



Analyzing the impact of demand management in rural shared mobility-on-demand systems

Fabian Anzenhofer¹ · David Fleckenstein¹ · Robert Klein¹ ·
Claudius Steinhardt²

Received: 2 January 2024 / Accepted: 10 December 2024
© The Author(s) 2025

Abstract

In rural areas, shared mobility-on-demand services can improve the sustainability of public transport. However, bundling customer rides is challenging due to an unfavorable spatial and temporal demand distribution. As one potential solution, service providers could apply demand management. By controlling the availability of offered rides on an operational level, they could try to influence the resulting orders to allow more bundling. In practice, however, the introduction of demand management, which is a strategic decision, is often impeded by the inability of stakeholders to assess the exact impact on system performance in advance. In this paper, we tackle this issue by developing a methodology that serves as a basis for the strategic decision on how to implement operational demand management by realizing different types of demand control policies. More precisely, we propose a methodology that evaluates different policies by applying them to a model of the operational planning problem, which itself has not been considered in the existing literature. For this purpose, we first formulate the operational planning problem as a Markov decision process. Second, we apply practical solution algorithms representing different control policies on a model variant supporting the strategic decision. Finally, drawing on real-world data from FLEXIBUS, a rural provider in Germany, we conduct a computational study and present managerial insights into the impact of different control policies on the system performance in terms of profit, which the provider aims at maximizing, and other sustainability-oriented objectives of municipal contracting authorities.

Keywords Mobility-on-demand · Rural areas · Demand management · Availability control · Routing

1 Introduction

Sustainable mobility is a key societal goal, where there are still critical issues despite all efforts: severe environmental impacts of the transport sector (EEA 2023; EPA 2023), a lack of social equity in terms of access to mobility (Banister 2011, European Parliament 2021), and economic inefficiencies that lead to significant macroeconomic costs (Allcott 2013; Gössling et al. 2022). A significant cause of these issues is the high reliance on individual motorized transportation, accounting for about 80% of the modal split during the period from 2010 to 2020 (Eurostat 2022). This has pushed many countries to pursue a shift towards more sustainable public transportation options.

In rural areas, traditional scheduled public transport often enters a vicious circle of limited demand and supply (Bar-Yosef et al. 2013). This leads to inefficiencies in mobility provision, unattractive schedules, and, therefore, a low modal split owing to its non-competitiveness with motorized individual transport (Nobis and Kuhnimhof 2018). To overcome these issues, shared mobility-on-demand (SMOD) represents one of the most promising concepts (Alonso-González et al. 2018; Poltimäe et al. 2022; Sörensen et al. 2021). Also known as demand-responsive transportation (Schasché et al. 2022) or (shared) ride-hailing (Gilibert et al. 2020), SMOD refers to a flexible, demand-responsive passenger transportation system in which rides can be booked on request and shared by unrelated individuals through pooling.

In practice, SMOD systems have been successfully implemented in rural areas across various countries and by different providers, such as ioki (ioki 2024), Padam Mobility (Padam Mobility 2024), and Via (Via 2024). Several publications have also demonstrated the substantial sustainability benefits of SMOD in ecological (Coutinho et al. 2020; Prud'homme et al. 2011), social (Asatryan et al. 2023; Ma and Koutsopoulos 2022), and economical (Bischoff et al. 2017; Vazifeh et al. 2018) terms. There are also promising results regarding the comparison to scheduled public transport. E.g., Asatryan et al. (2023) compare scheduled buses with an SMOD system operating during late evening hours in Wuppertal (Germany), and find that the SMOD system improves the service quality.

However, trade-offs exist between improving traditional public transport and introducing SMOD (e.g., Viergutz and Schmidt 2019; Sieber et al. 2020). Overall, while cost-efficiency is highly dependent on the region, SMOD systems have a high potential in rural areas to improve service quality, accessibility, and environmental sustainability, especially if integrated with scheduled services (Mortazavi et al. 2024).

Still, many providers face operational challenges that lead to failure (Currie and Fournier 2020). This highlights the critical importance of operational planning. The general operational planning problem in SMOD systems has two main components: demand management, which refers to the operational decision on which rides to offer to a customer requesting service, and vehicle routing, which refers to the decision on how to fulfill the collected orders (Arian et al. 2022; Atasoy et al. 2015; Haferkamp and Ehmke 2022).

Existing literature concludes that an SMOD system's performance improves in different ways by using either more advanced demand management or more advanced vehicle routing (Haferkamp and Ehmke 2022). However, in rural areas where compatible requests are scarce, demand management seems to be more effective. It can potentially "generate" more compatible requests that can be successfully pooled. Despite this potential, rural providers in practice to date usually do not actively manage demand. Instead, rides are typically offered in a first-come-first-served manner. One reason for this is that the precise effects of implementing specific, more sophisticated approaches are hard to assess in advance.

In this paper, we tackle this issue by developing a methodology for supporting a rural SMOD provider's strategic decision on how to implement demand management. On the operational planning level, this requires selecting some type of demand control, which is realized by applying a control policy. To allow for a sound strategic decision, our methodology incorporates a precise model of the operational planning problem along with appropriate practical solution algorithms for possible control policies. We take into account the following unique characteristics, which have not yet been considered in the literature on demand management for SMOD systems:

- First, pricing is integrated and harmonized with scheduled public transport resulting in a static pricing scheme (Schasché et al. 2022). For such a pricing scheme, demand control is restricted to availability control which is a concept from the area of revenue management and which is based on the definition of (virtual) products (Klein et al. 2020; Strauss et al. 2018). In the context of rural SMOD, such products may correspond to fulfillment options, i.e., different pick-up or drop-off times, in response to a specific request. During the booking process, availability control then decides on which products to offer to each requesting customer.
- Second, due to the limited scheduled public transport alternatives in rural areas, both long-term planning reliability and short-term service availability are crucial features for rural SMOD systems to compete with motorized individual transport. To meet these requirements, any customer can place a ride request for a future service day (advance request), for a time later in the current service day (same-day request), and for the current point in time (ad-hoc request).
- Third, in rural areas, demand for SMOD services is often sparse and spread over a wide geographical area (Imhof and Blättler 2023; Wang et al. 2015), which makes efficient pooling of requests challenging.

Given these unique characteristics, i.e., the application of availability control to advance requests, same-day requests, and ad-hoc requests in a setting with dispersed demand, the operational planning problem for SMOD in rural areas is a novel dynamic and stochastic optimization problem which we refer to as the rural Shared Mobility-on-Demand Control Problem (r-SMCP).

Our work differs from existing literature on SMOD as we are the first to analyze the combination of availability control and advance requests in a rural context. In terms of the methodology, our approach is the only one guiding the strategical selection of availability control policies. The only other work taking a strategic view

on demand management is Haferkamp und Ehmke (2022), which introduces only a simple accept/reject policy and does not account for the three unique characteristics in rural areas. In addition, our methodology preserves the stochasticity of request arrivals, features solution algorithms that are readily applicable also at the operational planning level, and analyzes sustainability-oriented objectives. Apart from this, the only other works analyzing multi-option availability control in SMOD systems are by Sharif Azadeh et al. (2022) and Atasoy et al. (2015), both in an urban context. Finally, Arian et al. (2022) also consider a rural setting but apply dynamic pricing instead of availability control.

We consider availability control policies based on three characteristics: First, they can employ different mechanisms, i.e., rejections of a request (not offering a ride at all) or utilizing time shifts (offering alternative times to the originally desired time). Second, they can use two different criteria – feasibility or profitability – for decision-making. Third, they may differ in their use of information, resulting in myopic or anticipatory decision-making. Our methodology then incorporates a model variant supporting the final strategic decision, which we call semi-perfect information model. It serves as the basis for our computational analyses and carefully trades-off model accuracy and data availability. To maintain the focus on demand management and isolate its performance impact, we use a uniform approach for making vehicle routing decisions.

In summary, our work makes the following scientific contributions:

- We develop a methodology for analyzing the impact of different availability control policies on performance metrics reflecting the provider's and the municipal contracting authority's objectives. Transferred into practice, our methodology can be applied at the strategic planning level to evaluate in advance whether, and if so, which policy fits best for their specific system.
- As part of the methodology, we are the first to present a model and solution algorithms for the novel operational planning problem of rural SMOD providers (r-SMCP).
- We apply our methodology to one year of real-world data from our industry partner FLEXIBUS who have been operating an SMOD system since 2009, and therefore, belong to the most experienced providers in Germany. Therefore, this case study not only serves as a proof-of-concept for the methodology, but it also yields structural insights into the system performance in a typical, mature rural SMOD system.

The remainder of this work is structured as follows: In Sect. 2, we first review the literature and distinguish our work from the existing publications. In Sect. 3, we present the methodology, comprising models and solution concepts, for analyzing the impact of demand management in rural SMOD systems. The computational study using real-world data from FLEXIBUS, which serves as a proof-of-concept for the methodology and yields managerial insights, follows in Sect. 4. Section 5 summarizes the key managerial insights and includes a discussion of promising research opportunities.

2 Literature review

In this section, we delve into the existing literature and review publications that consider an SMOD system with similar basic characteristics. These basic characteristics include the following: First, customers place requests dynamically and must receive an immediate offer. Second, the provider cannot decide on the pick-up and drop-off stop of requests and has full control over the fleet. Third, the system allows ridepooling and is accessible for the general public. Fourth, the SMOD system is controlled independently without explicitly considering multi-modal interdependencies. Table 1 lists all currently existing publications that—to the best of our knowledge—meet these criteria. With the exception of Haferkamp and Ehmke (2022), all of them only consider the operational planning level. To show that we cannot directly draw on the models, solution algorithms, and computational results from these publications for evaluating operational demand control policies for the r-SMCP, we compare them to our work regarding three dimensions: the considered operational problem and instance structure, the solution concept, and the data used in the computational study.

Columns 2 to 6 compare problem and instance structure. Regarding the type of demand control (Column 2), we distinguish between feasibility-based control (FE), accept/reject decisions (AR), availability control (AV), or dynamic pricing (PR). Columns 3 to 5 indicate whether each of the three types of requests is considered. Column 6 indicates whether instances resembling the demand structure in rural areas are considered. To characterize the solution concept, Column 7 indicates whether the analysis is based on a Markov decision process (MDP), and Column 8 categorizes the solution concept as myopic (M) or anticipatory, more precisely, sampling-based (S) or learning-based (L). Finally, in Column 9, we distinguish between computational experiments based on an artificially generated data set (A), a data set sampled from real-world demand distributions (D), and a data set comprising original real-world requests (O). Based on Table 1, we discuss the assumptions and contributions of existing publications and delineate them from our work, and then summarize the resulting research gap. The discussion is loosely grouped along the problem and instance structure.

The only two works using availability control in SMOD systems are Sharif Azadeh et al. (2022) and Atasoy et al. (2015), albeit in an urban context. Consequently, the respective booking processes exclude advance requests, and the instance structure resembles an urban setting, which is more favorable for ridepooling. Another difference to our work is that both systems do not operate exclusively in a ridepooling mode, as customers are additionally offered a taxi-like service. In the case of Sharif Azadeh et al. (2022), there are further differences, since the authors integrate discrete pricing and generate fulfillment options by varying the length of the pick-up time window instead of varying the pick-up or drop-off time. There are also significant methodological differences to our work as both papers neither include an MDP formulation nor an anticipatory solution concept.

Arian et al. (2022) is the only existing publication considering demand control specifically suited for a rural SMOD system. However, their approach involves

Table 1 Literature overview

| Authors | Problem and instance structure | | | | | Solution concept | | Computational study |
|-----------------------------|--------------------------------|------------------|-------------------|-----------------|-----------------|------------------|------------------|---------------------|
| | Demand control | Advance requests | Same-day requests | Ad-hoc requests | Rural instances | MDP | Solution concept | |
| Araldo et al. (2019) | FE | X | X | ✓ | X | X | M | A |
| Arian et al. (2022) | PR | X | X | ✓ | ✓ | ✓ | L | A |
| Atasoy et al. (2015) | AV | X | ✓ | ✓ | X | X | M | A |
| Atanasio et al. (2004) | FE | (✓) | ✓ | X | X | X | M | D |
| Bischoff et al. (2017) | FE | X | X | ✓ | X | X | M | O |
| Elting and Ehmke (2021) | FE | ✓ | ✓ | ✓ | ✓ | X | M | A |
| Haferkamp and Ehmke (2020) | FE | X | X | ✓ | X | X | M | A |
| Haferkamp and Ehmke (2022) | AR | X | X | ✓ | X | ✓ | (S) | D |
| Heitmann et al. (2023) | AR | X | X | ✓ | X | ✓ | L | D |
| Horn (2002) | FE | X | ✓ | ✓ | X | X | M | O |
| Hosni et al. (2014) | AR | X | X | ✓ | X | X | M | A |
| Hungerländer et al. (2021) | FE | X | X | ✓ | ✓ | X | M | O |
| Jung et al. (2012) | FE | X | X | ✓ | X | X | M | D |
| Jung et al. (2016) | AR | X | X | ✓ | X | X | M | D |
| Loffi and Abdelghany (2022) | AR | X | X | ✓ | X | X | M | A |
| Lotze et al. (2023) | (FE) | X | ✓ | ✓ | ✓ | X | M | D |
| Lu et al. (2023) | FE | (✓) | X | ✓ | ✓ | X | M | O |
| Qiu et al. (2018) | PR | X | X | ✓ | X | ✓ | L | D |
| Sharif Azadeh et al. (2022) | AV/PR | X | ✓ | ✓ | X | X | M | O |
| Our work | AV | ✓ | ✓ | ✓ | ✓ | ✓ | S | O |

Abbreviations: Demand management: FE (feasibility control), AR (accept/reject), AV (availability control), PR (dynamic pricing); Solution concept: M (Myopic), S (Sampling-based), L (Learning-based); Data: A (Artificial), D (Real-world Distribution), O (Original requests)

dynamic pricing and only allows customers to place ad-hoc requests. This also results in a different definition of fulfillment options, since only one option per vehicle is generated based on the time it becomes available next. Similar to our methodology, they formulate an MDP and present an anticipatory solution algorithm. Qiu et al. (2018) investigate a comparable urban problem setting.

If requests can only be answered by offering a single fulfillment option for a static price, demand control is still possible by completely rejecting requests. The resulting control problem, involving only ad-hoc-requests and in an urban context, is investigated, e.g., by Haferkamp and Ehmke (2022). Their work is the closest to ours in terms of the methodology, as their goal is to analyze the performance impact of applying different demand control and vehicle routing policies from a strategic perspective. To this end, they also formulate the problem as an MDP and perform analyses based on a corresponding perfect information model, which provides results independent of the quality of the available data on customer choice behavior. In comparison, the semi-perfect information model we propose is more refined in that it preserves the stochasticity of request arrivals. Further, we use algorithms representing each control policy that can readily be applied to solve the actual fully stochastic operational planning problem.

There are several other publications considering accept/reject control in urban settings with ad-hoc requests. Among them, Heitmann et al. (2023) is the only one presenting an MDP formulation and an anticipatory policy. Hosni et al. (2014), Jung et al. (2016), and Lotfi and Abdelghany (2022) only apply myopic policies.

Finally, there are papers investigating purely feasibility control. The one most closely related to the paper at hand is Elting and Ehmke (2021), since their work is the first to investigate a problem with all three types of requests in a rural context. Hence, their analysis focuses on the performance of feasibility control depending on the share of advance requests. Other works investigating feasibility control include Hungerländer et al. (2021), Lotze et al. (2023), and Lu et al. (2023), each of which analyzes a data set with original requests from a rural SMOD provider. Araldo et al. (2019), Attanasio et al. (2004), Bischoff et al. (2017), Haferkamp and Ehmke (2020), Horn (2002), and Jung et al. (2012) also employ feasibility control but analyze urban settings.

In addition to the publications listed in Table 1, there is literature on (S)MOD services that are less closely related, which we briefly summarize for the sake of completeness:

- SMOD systems that allow providers to process and consolidate requests in batches (e.g., Alonso-Mora et al. 2017), which is impractical in rural areas due to the sparse demand.
- Ride-hailing systems that provide a taxi-like service and exclude ridepooling (e.g., Bertsimas et al. 2019).
- SMOD systems that control the assignment of pick-up and drop-off stops, which is also a form of demand control (e.g., Melis and Sörensen 2022).
- SMOD systems dedicated to a specific group of users (e.g., Schilde et al. 2011).
- Ex-post analyses based on the static demand control problem (e.g., Gaul et al. 2022).

- Evaluation of empirical data through descriptive analyses (e.g., Coutinho et al. 2020).
- Analyses from a system-oriented perspective using multi-agent simulation to explicitly model the interplay between different modes of transportation (e.g., Zwick et al. 2021).

In summary, the review of the existing literature reveals a significant research gap. Although there is some literature on availability control for SMOD systems (Atasoy et al. 2015; Sharif Azadeh et al. 2022) and on controlling advance requests with a pure feasibility control (Elting and Ehmke 2021), we are the first to analyze the combination of availability control and advance requests. In conclusion, the r-SMCP itself is a novel optimization problem that rural SMOD providers face on the operational level. Methodologically, our approach differs from the bulk of existing work, which all aim to solely develop specific solution algorithms for the operational planning level. Opposed to that, our work aims at comparing the impact of selecting different availability control policies from the strategic perspective similar to Haferkamp and Ehmke (2022).

To solve the semi-perfect information model of the r-SMCP, we transfer and adapt algorithms from literature on closely related attended home delivery problems (Campbell and Savelsbergh 2005, Yang et al. 2016, and Koch and Klein 2020) for two reasons: First, these algorithms are tailored to controlling advance requests, whereas there are no SMOD-specific solution algorithms for this purpose. Second, the existing specific anticipatory algorithms for similar SMOD control problems are all learning-based (see Table 1), which limits explainability and requires extensive training and tuning, unlike the sampling-based algorithms designed for attended home delivery problems.

3 Methodology to analyze the impact of demand management

In this section, we outline our methodology for conducting an impact analysis of demand management on the SMOD system performance at the strategic planning level. We first provide a brief overview in Sect. 3.1. Then, we elaborate on the two main components of our methodology, each of which we explain in a dedicated subsection: The problem formalization and modeling in Sect. 3.2 and the solution concept in Sect. 3.3.

3.1 Overview

At the strategic planning level, an SMOD provider may decide whether to apply demand management at all, and in case of static pricing, which availability control policy to select. This requires an explicit evaluation of the impact of applying different availability control policies at the operational planning level compared to the status quo (feasibility control).

Figure 1 provides an overview of the methodology. To support this strategic decision, the methodology evaluates different policies by applying a specific algorithm of it to a modelling variant of the operational control problem (r-SMCP), in our case, to the semi-perfect information model. Following the framework presented in Fleckenstein et al. (2023), the r-SMCP itself can be cast as a sequential decision problem for the provider, involving provider-side decisions (black boxes) and customer-side realizations of exogeneous information (gray boxes) at every decision epoch of the planning horizon: First, a request arrives. Second, the provider makes a offer decision. In the context of rural SMOD this means that the provider offers a restricted set of fulfillment options (availability control), more precisely, rides with different pick-up or drop-off times, in response to a specific request. The offer decision is determined by the availability control policy selected at the strategic planning level and an offer set results. Third, the customer's response to the offer set realizes in the order confirmation step, i.e., the customer either chooses one of the fulfillment options or abandons the booking process. Fourth, the provider makes a vehicle routing decision to dynamically plan the order fulfillment.

3.2 Modeling

In this section, we first develop the operational MDP formulation as a mathematical formalization of the r-SMCP in Sect. 3.2.1. Regarding this operational problem formulation, we then propose a modelling variant, the semi-perfect information model, which serves as a basis for the analyses of the impact of demand management at the strategic planning level.

3.2.1 Operational Markov decision process formulation

3.2.1.1 General notation and assumptions We formalize the r-SMCP as a Markov decision process (Puterman 2014), using a consider-then-choose discrete choice model (Aouad et al. 2021) to capture customer choice behavior. We follow this modeling approach because MDPs are suitable both as a concise mathematical problem definition as well as a formal basis for the solution concept we consider (Fleckenstein et al. 2023; Ulmer et al. 2020).

First, we introduce some general notation and assumptions:

- *Planning horizon:* Customers can place requests for a specific service horizon (operating day) over a multi-day booking horizon. We subdivide the booking horizon into a set of stages $T = \{1, \dots, t^s, \dots, T\}$ with $t \in T$ denoting each individual stage. t^s indicates the first stage within the corresponding service horizon, i.e., both horizons overlap.
- *Requests:* Customers can place requests for a ride between pairs pre-defined stops. All stops can be used as a pick-up or drop-off stops and are stored in the set \mathcal{H} . Formally, a request of type $c \in \mathcal{C}$ is characterized by the following attributes:

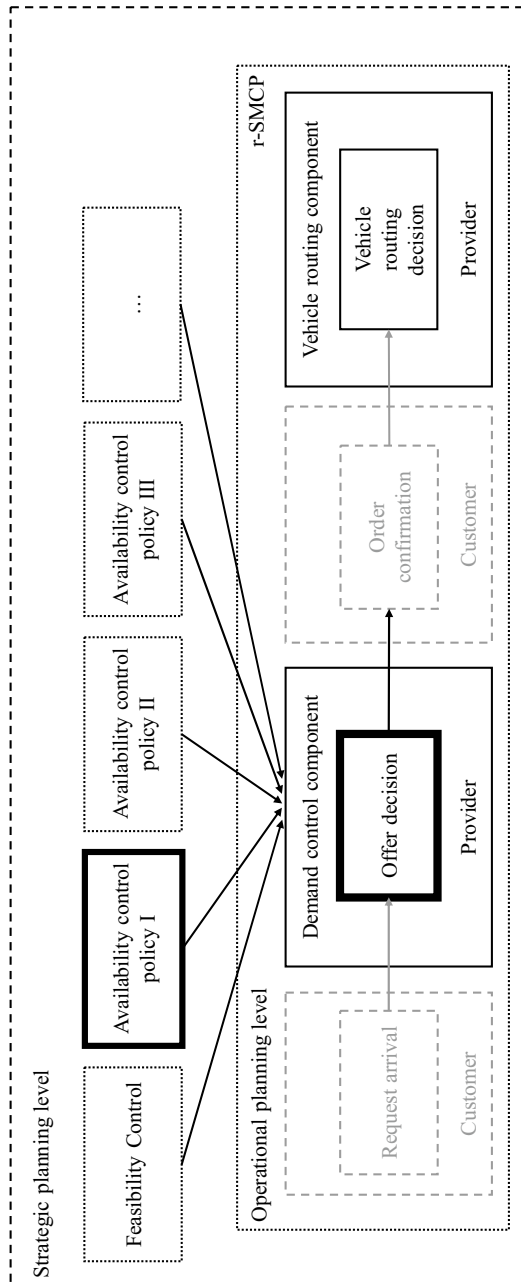


Fig. 1 An overview of the methodology

- Pick-up stop: $p_c \in \mathcal{H}$
 - Drop-off stop: $d_c \in \mathcal{H}$
 - Number of passengers: m_c
 - Desired time: $t_c \in \{t^s, \dots, T\}$
 - Desired time type: $f_c \in \{0, 1\}$ encoding whether t_c is a pick-up ($f_c = 0$) or a drop-off time ($f_c = 1$)
- Each request type c is associated with a fixed revenue r_c . Please note that a request can be placed for a single passenger ($m_c = 1$) or a group of multiple passengers ($m_c > 1$).
 - Individual requests $i \in \mathcal{I}$, with \mathcal{I} denoting the set of all individual requests, are defined by the underlying request type c_i . This implicitly defines the request's pick-up stop p_{c_i} , drop-off stop d_{c_i} , number of passengers m_{c_i} , desired time t_{c_i} , and type of the desired time f_{c_i} .
 - Finally, to model the case of no request arrival, we introduce a dummy request type $c = 0$.
- *Fulfillment options:* A fulfillment option $o \in \mathcal{O}_c$ represents a certain pick-up or drop-off time that the service provider offers in response to a request of type c with desired time t_c . The set \mathcal{O}_c includes all fulfillment options that can potentially be offered. When a customer places a request with a desired time t_c , the provider can respond in several ways:
 - Desired option: $o = t_c$ (the provider offers the exact desired pick-up or drop-off time)
 - Alternative option: $o \neq t_c$ (the provider offers a pick-up or drop-off time o deviating from the desired time)
 - No-purchase option: $o = 0$ (the provider allows the customer to abandon the booking process)

Thus, the set of potential fulfillment options \mathcal{O}_c can include the desired option and multiple alternative options but must include the no-purchase option.

Note that each fulfillment option o for a request of type c can be converted into a pair of time windows for pick-up and drop-off based on the direct ride time between pick-up and drop-off, the waiting time, and the maximum added ride time (see Appendix C or Jaw et al. (1986) for an in-depth explanation).

The provider then decides to present an offer set $g \subseteq \mathcal{O}_c$, which comprises a subset of the potential fulfillment options.

- *Order confirmation:* When faced with an offer set g , a customer with a request of type c chooses an option $o \in g$ based on probabilities $P_{c,o}(g)$ reflecting their time preferences. If an option $o \neq 0$ is chosen, the request $i \in \mathcal{I}$ becomes an order $j \in \mathcal{J}$ with $i = j$. Thus, the set of orders \mathcal{J} is a subset of the set of requests \mathcal{I}

($\mathcal{J} \subseteq \mathcal{I}$). In addition to the attributes of the corresponding request, an order j is further characterized by its associated fulfillment option $o_j \in \mathcal{O}_c$.

- **Order fulfillment:** To fulfill the orders, the provider deploys vehicles $v \in \mathcal{V}$ from a given fleet \mathcal{V} . Each vehicle has the following attributes:
 - Seat capacity: Q_v (maximum number of passengers per vehicle)
 - Start time: $t_v^b \in \{t^s, \dots, T\}$ (start of the service horizon)
 - End time: $t_v^e \in \{t^s, \dots, T\}$ (end of the service horizon)
 - Start and end location: $h = 0$ (the vehicle starts and ends its route at the depot)
- The planned route of each vehicle is encoded as the set $\theta_v = \left\{ (j_1, h_{j_1}, a_{j_1}^-, a_{j_1}^+), (j_2, h_{j_2}, a_{j_2}^-, a_{j_2}^+), \dots, (j_n, h_{j_n}, a_{j_n}^-, a_{j_n}^+) \right\}$, where $h_{j_n} \in \mathcal{H}$ is the n -th stop of the route. $j_n \in \mathcal{J}$ encodes the corresponding order the vehicle picks up or drops off. $a_{j_n}^- \in \{t^s, \dots, T\}$ and $a_{j_n}^+ \in \{t^s, \dots, T\}$ encode the vehicle's arrival time at and the departure time from the stop, respectively.

3.2.1.2 Markov Decision Process Now, with the general notation at hand and drawing on the modeling frameworks by Fleckenstein et al. (2023), Klein and Steinhardt (2023), and Ulmer et al. (2020), we formulate the MDP. Please note that the following explanations are complemented by two visual representations. Figure 2 depicts an intuitive visualization in the form of a decision tree. Figure 3 is more technical and provides a compact overview of the most important notation.

- **Decision epochs:** A decision epoch marks the beginning of each stage of the MDP, where the provider must make a decision. We adopt an incremental time-based definition (Puterman 2014). Each stage $t \in \mathcal{T} = \{1, \dots, t^s, \dots, T\}$ of the booking horizon is defined as a micro-period, with each period being equally and sufficiently short that the probability of more than one request arrival during the stage is negligible. Given this property, the arrival rate λ_c^t of the Poisson process underlying the request arrivals accurately approximates the probability of receiving exactly one request of type c in stage t .
- **States:** The post-decision state $s_t = (C_t, \phi_t)$ stores the information required for decision making at the subsequent decision epoch $t + 1$. The state definition of the r-SMCP comprises two elements:
 - Set of orders: C_t (storing all orders $j \in \mathcal{J}$ for which fulfillment has not yet been completed)
 - Current route plan: $\phi_t = \{\theta_{1,t}, \dots, \theta_{V,t}\}$ (where $\theta_{v,t}$ denotes the planned route of vehicle $v \in \mathcal{V}$)

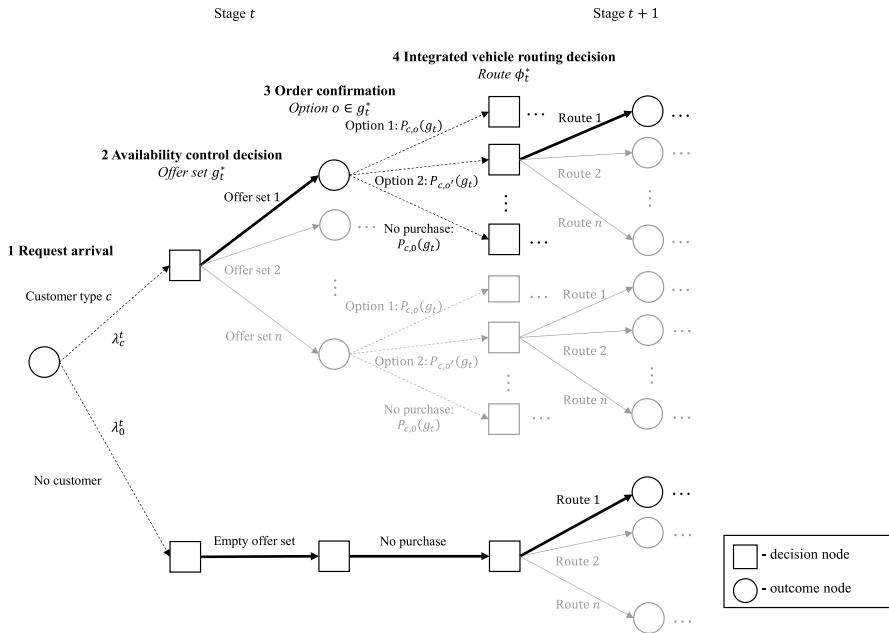


Fig. 2 Visualization of the Markov decision process with a decision tree

- Since in the r-SMCP, the provider makes decisions in response to a specific request arrival, defining s_t as a post-decision state simplifies the MDP formulation (Powell 2022). Note that decisions at epoch t are then made based on information stored in the preceding post-decision state s_{t-1} and the attributes of the request newly arrived in stage t (also see Fig. 3).
- **Actions:** An action is represented as $a_t = (g_t, (\phi_t(o))_{o \in g_t})$, and includes all operational decisions made at decision epoch t . In the r-SMCP, the provider applies demand management in form of an availability control, i.e., offers a limited set of fulfillment options, i.e., the pick-up or drop-off times. This *availability control decision* g_t is associated with integrated vehicle routing decisions $(\phi_t(o))_{o \in g_t}$ for any option $o \in g_t$ the customer could potentially choose when presented the offer set g_t . Note that both the control decision and the integrated vehicle routing decision are interdependent.

The availability control decision is encoded as an offer set $g_t \in \mathcal{G}(s_{t-1}, c) \subseteq 2^{\mathcal{O}_c} \setminus \emptyset$. The corresponding action space $\mathcal{G}(s_{t-1}, c)$ includes every feasible offer set, i.e., it is a subset of the power set of \mathcal{O}_c (excluding the empty set). An offer set is feasible if it only contains feasible fulfillment options.

A fulfillment option o is considered feasible if it satisfies the constraints of the integrated vehicle routing problem (see Appendix B for the corresponding model),

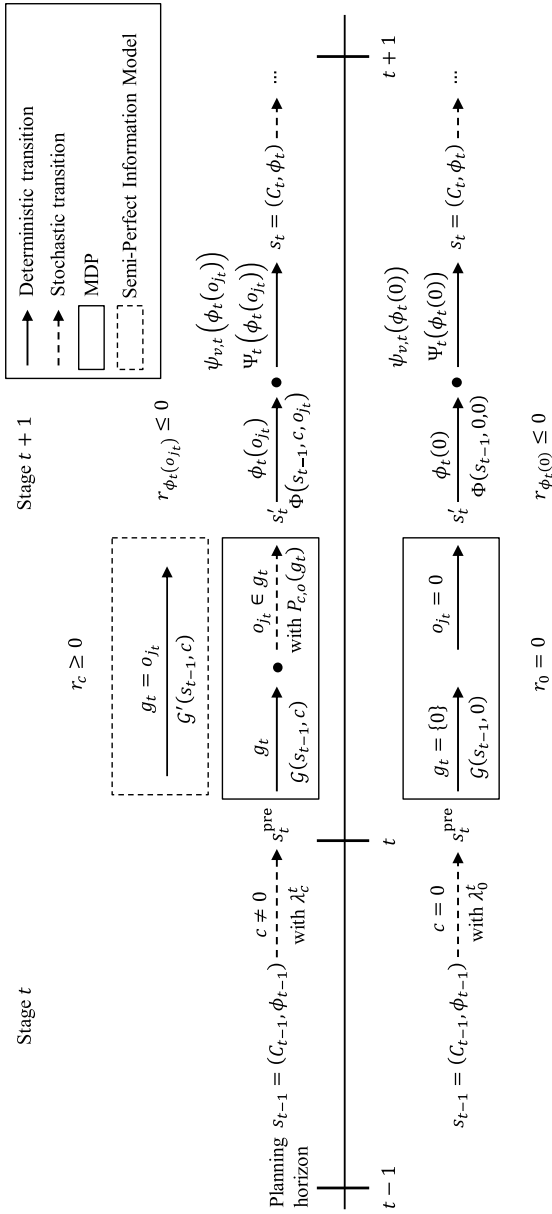


Fig. 3 Visualization of the Markov decision process and the semi-perfect information model

which is a standard dial-a-ride problem (DARP): More specifically, there must be at least one feasible route plan $\phi_t(o)$ that allows the provider to serve the following orders:

- All pending orders stored in C_{t-1} .
- The new potential order given option o is chosen.

We denote the set of feasible fulfillment options by $\mathcal{O}_c^f = \{o \in \mathcal{O}_c : \Phi(s_{t-1}, c, o) \neq \emptyset\}$. Thus, we can define the action space of the control decision more precisely as $\mathcal{G}(s_{t-1}, c) = 2^{\mathcal{O}_c^f} \setminus \emptyset$.

If there is no request, i.e., $c = 0$, the only feasible option is the no-purchase option, i.e., $\mathcal{G}(s_{t-1}, 0) = \{\{0\}\}$.

- For each feasible fulfillment option $o \in g_t$, the provider must make a tentative vehicle routing decision $\phi_t(o)$. The complete vehicle routing decision is encoded as a tuple of route plans $(\phi_t(o))_{o \in g_t}$. If the customer chooses option o , the corresponding route plan $\phi_t(o)$ is executed by the fleet until the subsequent decision epoch.
- The action space for the integrated vehicle routing decisions is $\prod_{o \in g_t} \Phi(s_{t-1}, c, o)$, which represents all combinations of feasible route plans.
- An important feature of this modeling approach is that vehicle routing decisions are predominantly tentative, meaning that large parts of the determined route plans $\phi_t(o)$ can still be adapted by future routing decisions. Only the parts of the route plans $\phi_t(o)$ that involve vehicle movements starting during stage $t + 1$ are definitive and must be executed.
- *Transitions:* Starting in a post-decision state s_{t-1} , a sequence of transitions caused by the provider's actions (availability control and routing decision) and customer-side stochasticity (request arrivals and customer choice behavior). The sequence leads to a successor state s_t . The transition process can be broken down into four steps: request arrival, availability control decision, order confirmation, and integrated vehicle routing decision (see Fig. 2):
 - *Step 1 (Request arrival):* First, the system transitions stochastically to the pre-decision state s_t^{pre} when a new request $i_t \in \mathcal{I}$ arrives (also see Powell 2022). If a request arrives ($c_{i_t} \neq 0$), Step 2 follows. If no request arrives ($c_{i_t} = 0$), the transition continues with Step 4.
 - *Step 2 (Availability control decision):* Once a request arrives, the provider determines an offer set g_t , which contains the options (pick-up times or drop-off times) the customer can choose from. This decision deterministically leads to the next step, where the customer makes a choice from the available options.
 - *Step 3 (Order confirmation):* Given the offer set g_t , the customer either confirms their order j_t by choosing an option $o_{j_t} \in g_t$, or they abandon the booking

process. The order confirmation follows customer-specific choice probabilities $P_{c_i,o}(g_i)$. If the order j_i is confirmed, it is added to the set of orders C_{t-1} .

- *Step 4 (Integrated vehicle routing decision)*: Finally, the process reaches the succeeding post-decision state s_t by a deterministic update of the route plan. The route plan $\phi_t(o_{j_i})$ determined as part of the routing decision replaces the route plan ϕ_{t-1} in the system state and is potentially partly executed. As already defined, $\phi_t(o_{j_i})$ is the route plan, pre-determined in the vehicle routing decision, specifically for the case that the customer chooses fulfillment option o_{j_i} . If in $\phi_t(o_{j_i})$ vehicle movements are planned to start until decision epoch $t + 1$, the respective stops $\psi_{v,t}(\phi_t(o_{j_i})) = \left\{ (j, h_j, a_j^-, a_j^+) \in \theta_{v,t} : \theta_{v,t} \in \phi_t(o_{j_i}), a_j^+ = t + 1 \right\}$ are removed from the individual routes $\theta_{v,t} \in \phi_t(o_{j_i})$ to reflect the planning being executed, and hence, becoming irreversible. If, by these vehicle movements, the fulfillment of some orders from C_{t-1} is completed, these orders are removed from C_{t-1} . The respective orders are determined according to

$$\Psi_t(\phi_t(o_{j_i})) = \left\{ j \in C_{t-1} : (j, h_j, a_j^-, a_j^+) \in \bigcup_{v \in \mathcal{V}} \psi_{v,t}(\phi_t(o_{j_i})), h_j = d_{c_j} \right\},$$
 i.e., based on the drop-offs ($h_j = d_{c_j}$) removed from the route plan.
- In summary, the transition from $s_{t-1} = (C_{t-1}, \phi_{t-1})$ to $s_t = (C_t, \phi_t)$ can be described as follows:

$$C_t = (C_{t-1} \cup \{j_i\}) \setminus \Psi_t(\phi_t(o_{j_i})) \quad (1)$$

$$\phi_t = \{ \theta_{v,t} \setminus \psi_{v,t}(\phi_t(o_{j_i})) : \theta_{v,t} \in \phi_{t-1}(o_{j_i}) \} \quad (2)$$

- *Rewards*: The provider collects two types of rewards:
 - Availability-control-related rewards ($r_c \geq 0$): These rewards are collected when a request of type c converts into an order. The reward corresponds to the fare paid by the customer, based on a static pricing scheme (see, e.g., Appendix J).
 - Vehicle-routing-related rewards ($r_{\phi_t(o)} \leq 0$): These are costs (negative rewards) incurred for the irreversible vehicle movements planned in $\phi_t(o)$. The routing costs are calculated based on the set of stops that are removed from the route plan $\psi_{v,t}(\phi_t(o_{j_i}))$ as follows:

$$r_{\phi_t(o)} = \begin{cases} - \sum_{v \in \mathcal{V}} \sum_{h: (j, h, a_j^-, a_j^+) \in \psi_{v,t}(\phi_t(o))} \rho_{h,h'}, & \text{if } \exists \psi_{v,t}(\phi_t(o)) \neq \emptyset \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

with h' denoting the successor stop of h in $\theta_{v,t} \in \phi_t(o)$ and $\rho_{h,h'}$ denoting the routing cost for traveling from stop h to stop h' .

- *Bellman equation:* The provider's objective can be represented using the Bellman equation (e.g., Powell 2019), which recursively defines the value $V_t(s_t)$, i.e., the expected future reward, for each state s_t and decision epoch $t \in \mathcal{T}$.

$$V_{t-1}(s_{t-1}) = \sum_{c \in \mathcal{C}} \lambda_c^t \max_{g_t \in \mathcal{G}(s_{t-1}, c)} \left(\sum_{o \in g_t} P_{c,o}(g_t) \left[r_c \cdot \mathbf{1}_{o \neq 0} + \max_{\phi_t(o) \in \Phi(s_{t-1}, c, o)} (r_{\phi_t(o)} + V_t(s_t | s_{t-1}, c, \phi_t(o))) \right] \right) + \lambda_0^t \max_{\phi_t(0) \in \Phi(s_{t-1}, 0, 0)} (r_{\phi_t(0)} + V_t(s_t | s_{t-1}, 0, \phi_t(0))) \quad (4)$$

with boundary condition $V_T(s_T) = 0$.

- The Bellman equation (4) consists of two summands, which can be explained as follows:

- *Request arrival:* The first summand models the case in which a request of type c arrives. The probability of such an event is λ_c^t . If a request of type c arrives as part of the transition from $t - 1$ to t , an integrated demand management and vehicle routing decision is necessary, which is reflected by the two nested maximum operators: First, the availability control decision is encoded by the outer maximum operator $\max_{g_t \in \mathcal{G}(s_{t-1}, c)} (\cdot)$. The provider selects an offer set g_t that maximizes their expected profit. Therefore, the provider must determine the sum of the positive reward (the fare paid) and the negative reward resulting from the vehicle routing decision for each option o , weighted by the probability $P_{c,o}(g_t)$. Second, the vehicle routing decision is encoded by the inner maximum operation $\max_{\phi_t(o) \in \Phi(s_{t-1}, c, o)} (\cdot)$. For each option $o \in g_t$, the provider must also decide on a routing plan $\phi_t(o)$, for which an evaluation of the vehicle-routing-related reward $r_{\phi_t(o)}$ and the value of the resulting post-decision state $V_t(s_t | s_{t-1}, c, \phi_t(o))$ is necessary.

- *No request arrival:* The second summand addresses the case in which no request arrives, indicated by $c = o = 0$. Here, the provider only makes a vehicle routing decision $\phi_t(0)$ by analogously solving the maximum operator

$$\max_{\phi_t(0) \in \Phi(s_{t-1}, 0, 0)} (\cdot).$$

- Drawing on the interim state $s'_t | s_{t-1}, c, o$ introduced by Fleckenstein et al. (2024), we can transform (4) such that the availability control subproblem is separated from the routing control subproblem:

$$V_{t-1}(s_{t-1}) = \sum_{c \in \mathcal{C}} \lambda_c^t \max_{g_t \in \mathcal{G}(s_{t-1}, c)} \left(\sum_{o \in g_t} P_{c,o}(g_t) [r_c \cdot \mathbf{1}_{o \neq 0} + V'_t(s'_t | s_{t-1}, c, o)] \right) + \lambda_0^t \cdot V'_t(s'_t | s_{t-1}, 0, 0) \quad (5)$$

$$\text{with } V'_t(s'_t | s_{t-1}, c, o) = \max_{\phi_t(o) \in \Phi(s_{t-1}, c, o)} (r_{\phi_t(o)} + V_t(s_t | s_{t-1}, c, \phi_t(o))).$$

- With another transformation, we can reformulate the availability control subproblem in (5) based on the opportunity cost

$\Delta V_t(s_{t-1}, c, o) = V'_t(s'_t|s_{t-1}, c, 0) - V'_t(s'_t|s_{t-1}, c, o)$ of converting a request of type c with option o into an order:

$$V_{t-1}(s_{t-1}) = \sum_{c \in \mathcal{C}} \lambda_c^t \max_{g_t \in \mathcal{G}(s_{t-1}, c)} \left(\sum_{o \in g_t} P_{c,o}(g_t) [r_c \cdot \mathbf{1}_{o \neq 0} - \Delta V_t(s_{t-1}, c, o)] \right) + V'_t(s'_t|s_{t-1}, 0, 0) \quad (6)$$

- Thereby, we exploit that $V'_t(s'_t|s_{t-1}, c, 0) = V'_t(s'_t|s_{t-1}, 0, 0)$ for each $c \in \mathcal{C}$. Formulation (6) is important, because it provides the theoretical foundation for decomposition-based solution concepts for integrated demand management and vehicle routing problems, which we also draw on in this work (see Sect. 3.3).

Stochastic modeling: The MDP formulation introduced above must be complemented by a suitable customer choice model, which defines the choice probabilities $P_{c,o}(g_t)$. These probabilities represent the exogenous information process (Powell 2022), determining how customers choose from a given offer set. While the MDP can be combined with any choice model, we apply a consider-then-choose model (Aouad et al. 2021). Generally, models of this class assume a two-step choice process:

- *Consideration set:* Customers use simple decision rules to filter out alternatives that they are not willing to choose at all. The remaining options form the customer's individual consideration set.
- *Ranking:* Second, customers rank the options in the consideration set according to their preferences and choose the highest-ranked option from the offer set. This can be the no-purchase option.

To allow for heterogeneity in the customer behavior, we define a set of customer segments \mathcal{L} . Each segment $l \in \mathcal{L}$ has its own consideration set structure and preference ranking. Many empirical studies have shown that the consider-then-choose paradigm and the heuristic construction of consideration sets are typical components of customers' multi-product decision-making (Hauser 2014).

In our model, the consideration set $\mathcal{S}_{l,c}$ for each customer segment $l \in \mathcal{L}$ and request types $c \in \mathcal{C}$ is determined by two quality cut-offs: Δ_l^+ for positive flexibility and $-\Delta_l^-$ for negative flexibility. These cutoffs represent the deviation from the customer's desired time t_c that they are willing to accept. Thus, the consideration set is defined as $\mathcal{S}_{l,c} = \{o \in \mathcal{O}_c : o - t_c \leq \Delta_l^+ \wedge o - t_c \geq -\Delta_l^-\}$.

Further, we assume a unique ranking function ζ for all segments, which ranks the fulfillment options in non-decreasing order based on their difference from the desired time t_c . In this regard, our model is similar to the lowest-open-fare model (e.g., Talluri and van Ryzin 2004), which is one of the standard models used in revenue management.

Given an offer set g , a request of type c and the customer's segment affiliation l , we obtain the option the customer will choose by:

$$o_{cgl} = \underset{o \in \mathcal{S}_{l,c} \cap g}{\operatorname{argmin}} \{ \zeta(o) \} \quad (7)$$

This means that the customer chooses the option that is closest to their desired time among the options they are willing to consider.

Finally, the choice probabilities $P_{c,o}(g)$ can be calculated as.

$$P_{c,o}(g) = \sum_{l \in \mathcal{L}} \gamma_l \cdot \mathbf{1}_{o=o_{cgl}} \quad (8)$$

with γ_l denoting the share of segment l in the customer population.

3.2.2 Semi-perfect information model

3.2.2.1 Motivation and outline While the operational MDP introduced in Sect. 3.2.1 accurately formalizes the r-SMCP, directly solving it to analyze the performance impact of demand management at the strategic planning level may not yield accurate results. This is because the quality of the results depends not only on the accuracy of the MDP formulation but also on how well the uncertain parameters of the MDP can be derived from historical real-world data to generate problem instances.

If these parameters are biased, the results will not properly reflect the real-world performance impact. In the case of the r-SMCP, there are two types of uncertain parameters:

- *Request arrivals*: It is uncertain what type of request c will arrive at each decision epoch $t \in \mathcal{T}$. This depends on the arrival rate λ_c^t .
- *Customer choice behavior*: It is also uncertain which fulfillment option o_{j_t} a customer chooses. This choice depends on the choice probabilities $P_{c,o}(g_t)$.

In practice, SMOD providers, such as our industry partner FLEXIBUS, can accurately track request arrivals because these are observable events. However, it is more challenging to capture customer choice behavior precisely as it involves complex, individual decision-making that is not easily observable.

To address this, we base our analyses on a semi-perfect information model, which is derived from the MDP formulation. It results from, on the one hand, preserving the stochasticity regarding request arrivals, but, on the other hand, deterministically modeling customer choice behavior. Hence, the solution algorithm is given perfect information about which fulfillment option a customer will choose from a certain offer set, but not about the requests that will arrive in the future. Thereby, the semi-perfect information model carefully trades off the accuracy of the model formulation against the accuracy of the parameter values obtainable from historical data. In the following, we explain in detail how this is achieved by the semi-perfect information model:

We assume that in SMOD systems, providers can track the arrival of a request at each decision epoch t , including the type of request c . This allows us to obtain the true realizations of demand reflecting the arrival rate λ_c^t from historical service days. Thus, no additional assumptions and modeling adjustments are required to model this source of uncertainty compared to the operational MPD formulation. To preserve the stochasticity of request, we simulate the request arrival process for each historical service day. That is, we generate a customer stream per day by using all original requests.

Unlike request arrivals, the true customer choice behavior cannot be directly observed. Historical data only reveals choices in response to the provider's historical demand control decisions. Approximating the true choice behavior by statistically estimating a choice model, which is the usual approach to modeling this uncertainty, would be particularly error-prone for the r-SMCP. Due to control decisions being made knowing the customer's desired time, providers usually try to offer options as close as possible to the desired time. Therefore, the historical data contains hardly any information on the true flexibility of customers, and there is a lack of exploration of the choice behavior.

To avoid having to rely on a potentially severely inaccurate customer choice model estimated on biased historical data, we instead consider customer choice deterministically. Since this assumption is strict, it is important to conduct sensitivity analyses to explore different customer choice behaviors in a systematic way (see Sect. 4.4).

3.2.2.2 Model formulation We now explain the resulting mathematical formulation of the semi-perfect information model and how it differs from the operational MDP formulation. The key difference is that while we retain the stochastic nature of request arrivals, we assume that the provider has perfect information about the customer's segment affiliation l_i for each request i_t for $t \in \mathcal{T}$. By assuming this, we can eliminate the need for estimating choice probabilities $P_{c,o}(g)$, which cannot be done reliably based on typically available data. Introducing perfect information on customer choice, the provider can deterministically steer the customer within their consideration set, which also changes the definition of actions and transitions. In Fig. 3, this corresponds to replacing the box with the solid frame by the box with the dashed frame.

- Formally, the following modifications occur compared to the operational MDP formulation:
- *Stochastic modeling*: In the semi-perfect information model, requests still arrive stochastically, following an arrival rate λ_c^t for each type of request c , just like in the operational MDP. However, customer choice behavior is now modeled deterministically. Since the segment affiliation l of a customer placing a request of type c is known, the provider can predict with certainty which option the customer will choose from the offer set g . Specifically, the choice probabilities are defined as:

$$P_{c,o}(g) = \begin{cases} 1, & \text{if } o = \operatorname{argmin}_{o' \in S_{l,c} \cap g} \{\zeta(o')\} \\ 0, & \text{otherwise} \end{cases} \quad (9)$$

In more detail, the segment affiliation yields the customer's consideration set $S_{l,c}$ which, according to the consider-then-choose paradigm, reveals which options $S_{l,c} \cap g$ from the offer set g the customer is generally willing to consider. Then, the known (uniform) ranking function ζ yields the most preferred among all considered options, which is the option the customer chooses with certainty. Therefore, the choice probabilities are effectively eliminated from the MDP. We illustrate the differences between the stochastic modeling component of the operational MDP and of the semi-perfect information model in Appendix D.

- **Actions:** The action space of the availability control subproblem can be reduced to $\mathcal{G}'(s_{t-1}, c) = (\mathcal{O}_c^f \cap S_{l,c}) \cup \{0\}$. Instead of determining an offer set that the customer chooses from, the provider deterministically assigns a feasible option from the consideration set. The assigned option becomes the confirmed order ($o_{j_t} = g_t \in (\mathcal{O}_c^f \cap S_{l,c})$), or the customer is rejected ($o_{j_t} = g_t = 0$). Thereby, the provider can fully exploit the flexibility provided by the customer.
- **Transitions:** Since the order directly results from the availability control decision ($o_{j_t} = g_t$), the originally stochastic transition from the pre-decision state s_t^{pre} to the interim state s_t' becomes deterministic.
- **Bellman equation:** Analogously to (6), the value function of the semi-perfect information model is then defined as:

$$V_{t-1}(s_{t-1}) = \sum_{c \in \mathcal{C}} \lambda_c^t \cdot \max_{g_t \in \mathcal{G}'(s_{t-1}, c)} (r_c \cdot \mathbf{1}_{g_t \neq 0} - \Delta V_t(s_{t-1}, c, g_t)) + V_t'(s_t' | s_{t-1}, 0, 0) \quad (10)$$

- Compared to the operational MDP (6), this formulation eliminates the need for choice probabilities $P_{c,o}(g_t)$ in the maximum operator $\max_{g_t \in \mathcal{G}'(s_{t-1}, c)} (\cdot)$. Since customer choice behavior is known with certainty, the reward for any availability control decision g_t consisting of the immediate reward r_c (if an order is confirmed) and the opportunity cost $\Delta V_t(s_{t-1}, c, g_t)$, becomes deterministic. Hence, determining the optimal control decision boils down to calculating the reward resulting from selling each of the feasible fulfillment options from the customer's consideration set, given by $\mathcal{G}'(s_{t-1}, c)$, and assigning the most profitable option as g_t . However, the stochasticity regarding request arrivals is preserved in the form of the arrival rate λ_c^t analogously to the value function of the operational model (6).

In summary, using the semi-perfect information model instead of the fully accurate operational MDP formulation, has two main advantages (Haferkamp and Ehmke 2022): First, our results represent an upper bound for the scenario of a certain average consideration set size, i.e., flexibility, in the customer population. Second, the

control policies' decision-making is only driven by the (accurately observable) customer stream and the general level of flexibility rather than a specific choice model. Thus, we obtain a clear picture of their respective control behavior and the performance impact, which is not distorted by the influence of a biased model of the customer choice behavior.

3.3 Solution concept

Solving the r-SMCP is equivalent to determining a policy, i.e., a function mapping each state to a decision, with a specific solution algorithm. To compute the optimal policy, it would be necessary to solve the Bellman Eq. (10), e.g., by backwards recursion. However, this is not possible for real-world instances since even the semi-perfect information model still exhibits two of the three curses of dimensionality (Powell 2019), namely regarding state and exogenous information. Thus, as part of our methodology, we define heuristic availability control policies for the r-SMCP (Sect. 3.3.1) and compare the performance of state-of-the-art solution algorithms that are representative of each policy (Sect. 3.3.2). This allows us to attribute performance differences to basic characteristics of availability control, i.e., to certain ways of decision-making, for the r-SMCP.

3.3.1 Availability control policies

Availability control policies for the r-SMCP can be systematically distinguished based on three key characteristics of availability control decision-making: First, a policy can utilize different mechanisms of availability control, namely rejections (not offering a ride at all) and time shifts (offering alternative times to the originally desired time). Second, availability control can be based on different criteria, either feasibility or profitability. Third, different types of information can be used for decision-making, either myopic information or anticipatory information. In the following, we provide a more detailed explanation of each characteristic in the context of the semi-perfect information model, which allows a deterministic assignment of fulfillment options by the provider. Since it is closely related to the operational MDP, the policies are readily transferable to policies for the operational MDP involving stochastic customer choice behavior.

- *Mechanisms:* For the r-SMCP, a policy can use rejections as a control mechanism, meaning that no fulfillment option is offered. The second mechanism are time shifts, i.e., controlling the offered times for request such that it differs from the desired time (e.g., by incremental steps). In the operational MDP, this can be done by only offering a selected subset of feasible fulfillment options. Note that both mechanisms can be applied separately or combined.
- *Criteria:* Both rejections and time shifts can be applied based on different criteria: feasibility and profitability. In the former case, the policy assigns an alternative fulfillment option or entirely rejects the request to avoid an infeasible order.

In the latter case, the policy assigns an alternative option although there are other feasible fulfillment options preferred by the customer or rejects the request if the order cannot be “made” profitable. Here, a request is considered profitable if it does not decrease the expected profit after fulfillment. Please note, that in general, the objective function does not necessarily have to be monetary. For the semi-perfect information model, we can unambiguously distinguish the four combinations of mechanisms and criteria: A feasibility rejection is applied if the policy cannot identify any feasible option for a request within the customer’s consideration set. Conversely, if the policy rejects a request despite having identified at least one such option, it applies a profitability rejection. To distinguish the two types of time shifts, we can use the closest feasible option, which is defined as the feasible option with the smallest deviation from the desired time. If the policy assigns an option with a deviation from the desired time equal to that of the closest feasible option, it applies a feasibility time shift. If the deviation is greater, this difference is a profitability time shift.

- *Information:* Among policies considering profitability, we can further differentiate between myopic policies and anticipatory policies. While the former only draw on information from the current state, the latter incorporate information about future demand to make more accurate profitability rejections and profitability time shifts. Feasibility-based decisions are myopic by design since the feasibility of any fulfillment option can be exactly verified based on the current state.

From (meaningful) combinations of these characteristics, we obtain a set of seven control policies, which we briefly introduce in the following:

- Feasibility control (FC): A feasibility control does not consider profitability, and, thus, only applies feasibility rejections and feasibility time shifts. Given at least one feasible option can be identified within the customer’s consideration set, it always assigns the closest feasible option.
- Myopic control (MC): A myopic control uses both types of rejections and time shifts and makes decisions according to myopic information. In addition to this general myopic control, we consider two special cases:
 - A non-selective myopic control (NS-MC), which does not apply profitability rejections.
 - A non-time-shifting myopic control (NT-MC), which does not apply profitability time shifts.
- Anticipatory control (AC): An anticipatory control also uses both types of rejections and time shifts, but its decision-making is additionally based on probabilistic information on future demand. Analogously to the MC, we consider two special cases:
 - A non-selective anticipatory control (NS-AC), which does not apply profitability rejections.
 - A non-time-shifting anticipatory control, which does not apply profitability time shifts (NT-AC).

Since we aim at analyzing the impact of demand management, the policies use a myopic approach for making vehicle routing decisions that does not involve waiting strategies or empty relocations (see also Sect. 3.3.2).

3.3.2 Solution algorithms

For each policy, we design one solution algorithm that is representative of it. We do not compare several different algorithms per policy, which would go beyond the scope of this work. The selected solution algorithms do not require extensive efforts for training and tuning such that they are easily adoptable in practice. Furthermore, they yield interpretable results regarding the policies' control behavior, which is particularly important for analyzing the performance regarding the objective of equal accessibility. Since we are the first to consider the r-SMCP, we transfer and adapt elements of existing algorithms for related control problems. We introduce the algorithms such that they are suitable for the semi-perfect information model which we use for our analysis on the strategic planning level. However, they can readily be adapted such that they can be applied to the operational MDP as we explain at the end of this section.

```

1    $C_0^{act} := \emptyset$ 
2    $\phi_0^{act}, \phi_0^{sam} := \{\theta_v = \{(0,0,t_v^b, t_v^b), (0,0,t_v^r, t_v^r)\}: v \in \mathcal{V}\}$ 
3    $C_0^{sam} := draw\_sample(day\_type, AR^{sam})$ 
4    $\phi_0^{sam} := parallel\_insertion(\phi_0^{sam}, C_0^{sam})$ 
5   forall  $t \in \mathcal{T}$  do
6        $\phi_t^{act}, C_t^{act} := execute\_route\_plan(\phi_{t-1}^{act}, C_{t-1}^{act}, \tau_{i_t})$ 
7        $\phi_t^{sam}, C_t^{sam} := synchronize\_route\_plans(\phi_t^{act}, C_{t-1}^{sam})$ 
8        $\mathcal{O}_{c_{i_t}}^f := \emptyset$ 
9       forall  $o \in \mathcal{O}_{c_{i_t}}$  do
10           $\phi_t^{act}(o), \mathcal{O}_{c_{i_t}}^f := feasibility\_check(\phi_t^{act}, \mathcal{O}_{c_{i_t}}^f, i_t, o)$ 
11          if  $o \in \mathcal{O}_{c_{i_t}}^f$  do
12              $\Delta \tilde{V}_t(s_{t-1}, c_{i_t}, o), \phi_t^{sam}(o) := opportunity\_cost\_estimation(\phi_t^{act}, \phi_t^{act}(o), \phi_t^{sam}, (c_{i_t}, \tau_{i_t}, o))$ 
13           $o_{j_t} := demand\_control(i_t, \{\Delta \tilde{V}_t(s_{t-1}, c_{i_t}, o): o \in \mathcal{O}_{c_{i_t}}^f\})$ 
14          if  $o_{j_t} \neq 0$  then
15              $C_t^{act} := C_t^{act} \cup \{j_t\}$ 
16              $\phi_t^{act} := routing\_control(\phi_t^{act}(o_{j_t}), C_t^{act})$ 
17              $\phi_t^{sam}, C_t^{sam} := update\_sampled\_route\_plan(\phi_t^{sam}(o_{j_t}), C_t^{act})$ 

```

Fig. 4 Basic solution algorithm

To characterize the different algorithms, we introduce the general solution concept and basic algorithmic structure. Thereby, we draw on classification and terminology presented in Fleckenstein et al. (2023). Overall, we adopt a decomposition-based approximation as the general solution concept, which is used in most existing publications on solving integrated demand management and vehicle routing problems and builds on formulation (10) of the r-SMCP. There are four subproblems resulting from this decomposition, which are tackled by different algorithmic components: feasibility check, opportunity cost estimation, availability control, and routing control. This is directly reflected in the basic structure of the solution algorithm depicted as a pseudocode in Fig. 4.

Statements with italic line numbers are only needed for the AC. In statements with an underlined line number, the variables ϕ_t^{sam} are only required for the AC. All other statements are common to all policies.

Before the start of the booking horizon, the actual route plan ϕ_t^{act} , which encodes the routing decisions, and the set C_j^{act} of all orders $j \in \mathcal{J}$ for which fulfillment has not yet been completed, are initialized as empty (lines 1 and 2). At each decision epoch, it is first computed which part of the route plan determined at the previous decision epoch has been executed, and C_j^{act} and ϕ_t^{act} are updated (line 6). Then, for each fulfillment option, the feasibility check is performed (line 10). If the result is positive, the opportunity cost estimate for the option is determined (line 12). Based on the results from lines 9–12, a control decision is made (line 13). If it results in a newly confirmed order, C_t^{act} is updated (line 15). In line 16, the routing control decision is made.

Feasibility check: By solving the feasibility check subproblem, the action space $\mathcal{G}'(s_{t-1}, c)$ of the r-SMCP's control subproblem is determined. Hence, the subproblem must be solved separately for each fulfillment option $o \in \mathcal{O}_c \cap \mathcal{S}_{l,c}$ that is part of the customer's consideration set. To ensure short computation times, we solve the feasibility check subproblem heuristically using a parallel insertion heuristic for the DARP (Jaw et al., 1986) and maintain the (tentative) route plan from the preceding decision epoch. If the potential order defined by i_t and o can be feasibly inserted, we add o to the set of feasible options $\mathcal{O}_{c_{i_t}}^f$. We integrate this approach, as given in Appendix E, in identical form into each of the seven policies.

Opportunity cost estimation: By solving this subproblem, we aim at determining an accurate approximation $\Delta \tilde{V}_t(s_{t-1}, c_{i_t}, o)$ of each potential order's opportunity cost. It measures the loss of expected future profit due the consumption of logistical resources associated with the additional order. Consequently, the classification of policies for the r-SMCP into FC, MC, and AC depends on the opportunity cost estimation approach.

- The feasibility control is characterized by generally setting $\Delta \tilde{V}_T(s_{t-1}, c, o) := 0$. Thereby, the opportunity cost estimation problem is effectively omitted.
- Myopic policies determine a myopic opportunity cost estimate, which is solely based on information stored in s_{t-1} . For the specific MC, and its two variants NT-MC and NS-MC, that we apply to the r-SMCP, we again draw on the parallel insertion heuristic and use the value of the cheapest insertion cost as a myopic

opportunity cost estimate, i.e., $\Delta \tilde{V}_t(s_{t-1}, c_{i_t}, o) := \cos t(\phi_t^{act}(o)) - \cos t(\phi_t^{act})$. Starting with Campbell and Savelsbergh (2006), this approach has been used in many works on integrated demand management and vehicle routing problems.

- Anticipatory policies additionally draw on probabilistic information on future demand to determine an anticipatory opportunity cost estimate. For the opportunity cost estimation in our specific AC, and its variants NT-AC and NS-AC, we apply a sampling-based look-ahead algorithm, which combines elements from the algorithms developed by Koch and Klein (2020), Köhler et al. (2024), and Yang et al. (2016) for an attended home delivery problem with similar structure. The basic idea is to derive the cost estimate from the cheapest insertion position of each potential order in a skeletal route plan, which we call the sampled route plan ϕ_t^{sam} . At the beginning of the booking horizon, the sampled route plan ϕ_0^{sam} is initialized only with a set of sampled orders C_0^{sam} (line 3 and 4). We draw C_0^{sam} directly from the historical data. Compared to methods that sample from individual distributions or joint distributions of request attributes, this sampling method performs superior since it extracts more accurate information about future demand from the historical data set (Köhler et al. 2024). Furthermore, we consider all historical requests and not only those that resulted in a confirmed order to avoid the sample being biased by the policy the provider used at the time the requests were observed. At each decision epoch, the algorithm first synchronizes the sampled route plan ϕ_t^{sam} with the actual route plan ϕ_t^{act} based on the routing control decisions that are made (line 7). Then, it determines the opportunity cost estimate $\Delta \tilde{V}_t(s_{t-1}, c_{i_t}, o)$ by searching the cheapest insertion position in the sampled route plan for each feasible fulfillment option. In case a new order is confirmed, the algorithm again updates the sampled route plan by selecting a sampled order to be replaced by the new order (line 17). A more detailed description of this algorithm can be found in Appendix H.

Availability control: Once the action space $\mathcal{G}'(s_{t-1}, c)$ and opportunity cost estimates $\Delta \tilde{V}_t(s_{t-1}, c_{i_t}, o)$ for all options $o \in \mathcal{O}_{c_{i_t}}^f$ are computed, a control decision $g_t = o_{j_t}$ must be determined by solving the maximum operator
$$o_{j_t} := \max_{o \in \mathcal{G}'(s_{t-1}, c_{i_t})} (r_{c_{i_t}} \cdot \mathbf{1}_{o \neq 0} - \Delta \tilde{V}_t(s_{t-1}, c_{i_t}, o)).$$
 For the semi-perfect information model, this can be done in linear time by complete enumeration. This approach is used for similar integrated demand management and vehicle routing problems in the literature, if the action space is sufficiently small (Avraham and Raviv 2021; Klein and Steinhardt 2023). The resulting solution algorithm for the control problem comprises the three steps profitability evaluation (line 1), time shift evaluation (line 2), and a tie breaker (line 3) (see Appendix F for a formal definition):

Regarding the first step, i.e., profitability evaluation, the revenue net of the option's estimated opportunity cost must be maximal as well as non-negative. In other words, the profitability evaluation ensures that only the most profitable option(s) is (are) selected. It is the only step that differs between the general

policies that are both selective and time shifting (FC, MC, and AC) and their non-selective and non-time-shifting special cases. In the case of the FC, both conditions are non-restrictive since all cost estimates equal zero. In the case of non-selective policies, the second condition is omitted, such that the evaluation returns the least unprofitable option if no profitable option exists. Conversely, for non-time-shifting policies, the first condition is omitted, such that only unprofitable options are filtered out. Regarding the second step, the subset of options causing the smallest time shift is generated from the result of the profitability evaluation. Then, either one option or two options with the minimal time shift in both directions remain. In the latter case, we assign the option with the earlier alternative time to break the tie.

Routing control: Mathematically, the routing control subproblem is defined by the maximum operator $\max_{\phi_t(o_{j_t}) \in \Phi(s_{t-1}, c, o_{j_t})} (r_{\phi_t(o_{j_t})} + V_t(s_t | s_{t-1}, c, \phi_t(o_{j_t})))$. We solve it as follows in all policies (also see Appendix G): At the last decision epoch before the start of the service horizon, the actual route plan ϕ_t^{act} is re-optimized from scratch by solving the static DARP with the parallel insertion heuristic. During the service horizon, we draw on the route plan $\phi_t^{act}(o_{j_t})$ resulting from the feasibility check for the assigned option o_{j_t} . In the case of a rejection, $o_{j_t} = 0$ and $\phi_t^{act} = \phi_t^{act}(0)$, i.e., the route plan is not changed. Combining feasibility check and routing control in this way is common in the literature on integrated demand management and vehicle routing problems (Fleckenstein et al. 2023).

In the operational MDP, the customer's choice behavior and thus the availability control decision on the offer set is stochastic. Hence, an appropriate choice model must be selected and estimated that yields the choice probabilities $P_{c,o}(g)$ of all options $o \in \mathcal{O}_c$ for any possible offer set $g \in 2^{\mathcal{O}_c} \setminus \emptyset$. The availability control decision on the offer set for an individual request represents an assortment optimization problem (Heger and Klein 2024).

4 Computational results

In this section, we evaluate the policies presented in Sect. 3.3 using a real-world data set, provided by our industry partner FLEXIBUS. From the different service areas FLEXIBUS operates in, we consider the most mature service area established in 2009, which consists of the small town Krumbach and the surrounding peripheral area.

The service area counts almost 1300 users who requested trips during the one-year observation period from February 2022 to February 2023. For the sake of comparability, we consider only working days with a service horizon from 5:00 a.m. to 9:00 p.m., excluding Fridays and holidays. This results in a data set of 200 service days, which all show statistically significant similarity.

In Sect. 4.1, we start with a brief descriptive analysis of this data set. Then, we introduce the experimental setup and the parameters of the base scenario in Sect. 4.2. After that, we discuss the results for the base scenario (Sect. 4.3) as well

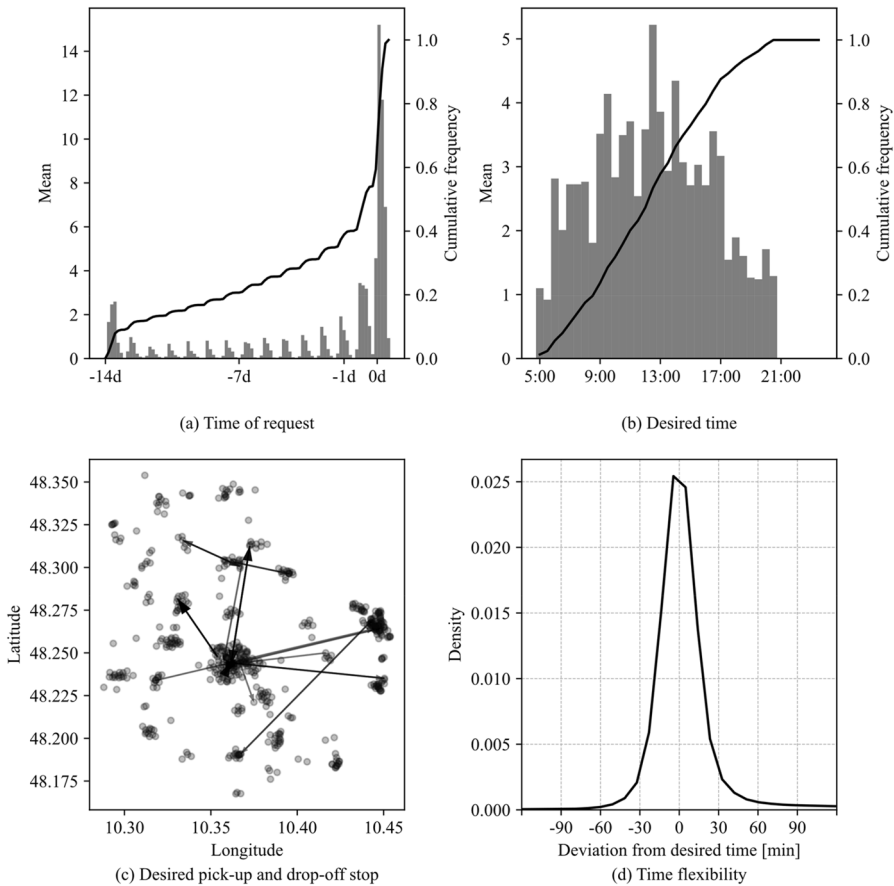


Fig. 5 Descriptive analysis of demand structure: **a** The horizontal axis shows the booking horizon, while the vertical axis displays the average number of requests per 4-h interval for a service day. **b** The horizontal axis represents the service horizon with specific desired times, while the vertical axis shows the average number of requests per 0.5-h interval. **c** The plot is a flow map, with arrows indicating the direction of OD pairs, and the arrowhead width representing the average frequency over the observation period. **d** The plot displays the empirical distribution (density) of observed average time flexibility

as the sensitivity analyses of demand-side flexibility (Sect. 4.4), level of profitability (Sect. 4.5), and supply-side capacity (Sect. 4.6). Finally, we demonstrate that our methodology can also be applied to support other strategic decisions (Sect. 4.7). We performed all computations on an Intel® Core® i7-6700 processor with 4 cores, 3.40 GHz, and 16 GB RAM. The algorithms were implemented in PYTHON (Version 3.9).

4.1 Descriptive analysis of demand structure

The descriptive analysis serves two purposes. First, we illustrate the key features of the r-SMCP, namely the relevance of advance requests, same-day requests, and

ad-hoc requests as well as the low, dispersed demand. Second, we analyze how much temporal flexibility customers have shown.

The booking curve depicted in Fig. 5a shows the relative proportions of advance requests (54.2%), same-day requests (22.8%), and ad-hoc requests (23%). Note that the booking curve does not show the advance booking time, but the actual time of request with data points grouped into 4-h bins. The fact that all three kinds of requests occur in relevant proportions highlights the importance of applying availability control to all of them, which is one of the distinguishing features of the r-SMCP. During an average booking horizon in the data set, 85.87 requests are received. Regarding the temporal perspective illustrated in Fig. 5b, the distribution of desired times is characterized by off-peak times at the beginning and the end of the service horizon (5 a.m. to

9 a.m. and 5 p.m. to 9 p.m.) and the relatively popular mid-day peak (9 a.m. to 5 p.m.). Looking at Fig. 5c, we observe a hub-and-spoke-type spatial distribution of demand with the town of Krumbach being both the most selected origin (54.9%) and the most selected destination (57.3%), and the peripheral area, with the exception of two larger villages, showing a very low density of demand. The arrows represent the direction of the OD-pairs and the width of the arrow heads correspond to the frequency, with which they are requested. We only show the most popular OD-pairs that are requested at least every two days. Figure 5d shows the empirical distribution (density) of the observed time flexibility measured by the difference of the desired time of the request and confirmed time of the order. In total, 79.20% of all orders show non-zero flexibility. The average request has a flexibility of 25.62 min. However, the data is likely heavily biased, because we can only observe the flexibility customers have shown based on the offer set they were presented. Since FLEXIBUS used a FC policy to control all booking processes in the data set, the observable flexibility should be viewed as a lower bound for customers' true flexibility, especially in case of advance requests that receive the best offers from an FC policy. That the observed lower flexibility bound already amounts to nearly half an hour, is a promising finding for the application of availability control.

4.2 Experimental setup

An instance of the r-SMCP is defined by two types of parameters: the request parameters and the scenario parameters. From each of the 200 historical service days contained in the dataset, we generate one customer stream by extracting the set of relevant request parameters for each request $i \in \mathcal{I}$. The scenario parameters describe the general setting of the SMOD system. Most scenario parameters result from the providers' strategic and tactical decision-making (system parameters), the remaining ones model the customer choice behavior (choice parameters). In the following, we describe the scenario parameters including their values in the base scenario, which is designed to resemble FLEXIBUS' real-world system as closely as possible.

The set of stops \mathcal{H} contains 563 stops across the service area, including the depot. Travel distance matrix and travel time matrix are calculated with Open Source Routing Machine (OSRM, n.d.) and include a constant service time of one minute. The fleet \mathcal{V} comprises a single vehicle deployed continuously during each service horizon. Since considering shift planning on a detailed, day-specific level would be out of scope for this study, we do not use the original shift plans from FLEXIBUS. We define fulfillment options analogously to FLEXIBUS and generate alternative fulfillment options with a step-size of $\epsilon = 10$ minutes starting from the desired time t_c . As an example, a request with a desired pick-up time of $t_c = 10 : 00$ could be offered an earlier pickup at $o = 09 : 50, 09 : 40, \dots$ or a later pick-up at $o = 10 : 10, 10 : 20, \dots$ as alternative options. To derive the time windows for pick-up and drop-off (Jaw et al., 1986), we use a uniform waiting time of $\omega = 10$ minutes and set the added ride time factor to $\mu = 0.5$. To determine the revenues r_c for all request types $c \in \mathcal{C}$, we use the original pricing scheme from FLEXIBUS' system. See Appendices I and J for a map of the service area, which highlights the different fare zones and the associated pricing scheme. Following this scheme, the revenue r_c depends on the number of fare zones, which a line connecting the stops p_c and d_c traverses, multiplied with the number of passengers m_c . The resulting revenues range from 2.4€ (one zone) to 9.9€ (eight zones) per passenger. To calculate the cost matrix $(\rho_{hh'})_{h,h' \in \mathcal{H}}$ from the travel distance matrix, we use a cost parameter of $0.3 \frac{\text{€}}{\text{km}}$, which is similar to the cost parameter FLEXIBUS assumes. For the choice parameters, we assume that customers belong to a single segment $l = 1$ with a consideration set of size $\Delta_1^+ = \Delta_1^- = 30$. Hence, all customers accept a maximum deviation of 30 minutes in both directions from their desired time, which is similar to the flexibility observable with descriptive analyses (Sect. 4.1) that can be viewed as a lower bound for the true flexibility.

We use a deterministic simulation framework to replay the original historical customer streams of the problem instances, which is based on the semi-perfect information model formulated. Each booking horizon begins at $t = 0$, which is 14 days prior to the service horizon. Each stage has the duration of one minute, such that $t^s = 20460$ and $T = 21420$.

We apply and compare the seven control policies introduced in Sect. 3.3: Feasibility control (FC), myopic control (MC), and anticipatory control (AC) as well as the non-selective variants NS-MC and NS-AC and the non-time-shifting variants NT-MC and NT-AC. All anticipatory policies require the selection of a sampling acceptance rate AR^{sam} , which we set to $AR^{sam} = 0.4$ based on preliminary tests. The other policies do not have any tunable hyperparameters.

We evaluate the performance of the availability control policies using the following additional metrics, each of which refers to one of the provider's or municipal contracting authorities' objectives:

- **Reliability:** The system should be reliable meaning that customers are shown a non-empty offer set as often as possible. To measure reliability, we consider the number of orders.
- **Environmental sustainability:** Having an SMOD system in place in a certain region should save emissions compared to not having the system. Hence, we

analyze the vehicle distance savings compared to motorized individual transport, i.e., booked passenger kilometers net of vehicle kilometers.

- **Service differentiation:** The SMOD service should be reasonably differentiated from other public transport modes regarding prices and the service characteristics. The authorities' aim behind this is to avoid undesirable cannibalization effects and create a level playing field. As a metric for this objective, we analyze the pooling rate, i.e., driven passenger kilometers divided by vehicle kilometers.
- **Equal accessibility:** Finally, equal accessibility is a prerequisite for that an SMOD system can provide mobility as a basic public service. It means that no discrimination should occur based on request characteristics. To evaluate the performance regarding equal accessibility, we analyze the policies' control behavior. More precisely, we consider the acceptance rate and the average time shift for different subsets of request types with certain characteristics (Sect. 4.3.2).

4.3 Base scenario

To analyze the impact of demand management in the base scenario, we apply the policies to 200 r-SMCP instances from the FLEXIBUS data set with the parameter setting of the base scenario introduced in Sect. 4.2. To account for the probabilistic nature of AC, NS-AC, and NT-AC, we calculate the (weighted) mean over 25 runs per r-SMCP instance for all metrics. In Sect. 4.3.1, we point out the general performance differences and discuss explanations for and implications from them. In Sect. 4.3.2, we deepen this analysis by investigating patterns in the control behavior of the seven policies with a focus on the objective of equal accessibility.

4.3.1 Overview

Table 2 summarizes each policy's performance regarding the objective metrics. To measure performance, we report the arithmetic mean (AM) and the coefficient of variation (CV). The average computation time per decision epoch lies between 0.003 s and 0.005 s for the non-anticipatory policies and only increases to around 0.007 s for the anticipatory policies. This indicates that our methodology can provide results

Table 2 Results overview of the base scenario

| Policy | Profit [€] | | Number of orders | | Distance savings [km] | | Pooling rate | |
|--------|--------------|-------------|------------------|-------------|-----------------------|-------------|--------------|-------------|
| | AM | CV | AM | CV | AM | CV | AM | CV |
| FC | 68.51 | 0.28 | 51.40 | 0.12 | −207.17 | 0.14 | 0.75 | 0.10 |
| MC | 93.82 | 0.25 | 44.12 | 0.21 | − 107.82 | 0.24 | 0.96 | 0.12 |
| AC | 98.42 | 0.21 | 49.84 | 0.14 | −124.27 | 0.13 | 0.93 | 0.08 |
| NS-MC | 80.20 | 0.26 | 52.72 | 0.13 | −180.15 | 0.15 | 0.83 | 0.10 |
| NS-AC | 80.74 | 0.23 | 53.05 | 0.11 | −182.14 | 0.12 | 0.82 | 0.08 |
| NT-MC | 86.85 | 0.26 | 43.34 | 0.21 | −124.54 | 0.24 | 0.86 | 0.13 |
| NT-AC | 83.95 | 0.24 | 45.40 | 0.17 | −137.05 | 0.15 | 0.81 | 0.09 |

even for considerably larger SMOD systems than that of FLEXIBUS in a reasonable time frame.

We observe a considerable profit gain of more than 35% due to availability control (MC and AC) compared to the FC. The revenue per order is comparable for all three policies (just above 4€ per order) and there are less orders for MC and AC. Hence, the profit gain can be attributed to a substantial cost reduction from 2.8€ per order to around 2.1€ per order, which is partially caused by a reduction of the orders' average OD-pair length from 5.4 km to 4.5 km. This finding shows that the main lever of improvement for availability control with uniform prices is increasing the routing efficiency and optimizing the length of orders' OD-pairs rather than exploiting the customers' willingness-to-pay to a larger extent or collecting more orders. The comparison to the non-selective and non-time-shifting policies shows that selectiveness contributes more to the performance gain.

The additional profit gain due to anticipation is only incremental (4.9%) but still statistically significant (p-value of Wilcoxon rank sum test: 0.01). The lower coefficient of variation indicates a more robust performance. Interestingly, there is no profit gain for NS-AC and NT-AC compared to their myopic counterparts. Hence, the benefit of anticipation only arises as a synergy benefit from combining both control mechanisms.

Although the revenues are sufficient to cover the variable routing cost for most types of requests, the potential of availability control with uniform prices is not large enough to achieve a positive operating result. For the base scenario, we observe a fleet productivity of around 6€ per shift hour for MC and AC, which is clearly not sufficient to cover the system's overhead cost, such as driver wages.

The number of orders, which is a measure of reliability from the customer perspective, is lower and less robust for MC (−14%) and AC (−3%) compared to the FC, notwithstanding the greater routing efficiency, which frees up shift capacity and would even allow more customers to be served. In fact, we observe a substantial reduction of fleet utilization from 79% (FC) to 54% (MC) and 62% (AC), which indicates that this capacity is not used. Hence, there is a subset of requests that are estimated to remain unprofitable, and are thus rejected, despite MC and AC being able to exploit their entire flexibility. Further evidence for this is provided by the results of the non-selective policies: Here, the freed-up capacity is used as the fleet utilization shows (78% for both policies), and both policies outperform the FC in terms of orders. In contrast to the profit gain, the increase in the number of orders between MC and AC is greater (13%), such that the AC gets close to the FC in terms of reliability. Although the vehicle is not fully utilized, FC, MC, and AC achieve rather low acceptance rates (FC: 59.9%, MC: 51.3%, AC: 58.0%), which can be improved by adding supply (Sect. 4.6).

The pooling rate indicates that MC and AC apply profitability time shifts extensively and thereby exploit the available demand-side flexibility to a larger extent compared to the FC. This additional flexibility is used to create consolidation opportunities, which results in substantially increasing pooling rates of MC and AC. Still, the pooling rates are relatively low, which underlines that rural areas are generally a challenging environment for SMOD services since demand is hard to consolidate.

When analyzing the distance savings, we find that the vehicle travel distance generally increases by several kilometres per order compared to the scenario in which all customers use their private cars to drive directly from their desired

origin to destination. Regarding this metric, availability control also leads to substantial performance improvements as it cuts the additional vehicle travel distance almost in half in absolute terms (MC: -48% , AC: -40%) as well as per order (MC: -39% , AC: -38%). As for the profit, this improvement is partially caused by accepting shorter OD-pairs on average.

4.3.2 Control behavior

In this section, we investigate the policies' control behavior toward different subsets of request types with certain characteristics. Thereby, we not only gain additional insights into the causes of the performance differences but can also assess the performance regarding the objective of equal accessibility. We first investigate the behavior of the policies in different phases of the booking horizon, i.e., depending on the time of request. Then, we conduct the same analysis depending on the desired time and the length of the requested OD-pair. To keep the plots clear, we only include FC, MC, and AC in the plots of this section, and refer the interested reader to Appendix K for plots including NS-MC, NS-AC, NT-MC, and NT-AC.

4.3.2.1 Time of request Since the total number of request arrivals varies over the 200 instances, we define the progress in the booking horizon based on the share of requests that has arrived already and group requests accordingly into 25 bins.

The MC consistently rejects more requests due to unprofitability since it only has information about consolidation opportunities with existing orders but not about future ones (Fig. 6c). By contrast, the AC correctly accepts additional advance requests that are unprofitable at their time of arrival but eventually become profitable when consolidated with future orders, which, in turn creates additional consolidation opportunities later. Forcing the myopic policy to accept any feasible order, as in the NS-MC, has a similar (but not equally beneficial) effect as anticipation, since consolidation opportunities with real orders start to arise at an earlier point in the booking process. The profitability rejection rates of the non-time-shifting policies are slightly higher, which indicates that for some of these additional requests, a profitability time shift is necessary to realize this consolidation.

Feasibility rejections show an inverse trend compared to profitability rejections (Fig. 6b). Because of the higher number of orders, the increase in feasibility rejections is greater for the AC, which to some extent thwarts the positive effect of more consolidation opportunities. The non-selective policies achieving lower feasibility rejection rates than the FC shows that profitability time shifts improve capacity utilization.

The resulting acceptance rates are generally decreasing almost monotonically until 80% of requests have arrived (Fig. 6a). The minimum corresponds to the maximum of feasibility rejections. In this phase, most customers request desired times within the mid-day demand peak, which also explains the subsequent small rise of acceptance rates, when desired times are again off-peak. By design, the FC starts with a 100% acceptance rate that drops over time with increasing slope. The same is true for NS-MC and NS-AC but on a higher level. In contrast, MC, AC, and their




Fig. 6 Use of rejections depending on the time of request: The horizontal axis plots the percentage of requests arrived. The vertical axis plots the rate of acceptances **a**, feasibility rejections, **b** and profitability rejections, **c** Each series corresponds to one of the policies FC, MC, and AC




Fig. 7 Use of time shifts depending on the time of request. The horizontal axis plots the percentage of requests arrived. The vertical axis plots the average total time shift, **a** feasibility time shift, **b** and profitability time shift per order in minutes. **c** Each series corresponds to one of the policies FC, MC, and AC

non-time-shifting variants achieve much more balanced acceptance rates, and thus, improve the performance regarding equal accessibility. Because of the fewer profitability rejections of advance requests, the AC initially achieves a much higher acceptance rate than the MC, while still maintaining a similar level for same-day requests and ad-hoc requests, which explains the gains in profit and confirmed orders.

Now, we investigate the use of time shifts. Until the arrival of one third of requests, we can observe clear differences in profitability time shifts between MC and AC (Fig. 7c). The MC hardly uses them initially, while the AC shows the maximal use since it anticipates later consolidation opportunities that can be realized with suitable time shifts. From then onward, MC and AC similarly show a decreasing use of profitability shifts since the available flexibility must increasingly be used for feasibility time shifts (Fig. 7b). NS-MC and NS-AC show a similar behavior but apply significantly more feasibility time shifts, probably to fulfill orders that can be consolidated poorly but that they are still forced to accept, leaving less flexibility for profitability time shifts.

The total time shift increases for all policies throughout the booking horizon (Fig. 7a). For the FC, these are all feasibility time shifts by design, which again indicates that it becomes increasingly difficult to find feasible options as more orders are confirmed. The non-time-shifting policies apply even less feasibility time shifts than the FC since they collect fewer orders, and more capacity is available consequently.

4.3.2.2 Desired time For this analysis, we group the requests into 1-h bins according to their desired time. We observe more profitability rejections by the MC for all types of desired times but the difference varies strongly (Fig. 8c). At the center of the mid-day demand peak (12 a.m. to 1 p.m.) as well as during the off-peak times in the early morning and late evening, the policies behave similarly. Around the boundary between peak times and off-peak times, the difference is much greater. This indicates that anticipation is especially beneficial when demand is moderate, and thus, some consolidation is possible but hard to identify. If demand is high and consolidation opportunities are easy to find, or if demand is low and consolidation is clearly almost impossible, the MC's inaccurate cost estimates do not lead to worse decisions compared to the AC.

As expected, feasibility rejection rates are roughly inversely proportional to the demand volume (Fig. 8b). At.

7 a.m. and 3 p.m., we observe local peaks, which are consistent with the peaks in the MC's profitability rejection rate. A possible explanation could be that around these times commuters request rides between the peripheral villages and the central town of the service region, which, if not consolidated well, consume a lot of logistical capacity and are unprofitable.

Overall, this results in three minima of the acceptance rate, which are more or less pronounced depending on the policy (Fig. 8a). The FC achieves the highest acceptance rate during off-peak times, due to MC and AC making maximal use of profitability rejections. During the mid-day demand peak, the AC shows the highest acceptance rate. Since acceptance rates range from 50 to 70% most of the time, the temporal discrimination in the off-peak periods is not particularly severe.




Fig. 8 Use of rejections depending on the desired time: The horizontal axis plots the desired time. The vertical axis plots the rate of acceptances, **a** feasibility rejections, **b** and profitability rejections. **c** Each series corresponds to one of the policies FC, MC, and AC




Fig. 9 Use of time shifts depending on the desired time: The horizontal axis plots the desired time. The vertical axis plots the average total time shift, **a** feasibility time shift, **b** and profitability time shift per order in minutes. **c** Each series corresponds to one of the policies FC, MC, and AC




Fig. 10 Use of rejections depending on the OD-pair length: The horizontal axis plots the direct distance between pick-up and drop-off in km. The vertical axis plots the rate of acceptances, **a** feasibility rejections, **b** and profitability rejections, **c** Each series corresponds to one of the policies FC, MC, and AC




Fig. 11 Use of time shifts depending on the OD-pair length. The horizontal axis plots the direct distance between pick-up and drop-off in km. The vertical axis plots the average total time shift, **a** feasibility time shift per order in minutes, **b** and profitability time shift per order in minutes. **c** Each series corresponds to one of the policies FC, MC, and AC

Regarding the use of time shifts, we do not find clear patterns aside from a slight increase in feasibility time shifts during the demand peak (Fig. 9). Thus, temporal discrimination by the use of time shifts is not an issue.

4.3.2.3 OD-pair length Now, we investigate the control behavior in spatial terms by considering requests with an OD-pair of similar length (1-km bins). Here, MC and AC show a sharp increase in the rate of profitability rejections for OD-pair distances between 8 and 11 km (Fig. 10c). Left and right of this interval, the rates are relatively stable. Rather than an inherent discriminatory behavior, this suggests an imbalance in the pricing scheme. Apparently, the revenue of many longer OD-pairs is not sufficient to make them equally profitable compared to shorter ones.

In terms of feasibility rejections, all three policies trend upward due to the higher logistical capacity consumption by the requests with longer OD-pairs (Fig. 10b). As a result, the FC shows a more balanced but still decreasing acceptance rate, while the acceptance rates of MC and AC are similarly unbalanced (Fig. 10a).

Considering the usage of time shifts, we do not find strong patterns, i.e., no systematic spatial discrimination (Fig. 11). Note that there is a limited number of data points, and consequently a high random variance, for requests with an OD-pair length greater than 15 km, which explains the outliers, in particular between 20 and 25 km.

4.4 Sensitivity analysis: demand-side flexibility

As discussed in Sect. 4.1, the FLEXIBUS data set only allows a limited descriptive analysis of the customer's willingness to accept time shifts, i.e., the demand-side flexibility. Hence, the true average consideration set size is uncertain and may be much greater than the empirically observed lower bound (around 30 minutes) we assume for the base scenario. Therefore, we analyze the impact of alternative consideration set sizes on the policies' performance in this section. For completeness, we start with a consideration set size of 0, i.e., entirely inflexible customers, and increase the consideration set size incrementally by 10 minutes until reaching a size of 180 minutes. The other scenario parameter values remain the same as in the base scenario.

Figure 12 shows the values of the objective metrics over the different scenarios. On the horizontal axis, we plot the consideration set size. Starting from the base scenario (size 30) and increasing flexibility, the profit of all policies improves at a diminishing rate (Fig. 12a). This improvement is quite considerable, especially in scenarios similar to the base scenario. E.g., until a consideration set size of 120, the AC achieves a 3% to 5% profit gain per 10 minutes of additional flexibility. For FC and MC, the slope is roughly equal, and thus, a constant profit gap slightly above 40% results. In contrast, the gap between MC and AC increases up to 19%. Hence, anticipation enables exploiting additional flexibility to a greater extent than already possible with feasibility control. Considering scenarios with very small consideration sets, we observe a sharp decrease in profit, especially for MC and AC. This finding provides further evidence for that time shifts, which are hardly possible in these scenarios, represent a more powerful control mechanism than rejections.

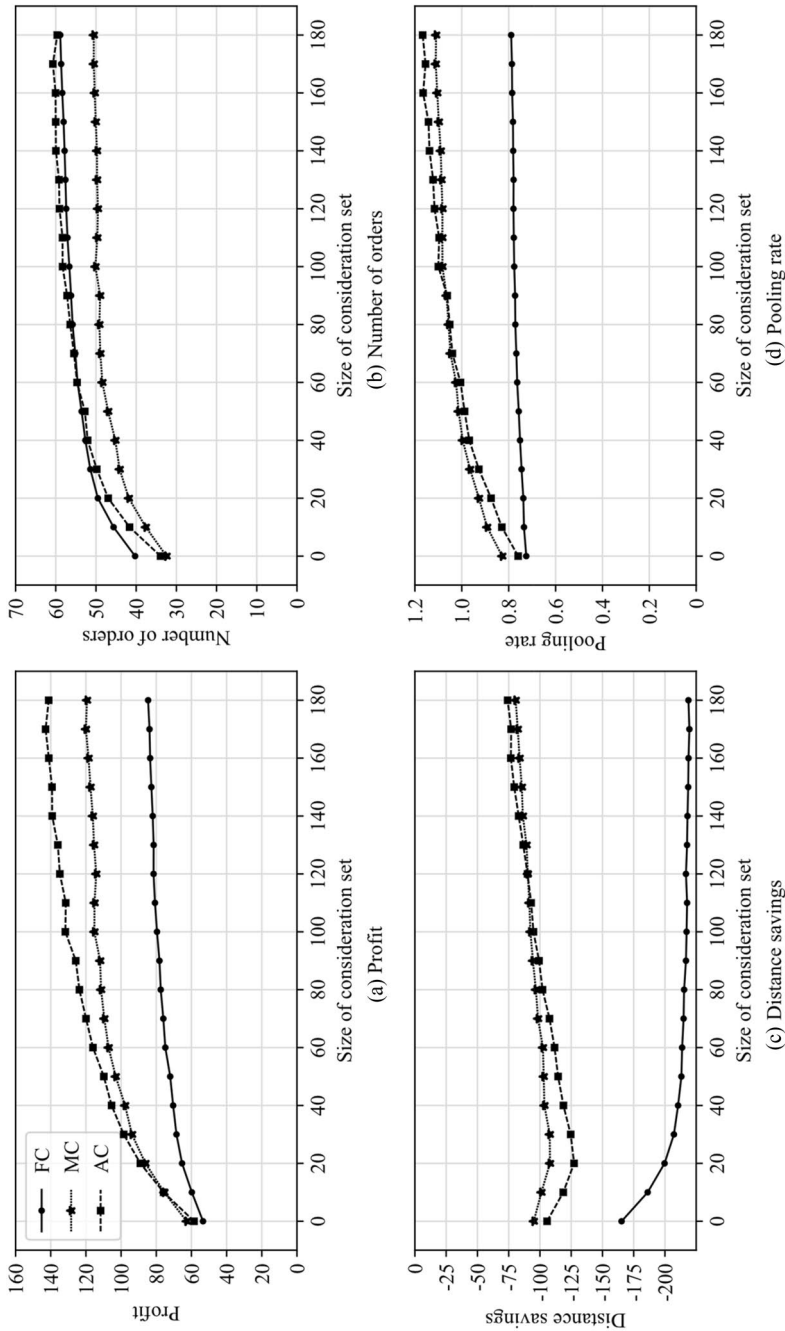


Fig. 12 Results of demand-side sensitivity analysis: The horizontal axis plots the size of the consideration set in terms of the maximum deviation from the desired time. The vertical axis plots the profit, **a** the number of orders, **b** the distance savings, **c** and the pooling rate. **d** Each series corresponds to one of the policies FC, MC, and AC

For the number of orders, we observe similar trends (Fig. 12b). The AC outperforms the FC from a size of 60 onward due to the effective use of profitability time shifts for making more requests profitable. The improvement drops to less than 2% per 10 minutes of additional flexibility for all policies at this point, such that only acceptance rates slightly above 70% are reached. This suggests that the vehicle supply increasingly becomes a constraining factor for the possible gains through exploiting the flexibility. Similarly, regarding the distance savings (Fig. 12c) and the pooling rate (Fig. 12d), the AC also becomes the best-performing policy at a certain point with FC and MC nearly stagnating.

4.5 Sensitivity analysis: level of profitability

When applying availability control with uniform fulfillment option prices, the provider needs to decide on a pricing scheme at the strategic planning level. The pricing scheme is yet another crucial input parameter for the r-SMCP because it determines the relation between the price level and variable fulfillment costs, i.e., the level of profitability. In this section, we investigate how the level of profitability impacts the system performance. To this end, we test a set of alternative pricing schemes that result from a change of the price for each number of zones by a certain percentage, i.e., a change of the general price level. The remaining scenario parameter values are identical to the base scenario.

Since we generate the instances with the original historical customer streams from the FLEXIBUS data set, we implicitly assume the demand to be completely price-inelastic. This represents a very strong assumption, which becomes less valid the more changes we make to the original FLEXIBUS pricing scheme used in the base scenario. Hence, it is not possible to draw meaningful insights from the policies' absolute performance expressed by the different metrics. Instead, we focus on the relative performance differences caused by the change in the level of profitability. When interpreting those differences, the assumption of inelastic demand is far less problematic since all policies have the same (deterministic) information about the customer choice behavior.

To analyze the impact of the level of profitability, we assume price reductions (lower level of profitability) and price increases (higher level of profitability) by up to 50% and generate scenarios in 5%-intervals. The results are plotted in Fig. 13. Comparing the policies' profit (Fig. 13a), we observe a declining gap between FC and MC/AC as the level of profitability increases. The underlying reason is that when requests become more profitable in general, the number of profitability rejections decreases, which more and more deprives MC and AC of one of their superior demand management mechanisms compared to the FC. If the level of profitability is very low, we observe that the FC even yields negative profits, which MC and AC can avoid by many profitability rejections.

Since the FC's control behavior is completely independent from pricing, its performance regarding the number of orders, the distance savings, and the pooling rate is constant over all scenarios. (Fig. 13b) clearly shows that MC and AC collect more orders when the system becomes more profitable overall since the number of profitability rejections decreases. Distance savings (Fig. 13c) and pooling rates (Fig. 13d)

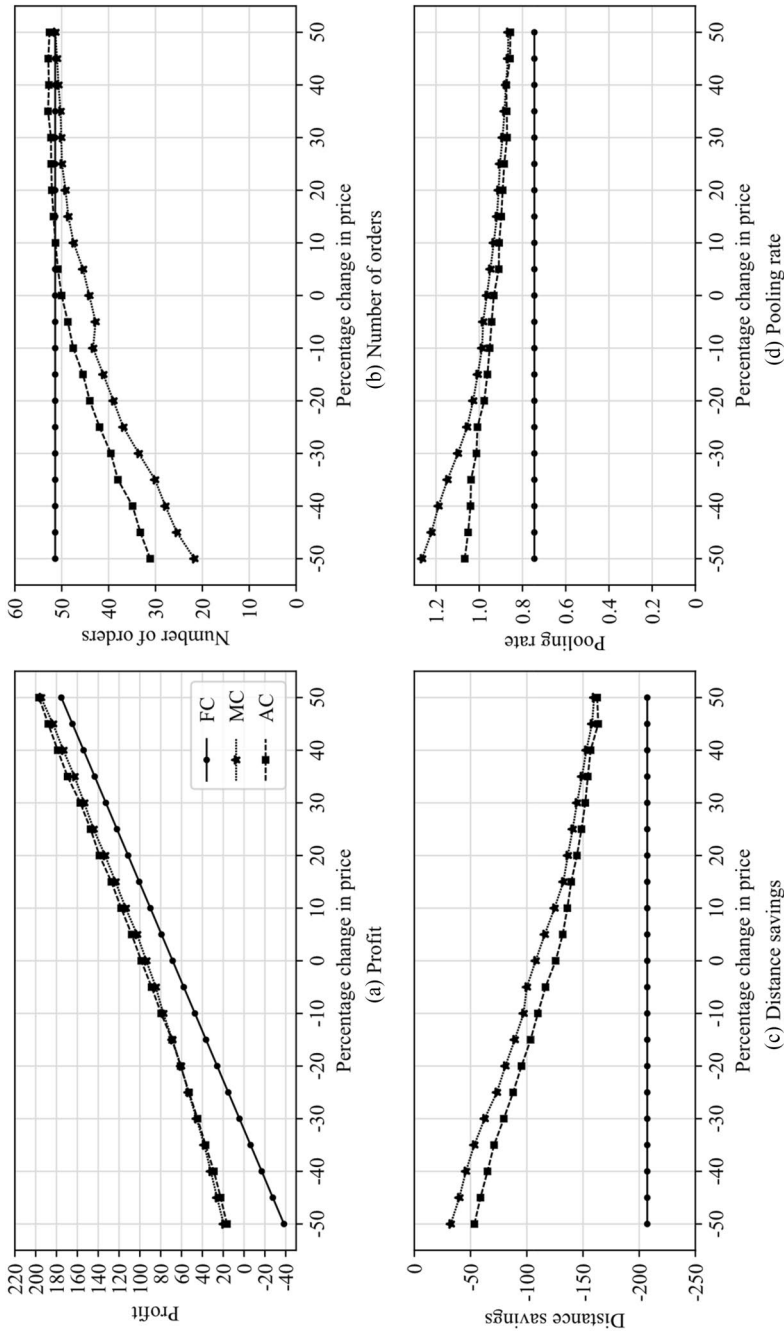


Fig. 13 Results of sensitivity analysis regarding the level of profitability: The horizontal axis plots the percentage change of the general price level. The vertical axis plots the profit, **a**, the number of orders, **b** the distance savings, **c** and the pooling rate. **d** Each series corresponds to one of the policies FC, MC, and AC

decrease with a higher level of profitability because, with higher prices, more and more orders for which the fulfillment is relatively inefficient become profitable.

4.6 Sensitivity analysis: supply-side capacity

The available vehicle fleet, which results from the provider's strategic and tactical planning, is a critical input parameter for the r-SMCP that constrains operational decisions. Given a certain demand, it determines the supply–demand ratio the SMOD system operates under in different phases of the service horizon. In practice, providers may target different supply–demand ratios. Hence, investigating how the policies perform under different supply–demand ratios is highly relevant.

In the base scenario, we use the minimum possible supply, i.e., a single vehicle over the entire service horizon (Index 0). We generate alternative scenarios by adding vehicles according to the following pattern: We start with an additional vehicle deployed for a two hour period in the center of the mid-day demand peak. We then increase its length successively until the second vehicle is also deployed over the entire service day (Indices 1 to 8). Applying this pattern once more yields scenarios for a fleet of three vehicles (Indices 9 to 16). Table 3 provides an overview of the scenarios. The other parameter values remain the same as in the base scenario.

Applying FC, MC, and AC to the scenarios with additional vehicle supply yields the results depicted in Fig. 14. The horizontal axis plots the indices of the scenarios from Table 3. For the total profit, we generally observe diminishing marginal gains for adding another vehicle (Fig. 14a). Deploying a second full-day vehicle, the FC's profit improves by 51%, MC and AC gain 33%. Adding a third full-day vehicle only leads to further improvements by 12% (FC), 5% (MC), and 6% (AC). The same accounts for prolonging the vehicles' time of deployment since, for the given scenario design pattern, the additional vehicle supply covers hours with less and less demand. Initially, the FC gains around 10% each time 2 vehicle hours are added, while MC and AC gain only around 6%. Thus, the FC's gap to MC (AC) decreases from 37% (42%) to 13% (19%), while the gap between MC and AC remains roughly constant. The performance of all policies appears to converge for high-supply scenarios. We observe similar results for the number of orders (Fig. 14b). Naturally, the acceptance rate of the FC converges to 100% and already reaches 95% for three full-day shifts. However, MC and AC only reach acceptance rates of 73% and 80%, respectively. This indicates that around 15% of all requests are inherently unprofitable given the flexibility and the revenue they provide in the base scenario parameter setting.

The results for the distance savings are inversely proportional to the number of orders, i.e., the distance savings per order are constant over all scenarios (Fig. 14c). Generally, we expect this metric to be influenced by two effects. First, the orders that

Table 3 Supply scenarios

| Index | 1/9 | 2/10 | 3/11 | 4/12 | 5/13 | 6/14 | 7/15 | 8/16 |
|------------------------|-------|-------|-------|-------|-------|-------|-------|-------|
| Start of added vehicle | 12:00 | 11:00 | 10:00 | 09:00 | 08:00 | 07:00 | 06:00 | 05:00 |
| End of added vehicle | 14:00 | 15:00 | 16:00 | 17:00 | 18:00 | 19:00 | 20:00 | 21:00 |




Fig. 14 Results of supply-side sensitivity analysis: The horizontal axis plots the index of the supply scenario as defined in Table 3. The vertical axis plots the profit, **a** the number of orders, **b** the distance savings, **c** and the pooling rate. **d** Each series corresponds to one of the policies FC, MC, and AC

can be collected additionally due to the growing supply should be increasingly less profitable to serve, which negatively impacts distance savings. Second, however, a growing number of orders entails more consolidation opportunities, which, together with the larger action space due to more vehicles, positively impacts distance savings. Based on our results, these effects seem to offset. Further evidence for this conclusion is provided by the pooling rate, which also remains roughly constant over all scenarios (Fig. 14d).

4.7 Further applications

Aside from supporting the strategic decision on how to implement operational demand management, which is the primary purpose of our methodology, it can also be applied to support other strategic decisions such as fleet sizing and the definition of pricing scheme or service areas. To demonstrate this, we exemplarily consider the fleet sizing decision, with which the provider mainly trades off the operating result and the reliability of the SMOD system. While Fig. 14 (Sect. 4.6) shows that more supply expectedly increases profit and number of orders, it also increases fixed cost (e.g., driver wages). Although we cannot directly measure the impact on the operating result due to a lack of data on fixed cost, we can derive insights from the profit per shift hour.

Figure 15 plots this metric against the number of orders, which measures the reliability. Each data point corresponds to the performance of one of the shift plans from Table 3 in combination with one of the demand control policies FC, MC, and AC. The different shades of grey indicate how much supply in terms of shift hours is available according to a certain shift plan. For all three policies, we observe that the profit per shift hour, and thus also the operating result, deteriorates with every additional




Fig. 15 Evaluation of potential shift plans: the horizontal axis plots the number of orders. The vertical axis plots the profit per shift hour. The marker shape indicates the policy of a data point, the shade of grey indicates the supply of shift hours

shift hour, even though the total profit increases (Fig. 14a). Hence, any increase of the fleet size can be seen as an investment into service quality. The specific numerical relation between the two metrics, i.e., the return of a certain investment, depends on the demand management policy. Thus, providers should optimize fleet sizing decisions such that they obtain a pareto-efficient shift plan reflecting their own weighting of economic efficiency and reliability. To this end, it is important to explicitly evaluate the potential shift plans regarding their operational consequences when the actual demand management policy is used. In this vein, evaluating all potential shift plans as depicted in Fig. 15, potentially taking into account other system parameters, can be viewed as a (brute-force) fleet sizing approach. Likewise, the provider can evaluate, e.g., the performance of different pricing schemes or service areas sizes.

5 Conclusion

In the following, we summarize the key findings our work. Based on the results from our computational study, we provide recommendations for implementing demand control in rural SMOD systems (Sect. 5.1). Furthermore, we address the limitations of our approach when applied to real-world settings, focusing on data availability, infrastructure requirements, and assumptions about customer behavior (Sect. 5.2). Finally, we outline possible future research directions (Sect. 5.3).

5.1 Recommendations

In this paper, we propose a methodological approach to evaluate at the strategic planning level whether and how an SMOD provider should control the availability of rides when performing demand management. We introduce a semi-perfect information model and representative algorithms for different availability control policies. These policies differ in three characteristics: applied mechanisms (rejections and/or time shifts), criteria (feasibility or profitability), and utilized information (myopic or anticipatory). The following recommendations summarize key insights from our computational study, focusing on the positive impact of demand management across multiple dimensions: economic efficiency (profit), reliability (number of orders), environmental sustainability (distance savings), service differentiation (pooling rate), and equal accessibility (control behavior):

- *Positive impact of availability control:* Overall, our experiments demonstrate that implementing availability control policies substantially improves the system performance across multiple objectives when compared to feasibility control. In the base scenario, myopic availability control already increases profit by about 37%. While there is a reduction in the number of orders (up to 14%), this is primarily due to the rejection of inherently unprofitable requests, leading to a far more efficient use of available resources. Distance savings show remarkable improvement, with availability control policies achieving around 48% greater savings compared

to feasibility control. Regarding service differentiation, the pooling rate increases by 28%. Lastly, equal accessibility also benefits, with more balanced acceptance rates for advance and same-day requests, although slight disparities for off-peak or long-distance requests may occur. Overall, the adoption of demand management through availability control policies offers considerable potential to improve the system performance across multiple objectives. Hence, we recommend its application in practice.

- *Advantages of anticipatory information:* While myopic availability control already yields substantial improvements, the use of anticipatory information further enhances system performance. In the base scenario, anticipatory control increases profit by an additional 5%. Moreover, it reduces the decline in the number of orders to -3% compared to myopic control. However, with anticipatory control, distance savings decline by -15% compared to myopic control, as it tends to generate more orders, leading to increased vehicle kilometers. The differences in pooling rate and accessibility are minor between myopic and anticipatory information utilized, with no significant patterns emerging. Thus, while the larger performance gain can be attributed to the applied criteria (profitability and not only feasibility), utilizing anticipatory information can bring additional benefits and is recommendable.
- *Synergy of mechanisms – rejections and time shifts:* Our results highlight that the combination of profitability rejections and time shifts unlocks the full potential of availability control policies, particularly when paired with anticipatory information. Profitability rejections ensure that unprofitable requests are excluded, while time shifts allow for more flexible and efficient order fulfillment. This synergy is especially pronounced in terms of profit, distance savings, and pooling rate. Therefore, we strongly recommend using both mechanisms in combination to fully exploit the potential of availability control policies.
- *Leveraging demand-side flexibility:* Our analysis shows that higher customer flexibility can improve system performance. Increasing the consideration set size from 30 to 120 min yields profit gains between 3% and 5% per 10 min of additional flexibility. However, after 120 min, gains begin to diminish as vehicle supply becomes a limiting factor. The number of orders, distance savings, and pooling rate also benefit from increased flexibility, with availability control outperforming feasibility control beyond a 60-min consideration set. Thus, providers should leverage as much of the customers' time flexibility as possible, e.g., by applying profitability time shifts.
- *Balancing supply and demand:* Adding more vehicles improves performance, but the marginal gains decrease as fleet size increases. E.g., adding a second full-day vehicle increases profit by 51% for feasibility control and 33 – 35% for anticipatory and myopic control. However, adding a third vehicle yields only 12% to 6% further improvement, and performance across all policies begins to converge as supply grows. The number of orders rises with vehicles, but distance savings and the pooling rate remain constant. These results show that providers should carefully balance vehicle supply with demand. Vehicle oversupply results in diminishing returns and does not improve key metrics like the distance savings.
- *Balancing profitability with environmental sustainability:* As prices increase, the profit gap between feasibility control and advanced control policies narrows,

with fewer requests being rejected as unprofitable. While this can boost profit given sufficiently high willingness-to-pay, it creates a trade-off with environmental sustainability. Higher profitability leads to more fulfilled requests, even if they result in inefficient vehicle usage, as routing costs become less critical. This can increase vehicle kilometers, reducing the system's environmental sustainability. Providers could adopt sustainability-oriented demand management approaches to ensure that profit maximization does not undermine environmental goals.

- *Computational times:* Importantly, our solution approaches are scalable for larger SMOD systems. The average computation time per decision epoch is 0.003–0.005 s for non-anticipatory policies and 0.007 s for anticipatory policies, ensuring that the proposed demand management strategies can be applied in real-time operational environments without causing performance bottlenecks.

In summary, we find a considerable improvement potential by applying availability control. A positive finding is that even “less sophisticated” forms of demand management, i.e., myopic, non-selective, and non-time shifting policies, already yield benefits compared to feasibility control, which makes a step-wise introduction viable.

5.2 Limitations

While our approach for strategic decision support on the selection of demand management policies yields promising results, there are several limitations that need to be considered when implementing our methodology in real-world settings:

- *Data availability*—Request arrival: To apply our approach, providers must have access to comprehensive historical request data. This data should include request attributes such as the time of request, desired pick-up and drop-off times, locations, and the number of passengers. It is essential that this data is uncensored: All customer requests, not just those that were successfully converted into orders, must be recorded, and the provider must not communicate any information on service availability before request placement.
- *Data availability*—Customer choice behavior: The semi-perfect information model assumes perfect information regarding customer segment affiliation, which is a strict assumption. It is necessary because typically, real-world data on customer preferences is incomplete or noisy, making it difficult for providers to estimate choice models accurately. To mitigate this limitation, it is important to conduct sensitivity analyses to account for the uncertainty around the true customer choice behavior (see Sect. 4.4).
- *External factors:* For anticipatory demand management, we assume that the provider uses information from historical request data. While this accounts for any external effects that occurred in the past, it may not fully account for future events affecting demand such as economic shifts. Hence, in systems with a volatile demand structure (e.g., newly established systems), the benefits of anticipatory demand management are likely lower than in the results of our computational study.

5.3 Future research directions

Existing research on the operation of rural SMOD systems is still scarce, and our results give rise to further novel research questions in this area:

- First, we see potential for developing anticipatory availability control policies that are tailored to the rural problem setting. Since algorithm development is not the focus of our work, we transferred existing algorithms that are practical, interpretable, and do not require parameter tuning. Also, we focused only on the availability control subproblem. Hence, e.g., by drawing on methods that involve statistical learning, explicitly consider displacement effects, or allow anticipatory routing decisions, there is still potential for algorithmic improvements.
- Second, the low-profitability rural environment makes subsidies by municipal contracting authorities a necessity. Since there are various possibilities for the design of subsidy schemes, future research could apply a similar methodological approach to investigate the impact of different subsidy schemes on the system performance. Such an approach could also provide decision support to authorities on how to design a subsidy scheme such that the profit-maximizing provider is incentivized to make operational decisions in a way that guarantees the authorities' sustainability objectives to be reached.
- Third, it could be investigated how demand management can be applied to SMOD systems with more complex, differentiated fulfillment option designs. Being able to offer and control, e.g., express rides, subscriptions, or group tickets could make the system more customer-centric and further increase its attractiveness compared to motorized individual transport.
- Fourth, future research is required conducting a holistic environmental sustainability assessment of rural SMOD systems. Besides the direct vehicle kilometer savings, we believe that it is particularly important to investigate indirect effects within the entire transportation system of rural areas that result from an SMOD system being in place. Examples are possible reductions of the private fleet size or line-based public transport, but also induced demand or cannibalization effects regarding more sustainable means of transport.
- Finally, while the work at hand focuses on demand management, we believe that it is promising to analogously analyze the impact of different policies for vehicle routing on the performance of rural SMOD systems. Building on these results, it is also of practical relevance to provide guidance on which combination of algorithmic elements provides the best performance depending on a limited computational budget that is available due to the requirement of real-time decision-making.

Appendix A: notation

See Tables 4, 5, 6 and 7.

Table 4 Notation Markov decision process model

| | |
|--|---|
| $t \in \mathcal{T} = \{1, \dots, t^s, \dots, T\}$ | Decision epoch |
| t^s | Start of the service horizon |
| $c \in \mathcal{C}$ | Customer request type |
| λ_c^t | Arrival rate of request type c in stage t |
| $p_c, d_c \in \mathcal{H}$ | Pick-up (drop-off) stop of request type c |
| m_c | Number of passengers of request type c |
| $t_c \in \{t^s, \dots, T\}$ | Desired time of request type c |
| $f_c \in \{0, 1\}$ | Indicator for time window type of request type c |
| r_c | Revenue of request type c |
| $i \in \mathcal{I}$ | Request |
| $\tau_i \in \mathcal{T}$ | Time of request for request i |
| $o \in \mathcal{O}_c$ | Fulfillment option defined for request type c |
| $(\tau_{c,o}^{e+}, \tau_{c,o}^{l+}), (\tau_{c,o}^{e-}, \tau_{c,o}^{l-})$ | Pick-up (drop-off) time window for request type c and option o |
| $g \subseteq \mathcal{O}_c$ | Offer set that can be presented to customer type c |
| $P_{c,o}(g)$ | Probability of customer placing a request of type c choosing option o when presented offer set g |
| $j \in \mathcal{J}$ | Order |
| $v \in \mathcal{V}$ | Vehicle |
| Q_v | Seat capacity of vehicle v |
| t_v^b, t_v^r | Start (end) of operations for vehicle v |
| ϕ_t | Route plan at decision epoch t |
| $\theta_{v,t} \in \phi_t$ | Planned route of vehicle v according to route plan ϕ_t |
| a_n^-, a_n^+ | Vehicle arrival (departure) time at the n -th stop in a route |
| s_t | Post-decision state at decision epoch t |
| C_t | Set of confirmed but not yet fulfilled orders at decision epoch t |
| a_t | Action at decision epoch t |
| $g_t \in \mathcal{G}(s_{t-1}, c)$ | Availability control decision at decision epoch t |
| $\phi_t(o)$ | Route plan at decision epoch t including a potential order resulting from combining the newly arrived request with option o |
| $(\phi_t(o))_{o \in g_t} \in \prod_{o \in g_t} \Phi(s_{t-1}, c, o)$ | Vehicle routing decision at decision epoch t for request type c availability control decision g_t |
| \mathcal{O}_c^f | Set of feasible fulfillment options for request type c |
| s_t^{pre} | Pre-decision state at decision epoch t |
| o_{j_t} | Fulfillment option chosen by the customer placing order j_t |
| s_t' | Interim state at decision epoch t |
| $\psi_{v,t}(\phi_t(o))$ | Stops that are visited definitively according to vehicle routing decision $\phi_t(o)$ |
| $\Psi_t(\phi_t(o))$ | Orders for which fulfillment is completed according to vehicle routing decision $\phi_t(o)$ |
| $r_{\phi_t(o)}$ | Vehicle-routing-related reward incurred by vehicle routing decision $\phi_t(o)$ |
| $\rho_{h,h'}$ | Routing cost for traveling from stop h to stop h' |

Table 5 Notation of Markov decision process model (continued)

| | |
|-----------------------------|--|
| $V_t(s_t)$ | Value of post-decision state s_t |
| $V'_t(s'_t)$ | Value of interim state s'_t |
| $\Delta V_t(s_{t-1}, c, o)$ | Opportunity cost of an order by request type c with fulfillment option o |
| $l \in \mathcal{L}$ | Customer segment |
| $S_{l,c}$ | Consideration set of customer segment l and request type c |
| $\Delta^+_l + \Delta^-_l$ | Total time flexibility provided by segment l |
| ζ | Ranking function over fulfillment options |
| o_{cgl} | Fulfillment option chosen by request type c and segment l from offer set g |
| γ_l | Share of segment l in the customer population |

Table 6 Notation of static Dial-a-Ride problem model

| | |
|--|--|
| $k \in \mathcal{P} \cup \mathcal{D} = \mathcal{N}$ | Pick-up/drop-off node |
| $\sigma_{kk'}$ | Distance for traveling from node k to node k' |
| $\delta_{kk'}$ | Time for traveling from node k to node k' |
| \mathcal{A} | Set of arcs |
| (τ_k^e, τ_k^l) | Time window of node k |
| k_v | Node at which vehicle v becomes available |
| q_k | Number of passengers picked-up or dropped off at node k |
| $x_{kk'v}$ | Binary decision variable indicating whether vehicle v travels from node k to node k' |
| B_{kv} | Time at which vehicle v stops at node k |
| Q_{kv} | Load of vehicle v when leaving node k |
| t_c^{mart} | Maximum added ride time for request type c |
| ω | Waiting time |
| μ | Maximum added ride time factor |

Table 7 Notation of solution algorithms

| | |
|---|---|
| C_j^{act} | Set of actual orders |
| ϕ_t^{act} | Actual route plan at decision epoch t |
| $\phi_t^{act}(o)$ | Actual route plan at decision epoch t including a potential order resulting from combining the newly arrived request with option o |
| C_t^{sam} | Set of sampled orders |
| ϕ_t^{sam} | Sampled route plan at decision epoch t |
| $\phi_t^{sam}(o)$ | Sampled route plan at decision epoch t including a potential order resulting from combining the newly arrived request with option o |
| AR^{sam} | Sampling acceptance rate |
| $\widetilde{\Delta V}_t(s_{t-1}, c, o)$ | Approximation of opportunity cost of an order by request type c with fulfillment option o |
| $\phi_{t,j}^{sam}$ | Sampled route plan at decision epoch t excluding sampled order j |
| j^* | Sampled order with the highest cost saving |
| $\mathcal{C}_{c_{ij}}^{pro}$ | Subset of profit-maximizing fulfillment options |
| $\mathcal{C}_{c_{ij}}^{clo}$ | Subset of profit-maximizing fulfillment options closest to the desired time |

Appendix B: mixed-integer-program for the static dial-a-ride problem

Determining the action space $\prod_{o \in g_t} \Phi(s_{t-1}, c, o)$ for the vehicle routing decisions $(\phi_t(o))_{o \in g_t}$ in the MDPs for the r-SMCP described in Sect. 3.2 corresponds to searching all solutions to $|\mathcal{O}_c|$ constraint satisfaction problems. Each of these problems (CS-DARP) has a structure similar to the static DARP. The instance is given by the set of unfulfilled orders C_{t-1} , the current vehicle positions stored in ϕ_{t-1} , and the potential order resulting from assigning the newly received request i_t a fulfillment option o . In the following, we present a mixed-integer programming model for the CS-DARP based on the DARP formulation by Cordeau (2006) and describe how its parameters can be determined from the information given in state s_{t-1} .

The model is based on a graph $G = (\mathcal{N}, \mathcal{A})$ consisting of a set of nodes \mathcal{N} and a set of arcs \mathcal{A} . The set of nodes $\mathcal{N} = \{0\} \cup \mathcal{P} \cup \mathcal{D} \cup \{2|\mathcal{J}| + 1\}$ contains a pick-up node $k = j \in \mathcal{P}$ and a drop-off node $k = (j + |\mathcal{J}|) \in \mathcal{D}$ for each order $j \in \mathcal{J}$ in addition to the origin depot node $k = 0$ and the destination depot node $k = 2|\mathcal{J}| + 1$. The geographical location of each node is given by the pick-up stop p_{c_j} and the drop-off stop d_{c_j} . This mapping of nodes to stops allows the computation of travel distances $\sigma_{kk'}$, travel times $\delta_{kk'}$, and travel costs $\rho_{kk'}$ between two nodes $k, k' \in \mathcal{N}$. The three parameters are weights of the respective arcs $(k, k') \in \mathcal{A} = \{(k, k') : k = 0, k' \in \mathcal{P} \vee (k, k' \in \mathcal{P} \cup \mathcal{D} \wedge k \neq k' \wedge k \neq k' + |\mathcal{J}|) \vee k \in \mathcal{D}, k' = 2|\mathcal{J}| + 1\}$. The time window of each node k is defined by two time points marking its start τ_k^e and its end τ_k^l . It is equal to the order's pick-up time window defined by the earliest pick-up time τ_k^{e+} and the latest pick-up time τ_k^{l+} in case of $k \in \mathcal{P}$ or the drop-off time window defined by the earliest drop-off time τ_k^{e-} and the latest drop-off time τ_k^{l-} in case $k \in \mathcal{D}$. Similarly, the number of passengers q_k that are picked-up or dropped-off at node k can be computed based on the number of passengers m_c .

The unfulfilled orders for the vehicle routing decision at decision epoch t are stored in C_{t-1} . A subset of these orders may be partly fulfilled, meaning that the passengers are already on board a vehicle. Hence, to derive the set of pick-up nodes \mathcal{P} , we remove these orders based on the route plan ϕ_{t-1} such that only indices of orders remain that have not yet been picked up:

$$\mathcal{P} = C_{t-1} \setminus \{k \in C_{t-1} : a_k^+ \leq t, (k, h_k, a_k^-, a_k^+) \in \theta_{v,t-1}, \theta_{v,t-1} \in \phi_{t-1}\} \quad (11)$$

The set of drop-off nodes \mathcal{D} contains one node for each unfulfilled order:

$$\mathcal{D} = \{j + |\mathcal{J}| : j \in C_{t-1}\} \quad (12)$$

Each vehicle $v \in \mathcal{V}$ starts its route from the origin node $k_v = k_1$ with $(k_1, h_{k_1}, a_{k_1}^-, a_{k_1}^+) \in \theta_{t-1,v}, \theta_{t-1,v} \in \phi_{t-1}$, at which it next becomes available according to the arrival times a_k^- stored in ϕ_{t-1} .

Similarly to the set of pick-up nodes, we can also derive the initial vehicle load q_{k_v} from the route plan ϕ_{t-1} :

$$q_{k_v} := q_{k_v} + \sum_{\{k \in C_{t-1} : a_k^+ \leq t, (k, h_k, a_k^-, a_k^+) \in \theta_{t-1, v}\}} m_{c_k} \quad (13)$$

In summary, we have the set of nodes $\mathcal{N} = \mathcal{P} \cup \mathcal{D} \cup \{2|\mathcal{J}| + 1\}$, which also includes the destination depot node $k = 2|\mathcal{J}| + 1$ where all vehicles must finish their route. The binary decision variables $x_{kk'v}$ encode whether vehicle v drives directly from node k to node k' ($x_{kk'v} = 1$) or not ($x_{kk'v} = 0$). Further, decision variables B_{kv} encode the time at which vehicle v stops at node k , and decision variables Q_{kv} encode the load of vehicle v when leaving node k .

$$\sum_{v \in \mathcal{V}} \sum_{k' \in \mathcal{N}} x_{kk'v} = 1 \quad \forall k \in \mathcal{P} \quad (14)$$

$$\sum_{k' \in \mathcal{N}} x_{kk'v} - \sum_{k' \in \mathcal{N}} x_{|J|+k, k'v} = 0 \quad \forall k \in \mathcal{P}, v \in \mathcal{V} \quad (15)$$

$$\sum_{k' \in \mathcal{N}} x_{k_v, k'v} = 1 \quad \forall v \in \mathcal{V} \quad (16)$$

$$\sum_{k' \in \mathcal{N}} x_{k'kv} - \sum_{k' \in \mathcal{N}} x_{kk'v} = 0 \quad \forall k \in (\mathcal{P} \cup \mathcal{D}) \setminus \mathcal{K}, v \in \mathcal{V} \quad (17)$$

$$\sum_{k \in \mathcal{N}} x_{k, 2|\mathcal{J}|+1, v} = 1 \quad \forall v \in \mathcal{V} \quad (18)$$

$$B_{k'v} \geq (B_{kv} + \delta_{kk'})x_{kk'v} \quad \forall k \in \mathcal{N}, k' \in \mathcal{N}, v \in \mathcal{V} \quad (19)$$

$$Q_{k'v} \geq (Q_{kv} + q_{k'})x_{kk'v} \quad \forall k \in \mathcal{N}, k' \in \mathcal{N}, v \in \mathcal{V} \quad (20)$$

$$B_{|J|+k, v} - B_{k, v} \geq \delta_{k, |J|+k} \quad \forall k \in \mathcal{P}, v \in \mathcal{V} \quad (21)$$

$$\tau_k^e \leq B_{kv} \leq \tau_k^l \quad \forall k \in \mathcal{N} \setminus \mathcal{K}, v \in \mathcal{V} \quad (22)$$

$$a_v \leq B_{k_v, v} \quad \forall v \in \mathcal{V} \quad (23)$$

$$\max\{0, q_k\} \leq Q_{kv} \leq \min\{\kappa_v, \kappa_v + q_k\} \quad \forall k \in \mathcal{N}, v \in \mathcal{V} \quad (24)$$

$$x_{kk'v} \in \{0, 1\} \quad \forall k \in \mathcal{N}, k' \in \mathcal{N}, v \in \mathcal{V} \quad (25)$$

Constraints (14) make sure that each order's pick-up node is visited by exactly one vehicle. Constraints (15) enforce that this vehicle also visits the corresponding drop-off node. Flow conservation is guaranteed by Constraints (16)–(18) for the vehicles' origin nodes, the remaining pick-up nodes and drop-off nodes, and the destination depot node, respectively. Consistency regarding time flow and loads is guaranteed

by Constraints (19) and (20), respectively, which can be straightforwardly linearized. Constraints (21) ensure that pick-up nodes are visited before drop-off nodes for all orders. Constraints (22) prevent any time window violations and thereby also prevent violations of the maximum ride time, since it is included in the time window definition (see Appendix C). The vehicles' time of availability is considered through Constraints (23). Finally, Constraints (24) prevent violations of the seat capacity.

Appendix C: time window generation for the dial-a-ride problem

In the following, we briefly define how the time windows for pick-up $\left(\tau_{c,o}^{e+}, \tau_{c,o}^{l+}\right)$ and drop-off $\left(\tau_{c,o}^{e-}, \tau_{c,o}^{l-}\right)$ are computed based on the desired time t_c (Jaw et al., 1986). To guarantee a certain service level for all types of requests $c \in \mathcal{C}$, we define a maximum added time t_c^{mart} to the direct ride time δ_{p_c, d_c} from the pick-up stop p_c to the drop-off stop d_c . It consists of a constant waiting time ω and a certain fraction μ of the direct ride time. Thus, $t_c^{mart} = \omega + (1 + \mu)\delta_{j,j+|\mathcal{I}|}$. Further, each node is assigned a time window $\left(\tau_k^e, \tau_k^l\right)$ resulting from $\delta_{j,j+|\mathcal{I}|}$, ω , μ , and the trip type f_c . The respective formulae are given in Table 8. Finally, each node is associated with a weight q_k indicating the number of passengers to be picked up ($q_k > 0$) or dropped off ($q_k < 0$), which is given by m_{c_j} for each order j .

See Table 8.

Appendix D: example for the stochastic modeling component of the operational MDP and semi-perfect information model

Consider a brief numerical example that illustrates the differences between the stochastic modeling component of the operational MDP and the semi-perfect information model. For simplicity, we assume a three-period time horizon $\mathcal{T} = \{1, 2, 3\}$ with an incoming request of type $c = 1$ and a desired time $t_1 = 2$. The feasible fulfillment options $o \in \mathcal{O}_c^f = \{0, 1, 2, 3\}$ represent (alternative) times that the provider can offer in response to the customer's desired time $t_1 = 2$, specifically $o = 1$ and $o = 3$, as well as the no-purchase option $o = 0$.

We assume that the customer population is divided into two customer segments $\mathcal{L} = \{1, 2\}$ with shares of $\gamma_1 = 0.4$ for segment $l = 1$ and $\gamma_2 = 0.6$ for segment $l = 2$: Customers of segment 1 will only accept alternative times earlier than their desired time, and customers of segment 2 will only accept alternative times later than their

Table 8 Computation of time windows

| | $f_c = 0$ (outbound trip) | $f_c = 1$ (inbound trip) |
|-------------------------------------|--|---|
| $k \in \mathcal{P}$ (pick-up node) | $\tau_{c,o}^{e+} = o - \delta_{j,j+ \mathcal{I} } - t_c^{mart}$, $\tau_{c,o}^{l+} = o - \delta_{j,j+ \mathcal{I} }$ | $\tau_{c,o}^{e+} = o$, $\tau_{c,o}^{l+} = o + \omega$ |
| $k \in \mathcal{D}$ (drop-off node) | $\tau_{c,o}^{e-} = o - \omega$, $\tau_{c,o}^{l-} = o$ | $\tau_{c,o}^{e-} = o + \delta_{j,j+ \mathcal{I} }$, $\tau_{c,o}^{l-} = o + \delta_{j,j+ \mathcal{I} } + t_c^{mart}$ |

desired time. Both, of course, consider the no-purchase option. This leads to the following consideration sets: $\mathcal{S}_{1,1} = \{0,1,2\}$, $\mathcal{S}_{2,1} = \{0,2,3\}$.

We assume a ranking function $\zeta = (2,1,3,0)$ that represents the general preferences within consideration sets, with the desired time being the most preferred option, followed by alternative times that are less preferred with increasing deviation, and lastly the no-purchase option. For alternative times with the same deviation, e.g., $o = 1$ and $o = 3$, alternative times earlier than the desired time are favored.

Given an offer set $g = \{0,1\}$, we can illustrate the differences of the stochastic modeling component for the operational MDP and the semi-perfect information model.

Operational MDP: According to the shares of the segments γ_1 and γ_2 , we can calculate the probabilities that a customer with request type $c = 1$ will choose an option. A customer from segment $l = 1$ will choose option 1, since the customer is willing to accept a time deviation before $t_1 = 2$. A customer of segment $l = 2$ will choose the no-purchase option. This results in the following probabilities, derived from the segments' share of the customer population:

$$P_{1,0}(\{0,1\}) = 0.6$$

$$P_{1,1}(\{0,1\}) = 0.4$$

$$P_{1,2}(\{0,1\}) = P_{1,3}(\{0,1\}) = 0$$

To summarize, for the operational MDP we first observe a deterministic transition by the provider choosing an offer set $g = \{0,1\}$, followed by a stochastic transition according to the customer's choice behavior, resulting in the choice probabilities $P_{1,o}(\{0,1\})$ (see also Fig. 3).

Semi-perfect information model: Contrary to the operational MDP, in the semi-perfect information model, the provider knows the true segment affiliation of each customer request. Assuming that the customer belongs to segment $l = 1$, the stochastic component of the choice behavior is eliminated: The provider knows with certainty ($P_{1,1}(\{0,1\}) = 1$), that the customer will choose option $o = 1$ and we can omit the choice probabilities for other options ($P_{1,0}(\{0,1\}) = P_{1,2}(\{0,1\}) = P_{1,3}(\{0,1\}) = 0$). This implies that this transition in the semi-perfect information model is completely deterministic as the provider can effectively assign any fulfillment option from the customer's consideration set by choosing an offer set that makes this option the customer's most preferred one in the offer set (see also Fig. 3).

Appendix E: feasibility check

See Fig. 16.

| | |
|---|---|
| Input: Current actual route plan ϕ_t^{act} , set of feasible options $\mathcal{O}_{c_{i_t}}^f$, current request i_t , option o | |
| <hr/> | |
| 1 | $\phi_t^{act}(o) := \text{parallel_insertion}(\phi_t^{act}, (i_t, o))$ |
| 2 | if (i_t, o) in $\phi_t^{act}(o)$ do |
| 3 | $\mathcal{O}_{c_{i_t}}^f := \mathcal{O}_{c_{i_t}}^f \cup \{o\}$ |

Fig. 16 Feasibility check function

Appendix F: availability control

See Fig. 17.

| | |
|--|--|
| Input: Current customer request i_t , set of cost estimates $\{\Delta \tilde{V}_t(s_{t-1}, c_{i_t}, o) : o \in \mathcal{O}_{c_{i_t}}^f\}$ | |
| <hr/> | |
| 1 | $\mathcal{O}_{c_{i_t}}^{pro} := \left\{ o' \in \mathcal{O}_{c_{i_t}}^f : o' = \underset{o \in \mathcal{O}_{c_{i_t}}^f}{\operatorname{argmax}} \{r_{c_{i_t}} - \Delta \tilde{V}_t(s_{t-1}, c_{i_t}, o)\}, r_{c_{i_t}} - \Delta \tilde{V}_t(s_{t-1}, c_{i_t}, o') \geq 0 \right\}$ |
| 2 | $\mathcal{O}_{c_{i_t}}^{clo} := \left\{ o' \in \mathcal{O}_{c_{i_t}}^{pro} : o' = \underset{o \in \mathcal{O}_{c_{i_t}}^{pro}}{\operatorname{argmin}} \{ o - t_{c_{i_t}} \} \right\}$ |
| 3 | $o_{j_t} := \min_{o \in \mathcal{O}_{c_{i_t}}^{clo}} o$ |

Fig. 17 Availability control function

Appendix G: routing control

See Fig. 18.

| | |
|--|---|
| Input: Actual route plan including new order $\phi_t^{act}(o_{j_t})$, set of actual orders C_t^{act} | |
| <hr/> | |
| 1 | $\phi_t^{act} := \phi_t^{act}(o_{j_t})$ |
| 2 | if $t = t^s - 1$ do |
| 3 | $\phi_t^{act} := \{\theta_v = \{(0, 0, t_v^b, t_v^b), (0, 0, t_v^r, t_v^r)\} : v \in \mathcal{V}\}$ |
| 4 | $\phi_t^{act} := \text{parallel_insertion}(\phi_t^{act}, C_t^{act})$ |

Fig. 18 Routing control function

Appendix H: additional functions for anticipatory control

In the following, we first describe the initialization and second the additional computations of the anticipatory policies at each decision epoch in a more detailed fashion.

Initialization: The basic steps of the initialization are given in lines 2–4 of Fig. 4. First the sampled route plan is initialized as empty (line 2). Then, the set of orders C_0^{sam} for the skeletal route planning is sampled (line 3). The pool of historical requests from which we sample is defined hierarchically by three attributes of the service day that is controlled, which are encoded by *day_type*. Primarily, we check whether the service day is a public holiday and, if so, sample from a pool of all public holidays. On the second level, we check whether the service day is a school vacation day. Finally, we evaluate the day of week. To determine the size of the sample, we multiply the average number of requests received per day over all days in the sampling pool with a fictive acceptance rate AR^{sam} , which is the only parameter of the AC. From each of the sampled requests, we generate a sampled order by assigning the desired time as the fulfillment option. In the final step of the initialization (line 4), we generate the initial sampled route plan from the set of sampled orders C_0^{sam} by means of the parallel insertion heuristic.

Iterations: At each decision epoch, there are three basic steps associated with solving the opportunity cost estimation subproblem in the AC (lines 7, 12, and 17). First, the sampled route plan is synchronized with the actual route plan (Fig. 19). This step ensures that the cost estimate is based on the actual positions of the vehicles and the actual (tentative) routing decisions. Consequently, it is only required at decision epochs within the service horizon (line 1). The synchronized sampled route plan ϕ_t^{sam} is initialized as a copy of the actual route plan ϕ_t^{act} . Then, the sampled orders remaining at the preceding decision epoch C_{t-1}^{sam} are inserted into ϕ_t^{sam} (line 3). Those sampled orders that cannot be feasibly inserted any more, e. g., because they are expired, are not included in the updated set of sampled orders C_t^{sam} (lines 4–6).

See Fig. 19.

The computation of the cost estimate in the second basic step is shown in Fig. 20. First, the potential order resulting from the assignment of option o to request i_t is inserted into the current sampled route plan $\phi_t^{sam}(o)$ (line 1). If the insertion is not

Input: Current actual route plan ϕ_t^{act} , current set of sampled orders C_{t-1}^{sam}

```

1   if  $t \geq t^s$  do
2        $\phi_t^{sam} := \phi_t^{act}$ 
3        $\phi_t^{sam} := \text{parallel\_insertion}(\phi_t^{sam}, C_{t-1}^{sam})$ 
4       forall  $j \in C_{t-1}^{sam}$  do
5           if  $j$  in  $\phi_t^{sam}$  do
6                $C_t^{sam} := C_t^{sam} \cup \{j\}$ 

```

Fig. 19 Route plan synchronization function

Input: Current sampled route plan ϕ_t^{sam} , current request i_t , option o

```

1   $\phi_t^{sam}(o) := parallel\_insertion(\phi_t^{sam}, (i_t, o))$ 
2  if  $\phi_t^{sam}(o) = \phi_t^{sam}$  do
3       $\Delta \tilde{V}_t(s_{t-1}, c_{i_t}, o) := \infty$ 
4  else do
5       $\Delta \tilde{V}_t(s_{t-1}, c_{i_t}, o) := cost(\phi_t^{sam}(o)) - cost(\phi_t^{sam})$ 

```

Fig. 20 Opportunity cost estimation function (anticipatory control)

feasible, the cost estimate is set to a sufficiently high value such that the option is guaranteed to not be offered (line 3), which is similar to the approach by Yang et al. (2016). Otherwise, the cost estimate is equal to the cheapest insertion cost in the sampled route plan (line 5).

See Fig. 20.

As a third step at each decision epoch, the sampled route plan is updated such that the newly confirmed order j_t replaces one of the sampled orders (Fig. 21). This approach is also used by Koch and Klein (2020). If a new order is confirmed, the sampled route plan including this order $\phi_t^{sam}(o_{j_t})$ already determined in the preceding step is used further (line 2). Then, each sampled order is preliminary removed from the sampled route plan to evaluate the associated cost savings (line 4). Finally, the algorithm permanently removes the sampled order j^* with the greatest cost saving (lines 5 and 6). The reasoning behind this rule is as follows: The greater the cost estimate for a sampled order, the more likely a similar actual request potentially arriving in the future would be rejected. Therefore, sampled orders with a high cost estimate forecast consolidation opportunities that are very unlikely to realize and would distort the cost estimates for the arriving request. By incorporating this rule, we not only anticipate future demand but also future decision-making.

See Fig. 21.

Input: Sampled route plan including new order $\phi_t^{sam}(o_{j_t})$, set of sampled orders C_t^{sam}

```

1  if  $o_{j_t} \neq 0$  do
2       $\phi_t^{sam} := \phi_t^{sam}(o_{j_t})$ 
3  forall  $j \in C_t^{sam}$  do
4       $\phi_{t,j}^{sam} := remove(\phi_t^{sam}, j)$ 
5       $j^* := \operatorname{argmin}_{j \in C_t^{sam}} \{cost(\phi_{t,j}^{sam})\}$ 
6       $\phi_t^{sam} := \phi_{t,j^*}^{sam}$ 

```

Fig. 21 Update sampled route plan function (anticipatory control)




Fig. 23 Use of rejections depending on the time of request: The horizontal axis plots the percentage of requests arrived and the vertical axis plots the rate of acceptances, **a** feasibility rejections, **b** and profitability rejections. **c** Each series corresponds to one of the policies NS-MC, NT-MC, NS-AC, and NT-AC




Fig. 24 Use of time shifts depending on the time of request: The horizontal axis plots the percentage of requests arrived and the vertical axis plots the average total time shift, **a** feasibility time shift, **b** and profitability time shift per order in minutes. **c** Each series corresponds to one of the policies NS-MC, NT-MC, NS-AC, and NT-AC




Fig. 25 Use of rejections depending on the desired time: The horizontal axis plots the desired time and the vertical axis plots the rate of acceptances, **a** feasibility rejections, **b** and profitability rejections. **c** Each series corresponds to one of the policies NS-MC, NT-MC, NS-AC, and NT-AC




Fig. 26 Use of rejections depending on the OD-pair length: The horizontal axis plots the direct distance between pick-up and drop-off in km and the vertical axis plots the average total time shift, **a** feasibility time shift per order in minutes, **b** and profitability time shift per order in minutes. **c** Each series corresponds to one of the policies NS-MC, NT-MC, NS-AC, and NT-AC

Acknowledgements We sincerely thank the FLEXIBUS KG, in particular Josef Brandner and Daniel Mayer, for providing access to the real-world data set used in this study. In addition, we highly appreciate their valuable support, the insightful discussions, and the practical experiences they shared with us.

Data availability The real-world data set provided by the FLEXIBUS KG is not publicly available since it contains confidential company data.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Allcott H (2013) The welfare effects of misperceived product costs: Data and calibrations from the automobile market. *Am Econ J Econ Pol* 5(3):30–66. <https://doi.org/10.1257/pol.5.3.30>
- Alonso-González MJ, Liu T, Cats O, Van Oort N, Hoogendoorn S (2018) The potential of demand-responsive transport as a complement to public transport: an assessment framework and an empirical evaluation. *Transp Res Rec* 2672(8):879–889. <https://doi.org/10.1177/0361198118790842>
- Alonso-Mora J, Samaranayake S, Wallar A, Frazzoli E, Rus D (2017) On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment. *Proc Natl Acad Sci USA* 114(3):462–467. <https://doi.org/10.1073/pnas.1611675114>
- Aouad A, Farias V, Levi R (2021) Assortment optimization under consider-then-choose choice models. *Manage Sci* 67(6):3368–3386. <https://doi.org/10.1287/mnsc.2020.3681>
- Araldo A, Di Maria A, Di Stefano A, Morana G (2019) On the importance of demand consolidation in mobility on demand. *2019 IEEE/ACM 23rd International Symposium on Distributed Simulation and Real Time Applications (DS-RT)*. 1–8. <https://doi.org/10.1109/DS-RT47707.2019.8958669>
- Arian E, Bai X, Chen X (2022) Joint pricing and routing for a ride-sharing platform in low-density rural areas. Working paper, University of Illinois at Urbana-Champaign. <https://doi.org/10.2139/ssrn.4073828>
- Asatryan H, Gaul D, Gottschalk H, Klamroth K, Stiglmayr M (2023). Ridepooling and public bus services: A comparative case-study. Working paper, Bochum University of Applied Sciences.
- Atasoy B, Ikeda T, Song X, Ben-Akiva ME (2015) The concept and impact analysis of a flexible mobility on demand system. *Trans Res Part C: Emerg Technol* 56:373–392. <https://doi.org/10.1016/j.trc.2015.04.009>
- Attanasio A, Cordeau JF, Ghiani G, Laporte G (2004) Parallel tabu search heuristics for the dynamic multi-vehicle dial-a-ride problem. *Parallel Comput* 30(3):377–387. <https://doi.org/10.1016/j.parco.2003.12.001>
- Avraham E, Raviv T (2021) The steady-state mobile personnel booking problem. *Transp Res Part B: Methodol* 154:266–288. <https://doi.org/10.1016/j.trb.2021.10.008>
- Banister D (2011) Cities, mobility and climate change. *J Transp Geogr* 19(6):1538–1546. <https://doi.org/10.1016/j.jtrangeo.2011.03.009>
- Bar-Yosef A, Martens K, Benenson I (2013) A model of the vicious cycle of a bus line. *Transp Res Part B: Methodol* 54:37–50. <https://doi.org/10.1016/j.trb.2013.03.010>
- Bertsimas D, Jaillet P, Martin S (2019) Online vehicle routing: The edge of optimization in large-scale applications. *Op Res* 67(1):143–162. <https://doi.org/10.1287/opre.2018.1763>
- Bischoff J, Maciejewski M, Nagel K (2017) City-wide shared taxis: A simulation study in Berlin. In: 2017 IEEE 20th international conference on intelligent transportation systems (ITSC), 275–280. <https://doi.org/10.1109/ITSC.2017.8317926>.

- Campbell AM, Savelsbergh MWP (2005) Decision support for consumer direct grocery initiatives. *Transp Sci* 39(3):313–327. <https://doi.org/10.1287/trsc.1040.0105>
- Campbell AM, Savelsbergh MWP (2006) Incentive schemes for attended home delivery services. *Transp Sci* 40(3):327–341. <https://doi.org/10.1287/trsc.1050.0136>
- Cordeau JF (2006) A branch-and-cut algorithm for the dial-a-ride problem. *Op Res* 54(3):573–586. <https://doi.org/10.1287/opre.1060.0283>
- Coutinho FM, van Oort N, Christoforou Z, Alonso-González MJ, Cats O, Hoogendoorn S (2020) Impacts of replacing a fixed public transport line by a demand responsive transport system: case study of a rural area in Amsterdam. *Res Transp Econ* 83:100910. <https://doi.org/10.1016/j.retrec.2020.100910>
- Currie G, Fournier N (2020) Why most DRT/Micro-Transits fail—what the survivors tell us about progress. *Res Transp Econ* 83:100895. <https://doi.org/10.1016/j.retrec.2020.100895>
- EEA (2023) Greenhouse gas emissions from transport in Europe. Retrieved from <https://www.eea.europa.eu/en/analysis/indicators/greenhouse-gas-emissions-from-transport>. Last accessed: 15/12/2023.
- Elting S, Ehmke JF (2021) Potential of shared taxi services in rural areas—a case study. *Transp Res Proced* 52:661–668. <https://doi.org/10.1016/j.trpro.2021.01.079>
- EPA (2023) U.S. transportation sector greenhouse gas emissions. Retrieved from <https://www.epa.gov/system/files/documents/2023-06/420f23016.pdf>. Last accessed: 15/12/2023.
- Eurostat (2022) Key figures on European transport. Retrieved from <https://ec.europa.eu/eurostat/documents/15216629/15589759/KS-07-22-523-EN-N.pdf>. Last accessed: 15/12/2023.
- Fleckenstein D, Klein R, Steinhardt C (2023) Recent advances in integrating demand management and vehicle routing: a methodological review. *Eur J Oper Res* 306(2):499–518. <https://doi.org/10.1016/j.ejor.2022.04.032>
- Fleckenstein D, Klein R, Klein V, Steinhardt C (2024) On the concept of opportunity cost in integrated demand management and vehicle routing. Working paper, University of Augsburg.
- FLEXIBUS (2024). Der FLEXIBUS in Krumbach. Retrieved from <https://www.flexibus.net/krumbach>. Last accessed: 05/07/2024.
- Gaul D, Klamroth K, Stiglmayr M (2022) Event-based MILP models for ridepooling applications. *Eur J Oper Res* 301(3):1048–1063. <https://doi.org/10.1016/j.ejor.2021.11.053>
- Gilibert M, Ribas I, Rosen C, Siebeneich A (2020) On-demand shared ride-hailing for commuting purposes: comparison of Barcelona and Hanover case studies. *Transp Res Proced* 47:323–330. <https://doi.org/10.1016/j.trpro.2020.03.105>
- Gössling S, Kees J, Litman T (2022) The lifetime cost of driving a car. *Ecol Econ* 194:107335. <https://doi.org/10.1016/j.ecolecon.2021.107335>
- Haferkamp J, Ehmke JF (2020) An efficient insertion heuristic for on-demand ridesharing services. *Transp Res Proced* 47:107–114. <https://doi.org/10.1016/j.trpro.2020.03.074>
- Haferkamp J, Ehmke JF (2022) Effectiveness of demand and fulfillment control in dynamic fleet management of ride-sharing systems. *Networks* 79(3):314–337. <https://doi.org/10.1002/net.22062>
- Hauser JR (2014) Consideration-set heuristics. *J Bus Res* 67(8):1688–1699. <https://doi.org/10.1016/j.jbusres.2014.02.015>
- Heger J, Klein R (2024) Assortment optimization: a systematic literature review. *Or Spectr*, Online First. <https://doi.org/10.1007/s00291-024-00752-4>
- Heitmann RJO, Soeffker N, Ulmer MW, Mattfeld DC (2023) Combining value function approximation and multiple scenario approach for the effective management of ride-hailing services. *EURO J Transp Logist* 12:100104. <https://doi.org/10.1016/j.ejtl.2023.100104>
- Horn ME (2002) Fleet scheduling and dispatching for demand-responsive passenger services. *Transp Res Part C: Emerg Technol* 10(1):35–63. [https://doi.org/10.1016/S0968-090X\(01\)00003-1](https://doi.org/10.1016/S0968-090X(01)00003-1)
- Hosni H, Naoum-Sawaya J, Artail H (2014) The shared-taxi problem: formulation and solution methods. *Transp Res Part B: Methodol* 70:303–318. <https://doi.org/10.1016/j.trb.2014.09.011>
- Hungerländer P, Maier K, Pachatz V, Truden C (2021) Improving sharing rates of a dial-a-ride problem implemented for an austrian mobility provider. *Transp Res Proced* 52:525–532. <https://doi.org/10.1016/j.trpro.2021.01.062>
- Imhof S, Blättler K (2023) Assessing spatial characteristics to predict DRT demand in rural Switzerland. *Res Transp Econ* 99:101301. <https://doi.org/10.1016/j.retrec.2023.101301>
- ioki (2024) References. Retrieved from <https://ioki.com/en/references/>. Last accessed: 22/10/2024.
- Jung J, Jayakrishnan R, Park JY (2016) Dynamic shared-taxi dispatch algorithm with hybrid-simulated annealing. *Comput-Aided Civ Infrastruct Eng* 31(4):275–291. <https://doi.org/10.1111/mice.12157>
- Jung J, Jayakrishnan R, Choi K (2012) Shared-taxi operations with electric vehicles. Working paper, University of California.

- Klein V, Steinhardt C (2023) Dynamic demand management and online tour planning for same-day delivery. *Eur J Oper Res* 307(2):860–886. <https://doi.org/10.1016/j.ejor.2022.09.011>
- Klein R, Koch S, Steinhardt C, Strauss AK (2020) A review of revenue management: recent generalizations and advances in industry applications. *Eur J Oper Res* 284(2):397–412. <https://doi.org/10.1016/j.ejor.2019.06.034>
- Koch S, Klein R (2020) Route-based approximate dynamic programming for dynamic pricing in attended home delivery. *Eur J Oper Res* 287(2):633–652. <https://doi.org/10.1016/j.ejor.2020.04.002>
- Köhler C, Campbell AM, Ehmke JF (2024) Data-driven customer acceptance for attended home delivery. *Or Spectr* 46(2):295–330. <https://doi.org/10.1007/s00291-023-00712-4>
- Jaw JJ, Odoni AR, Psaraftis HN, Wilson NH (1986) A heuristic algorithm for the multi-vehicle advance request dial-a-ride problem with time windows. *Transp Res Part B: Methodol* 20(3):243–257. [https://doi.org/10.1016/0191-2615\(86\)90020-2](https://doi.org/10.1016/0191-2615(86)90020-2)
- Lotfi S, Abdelghany K (2022) Ride matching and vehicle routing for on-demand mobility services. *J Heuristics* 28(3):235–258. <https://doi.org/10.1007/s10732-022-09491-7>
- Lotze C, Marszal P, Jung F, Manik D, Timme M, Schröder M (2023) Identifying the threshold to sustainable ridepooling. Working paper, Technical University Dresden.
- Lu C, Maciejewski M, Wu H, Nagel K (2023) Demand-responsive transport for students in rural areas: a case study in Vulkaneifel, Germany. *Transp Res Part A: Policy Pract* 178:103837. <https://doi.org/10.1016/j.tra.2023.103837>
- Ma Z, Koutsopoulos HN (2022) Near-on-demand mobility. The benefits of user flexibility for ride-pooling services. *Transp Res Part C Emerg Technol* 135:103530. <https://doi.org/10.1016/j.trc.2021.103530>
- Melis L, Sörensen K (2022) The real-time on-demand bus routing problem: the cost of dynamic requests. *Comput Oper Res* 147:105941. <https://doi.org/10.1016/j.cor.2022.105941>
- Padam Mobility (2024). Our references. Retrieved from <https://www.padam-mobility.com/en/success-stories>. Last accessed: 22/10/2024.
- Mortazavi A, Ghasri M, Ray T (2024) Integrated demand responsive transport in low-demand areas: a case study of Canberra, Australia. *Transp Res Part D: Transp Environ* 127:104036. <https://doi.org/10.1016/j.trd.2023.104036>
- Nobis C, Kuhnimhof T (2018) Mobilität in Deutschland – Ergebnisbericht. Retrieved from https://www.mobilitaet-in-deutschland.de/archive/pdf/MiD2017_Ergebnisbericht.pdf. Last accessed: 18/12/2023.
- OSRM (n.d.). Open source routing machine. <https://project-osrm.org/>. Last accessed: 18/12/2023.
- European Parliament (2021) Transport infrastructure in low-density and depopulating areas. Retrieved from [https://www.europarl.europa.eu/RegData/etudes/STUD/2021/652227/IPOL_STU\(2021\)652227_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/STUD/2021/652227/IPOL_STU(2021)652227_EN.pdf). Last accessed: 15/12/2023.
- Poltimäe H, Rehema M, Raun J, Poom A (2022) In search of sustainable and inclusive mobility solutions for rural areas. *Eur Transp Res Rev* 14(1):13. <https://doi.org/10.1186/s12544-022-00536-3>
- Powell WB (2019) A unified framework for stochastic optimization. *Eur J Op Res* 275(3):795–821. <https://doi.org/10.1016/j.ejor.2018.07.014>
- Powell WB (2022) Reinforcement Learning and Stochastic Optimization: A unified framework for sequential decisions, 1st edn. John Wiley & Sons, Hoboken
- Prud'homme J, Josselin D, Aryal J (2011) Quantitative analysis of pollutant emissions in the context of demand responsive transport. *Computational Science and Its Applications-ICCSA 2011: International Conference. Proceedings, Part I* 11. 439–453. https://doi.org/10.1007/978-3-642-21928-3_31
- Puterman ML (2014) Markov Decision Processes: Discrete Stochastic Dynamic Programming. John Wiley and Sons, New York
- Qiu H, Li R, Zhao J (2018) Dynamic pricing in shared mobility on demand service. Working paper, Massachusetts Institute of Technology.
- Schasché SE, Sposato RG, Hampl N (2022) The dilemma of demand-responsive transport services in rural areas: conflicting expectations and weak user acceptance. *Transp Policy* 126:43–54. <https://doi.org/10.1016/j.tranpol.2022.06.015>
- Schilde M, Doerner KF, Hartl RF (2011) Metaheuristics for the dynamic stochastic dial-a-ride problem with expected return transports. *Comput Oper Res* 38(12):1719–1730. <https://doi.org/10.1016/j.cor.2011.02.006>
- Sharif Azadeh S, Atasoy B, Ben-Akiva ME, Bierlaire M, Maknoon MY (2022) Choice-driven dial-a-ride problem for demand responsive mobility service. *Transp Res Part B: Methodol* 161:128–149. <https://doi.org/10.1016/j.trb.2022.04.008>

- Sieber L, Ruch C, Hörl S, Axhausen KW, Frazzoli E (2020) Improved public transportation in rural areas with self-driving cars: A study on the operation of Swiss train lines. *Transp Res Part a: Policy Practice* 134:35–51. <https://doi.org/10.1016/j.tra.2020.01.020>
- Sörensen L, Bossert A, Jokinen JP, Schlüter J (2021) How much flexibility does rural public transport need?—Implications from a fully flexible DRT system. *Transp Policy* 100:5–20. <https://doi.org/10.1016/j.tranpol.2020.09.005>
- 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. <https://doi.org/10.1016/j.ejor.2018.01.011>
- Talluri K, Van Ryzin G (2004) Revenue management under a general discrete choice model of consumer behavior. *Manag Sci* 50(1):15–33. <https://doi.org/10.1287/mnsc.1030.0147>
- Ulmer MW, Goodson JC, Mattfeld DC, Thomas BW (2020) On modeling stochastic dynamic vehicle routing problems. *EURO J Transp Logist* 9(2):100008. <https://doi.org/10.1016/j.ejtl.2020.100008>
- Vazifteh MM, Santi P, Resta G, Strogatz SH, Ratti C (2018) Addressing the minimum fleet problem in on-demand urban mobility. *Nature* 557(7706):534–538. <https://doi.org/10.1038/s41586-018-0095-1>
- Via (2024) Case studies. Retrieved from <https://ridewithvia.com/resources/category/case-study>. Last accessed: 22/10/2024.
- Viergutz K, Schmidt C (2019) Demand responsive-vs. conventional public transportation: a MATSim study about the rural town of Colditz. *Ger Proced Comput Sci* 151:69–76. <https://doi.org/10.1016/j.procs.2019.04.013>
- Wang C, Quddus M, Enoch M, Ryley T, Davison L (2015) Exploring the propensity to travel by demand responsive transport in the rural area of Lincolnshire in England. *C Stud Transp Policy* 3(2):129–136. <https://doi.org/10.1016/j.cstp.2014.12.006>
- Yang X, Strauss AK, Currie CSM, Eglese R (2016) Choice-based demand management and vehicle routing in e-fulfillment. *Transp Sci* 50(2):473–488. <https://doi.org/10.1287/trsc.2014.0549>
- Zwick F, Kuehnel N, Moeckel R, Axhausen KW (2021) Ride-pooling efficiency in large, medium-sized and small towns-simulation assessment in the munich metropolitan region. *Proced Comput Sci* 184:662–667. <https://doi.org/10.1016/j.procs.2021.03.083>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Fabian Anzenhofer¹  · David Fleckenstein¹ · Robert Klein¹  ·
Claudius Steinhardt² 

✉ Robert Klein
robert.klein@uni-a.de

Fabian Anzenhofer
fabian.anzenhofer@uni-a.de

David Fleckenstein
david.fleckenstein@uni-a.de

Claudius Steinhardt
claudius.steinhardt@unibw.de

¹ Chair of Analytics and Optimization, University of Augsburg, Universitätsstraße 16, 86159 Augsburg, Germany

² Chair of Business Analytics and Management Science, University of the Bundeswehr Munich, Werner-Heisenberg-Weg 39, 85577 Neubiberg, Germany