeednia M (2012) One-way carsharing. *Transportation Res. Rec.* 2319(1):113–120.

Di Febbraro A, Sacco N, Saeednia M (2019) One-way car-sharing profit maximization by means of user-based vehicle relocation. *IEEE Trans. Intell. Transportation Syst.* 20(2):628–641.

Ferrero F, Perboli G, Vesco A, Caiati V, Gobbato L (2015a) Car-sharing services – Part A taxonomy and annotated review. Working paper, Istituto Superiore Mario Boella, Turin, Italy.

Ferrero F, Perboli G, Vesco A, Musso S, Pacifici A (2015b) Car-sharing services – Part B business and service models. Working paper, Istituto Superiore Mario Boella, Turin, Italy.

Fishman E, Washington S, Haworth N (2013) Bike share: A synthesis of the literature. *Transportation Rev.* 33(2):148–165.

Fricker C, Gast N (2016) Incentives and redistribution in homogeneous bike-sharing systems with stations of finite capacity. *Eur. J. Transportation Logist.* 5(3):261–291.

Gallego G, Van Ryzin G (1997) A multiproduct dynamic pricing problem and its applications to network yield management. *Oper. Res.* 45(1):24–41.

Ghosh S, Varakantham P (2017) Incentivizing the use of bike trailers for dynamic repositioning in bike sharing systems. *Proc. Internat. Conf. Automated Planning Scheduling* 27(1):373–381.

Golalikhani M, Oliveira BB, Carravilla MA, Oliveira JF, Antunes AP (2021a) Carsharing: A review of academic literature and business practices toward an integrated decision-support framework. *Transportation Res. Part E Logist. Trans. Rev.* 138:23–45.

Golalikhani M, Oliveira BB, Carravilla MA, Oliveira JF, Pisinger D (2021b) Understanding carsharing: A review of managerial practices toward relevant research insights. *Res. Transportation Bus. Management* 41:100653.

Gönsch J, Klein R, Neugebauer M, Steinhardt C (2013) Dynamic pricing with strategic customers. *J. Bus. Econom.* 83(5):505–549.

Haider Z, Nikolaev A, Kang JE, Kwon C (2018) Inventory rebalancing through pricing in public bike sharing systems. *Eur. J. Oper. Res.* 270(1):103–117.

Herrmann S, Schulte F, Voß S (2014) Increasing acceptance of free-floating car sharing systems using smart relocation strategies: A survey based study of car2go Hamburg. *Internat. Conf. Computat. Logist.*, 151–162.

Illgen S, Höck M (2019) Literature review of the vehicle relocation problem in one-way car sharing networks. *Transportation Res. Part B: Methodol.* 120:193–204.

Jorge D, Correia G (2013) Carsharing systems demand estimation and defined operations: A literature review. *Eur. J. Transportation Infrastruct. Res.* 13(3):201–220.

Kamatani T, Nakata Y, Arai S (2019) Dynamic pricing method to maximize utilization of one-way car sharing service. IEEE *Proc. IEEE Internat. Conf. Agent* (IEEE, Piscataway, NJ), 65–68.

Kanoria Y, Qian P (2019) Blind dynamic resource allocation in closed networks via mirror backpressure. Working paper, Columbia Business School, NY.

Ke J, Yang H, Li X, Wang H, Ye J (2020) Pricing and equilibrium in on-demand ride-pooling markets. *Transportation Res. Part B: Methodol.* 139:411–431.

Laporte G, Meunier F, Wolfler Calvo R (2015) Shared mobility systems. *4OR* 13(4):341–360.

Laporte G, Meunier F, Wolfler Calvo R (2018) Shared mobility systems: An updated survey. *Ann. Oper. Res.* 271(1):105–126.

Mareček J, Shorten R, Yu JY (2016) Pricing vehicle sharing with proximity information. *Proc. 3rd MEC Internat. Conf. Big Data Smart City* (IEEE, Piscataway, NJ), 1–7.

Mourad A, Puchinger J, Chu C (2019) A survey of models and algorithms for optimizing shared mobility. *Transportation Res. Part B: Methodol.* 123:323–346.

Neijmeijer N, Schulte F, Tierney K, Polinder H, Negenborn RR (2020) Dynamic pricing for user-based rebalancing in free-floating vehicle sharing: A real-world case. *Proc. Internat. Conf. Comput. Logist.*, Lecture Notes in Computer Science (Springer International Publishing, Basel, Switzerland), 443–456.

Niels T, Bogenberger K (2017) Booking behavior of free-floating carsharing users. *Transportation Res. Rec.* 2650(1):123–132.

Pfrommer J, Warrington J, Schildbach G, Morari M (2014) Dynamic vehicle redistribution and online price incentives in shared mobility systems. *IEEE Trans. Intell. Transportation Syst.* 15(4):1567–1578.

Powell WB (2011) *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, 2nd ed. (John Wiley & Sons, Hoboken, NJ).

Qiu H, Li R, Zhao J (2018) Dynamic pricing in shared mobility on demand service. Working paper, Cornell University, Ithaca, NY.

Reiss S, Bogenberger K (2016) Optimal bike fleet management by smart relocation methods: Combining an operator-based with an user-based relocation strategy. *Proc. IEEE Intelligent Transportation Systems Conf.* (IEEE, Piscataway, NJ), 2613–2618.

Ren S, Luo F, Lin L, Hsu SC, LI XI (2019) A novel dynamic pricing scheme for a large-scale electric vehicle sharing network considering vehicle relocation and vehicle-grid-integration. *Internat. J. Production Econom.* 218:339–351.

Ricci M (2015) Bike sharing: A review of evidence on impacts and processes of implementation and operation. *Res. Transportation Bus. Management* 15:28–38.

Ruch C, Warrington J, Morari M (2014) Rule-based price control for bike sharing systems. *2014 Eur. Control Conf. (ECC)* (IEEE, Piscataway, NJ), 708–719.

Share Now (2021) Overview countries and cities. *Share Now*. Accessed February 25, 2021, https://www.share-now.com/de/en/country-list/.

Singla A, Santoni M, Bartók G, Mukerji P, Meenen M, Krause A (2015) Incentivizing users for balancing bike sharing systems. *Proc. Twenty-Ninth AAAI Conf. Artificial Intell. Pattern* (AAAI, Washington, DC), 723–729.

Soppert M, Steinhardt C, Müller C, Gönsch J (2022) Differentiated pricing of shared mobility systems considering network effects. *Transportation Sci.* 45(5):1279–1303.

Stancu-Minasian IM (1997) *Fractional Programming: Theory, Methods and Applications*, 1st ed. (Springer, Boston).

Statista (2022) Number of carsharing users in Germany by variant from 2014 to 2022. Accessed August 29, 2022, https://de.statista.com/statistik/daten/studie/202416/umfrage/entwicklung-der-carsharing-nutzer-in-deutschland/.

Strauss AK, Klein R, Steinhardt C (2018) A review of choice-based revenue management: Theory and methods. *Eur. J. Oper. Res.* 271(2):375–387.

Talluri KT, Van Ryzin GJ (2004) *The Theory and Practice of Revenue Management*, 1st ed. (Springer, Boston).

Train K (2009) *Discrete Choice Methods with Simulation*, 2nd ed. (Cambridge University Press, Cambridge, UK).

Wagner S, Willing C, Brandt T, Neumann D (2015) Data analytics for location-based services: Enabling user-based relocation of carsharing vehicles. *Proc. Internat. Conf. Inform. Systems*, vol. 3 (Association for Information Systems, Atlanta), 279–287.

Wang H, Yang H (2019) Ridesourcing systems: A framework and review. *Transp. Res. Part B: Methodol.* 129:122–155.

Wang L, Ma W (2019) Pricing approach to balance demands for one-way car-sharing systems. *Proc. IEEE Intelligent Transportation Systems Conf.* (IEEE, Piscataway, NJ), 1697–1702.

Waserhole A, Jost V (2012) Vehicle sharing system pricing regulation: A fluid approximation. Working paper, ENSTA ParisTech, Palaiseau, France.

Waserhole A, Jost V (2016) Pricing in vehicle sharing systems: Optimization in queuing networks with product forms. *Eur. J. Transportation Logist.* 5(3):293–320.

Zhang J, Meng M, David Z (2019) A dynamic pricing scheme with negative prices in dockless bike sharing systems. *Transportation Res. Part B: Methodol.* 127:201–224.