# **Integrated Demand Management and Vehicle Routing**

# Theory and Applications in Sustainable Planning for Rural Shared Mobility-on-Demand

Kumulative Dissertation an der Wirtschaftswissenschaftlichen Fakultät der Universität Augsburg zur Erlangung des akademischen Grades eines Doktors der Wirtschaftswissenschaften (Dr. rer. pol.)

vorgelegt von

M. Sc. David Fleckenstein

Erstgutachter: Prof. Dr. Robert Klein

Zweitgutachter: Prof. Dr. Axel Tuma

Vorsitzender der mündlichen Prüfung: Prof. Dr. Manuel Ostermeier

Tag der mündlichen Prüfung: 06.03.2025

Für Alessandra

# **List of Articles**

This dissertation comprises the articles listed below. The categories stated for the articles that are published or submitted for publication refer to the journal ranking VHB-Rating 2024 of the Verband der Hochschullehrer für Betriebswirtschaft e. V. (German Academic Association for Business Research). The articles are ordered in accordance with the order of print in the dissertation at hand.

- A1 Published in *European Journal of Operational Research*, 306(2), 499-518 (category A)
  Fleckenstein, D., Klein, R., & Steinhardt, C. (2023). Recent advances in integrating demand management and vehicle routing: A methodological review.
- A2 Published in *Transportation Science*, *online first* (category A)
  Fleckenstein, D., Klein, R., Klein, V., & Steinhardt, C. (2024). On the concept of opportunity cost in integrated demand management and vehicle routing.
- A3 Submitted to *Transportation Science* (category A), 1<sup>st</sup> round Fleckenstein, D., Klein, R., Klein, V., & Steinhardt, C. (2024). From approximation error to optimality gap – Explaining the performance impact of opportunity cost approximation in integrated demand management and vehicle routing. Working paper, University of Augsburg.
- A4 Anzenhofer, F., & Fleckenstein, D. (2024). Extended booking horizons in rural shared mobility-on-demand systems: Insights and implications for demand management.
  Working paper, University of Augsburg.
- A5 To appear in OR Spectrum (category A)
  Anzenhofer, F., Fleckenstein, D., Klein, R., & Steinhardt, C. (2024). Analyzing the impact of demand management in rural shared mobility-on-demand systems.
- A6 Anzenhofer, F., Fleckenstein, D., Klein, R., & Steinhardt, C. (2024). Sustainable dynamic pricing for rural shared mobility-on-demand systems. Working paper, University of Augsburg.

# **Table of Contents**

I Introduction1
II Articles
Article A1: Recent Advances in Integrating Demand Management and Vehicle Routing: A Meth- odological Review
Article A2: On the Concept of Opportunity Cost in Integrated Demand Management and Vehicle
Routing
Article A3: From Approximation Error to Optimality Gap – Explaining the Performance Impact of Opportunity Cost Approximation in Integrated Demand Management and Vehicle Routing 
Article A4: Extended Booking Horizons in Rural Shared Mobility-on-Demand Systems: Insights
and Implications for Demand Management152
Article A5: Analyzing the Impact of Demand Management in Rural Shared Mobility-on-Demand Systems
Article A6: Sustainable Dynamic Pricing for Rural Shared Mobility-on-Demand Systems247
III Conclusion

# I Introduction

With more than 85% of the European population using the mobile internet (GSMA, 2023), practically anyone can order individualized logistical services anytime, anywhere. This proliferation of digital distribution channels has fundamentally changed the market environment in which providers of business-to-consumer logistical services operate. Among the most impactful changes are the following three:

- Breakthrough of new business models: Logistical services are affected by the general trend of service individualization (Lehrer et al., 2018). In the pre-digital era, private customers only had access to highly standardized logistical services. These include both goods transportation services like standard mail delivery or parcel delivery and passenger transportation services such as scheduled public transport with buses and railways. Via digital channels, customers can now easily configure, order, and pay for customized logistical services on short notice. More precisely, customers can choose the desired parameters such as timing, speed, or origin and destination of the transportation service. Before, such parameters were set by the provider alone. Now, these business models, which were previously confined to niche markets, often in the business-to-business segment, can be run on a large-scale for private customers.
- Increase in dynamism and stochasticity: Regarding the operational fulfillment planning problems arising for logistical service providers, the new business models are associated with a shift from static, deterministic problems to dynamic, stochastic problems. The main reasons for that are the individuality of orders and the ability to submit ad-hoc service requests in parallel to ongoing fulfillment operations. This means that planning decisions can no longer be made in batches as, e.g., in standard mail delivery (Irnich, 2008), or at the strategic-tactical planning level as, e.g., in scheduled public transport (Schöbel, 2012). Instead, vehicle routing decisions must be made in real-time for sequentially arriving customer requests. Likewise, stochasticity becomes much more relevant because when making one such decision, the characteristics of future orders, which influence the overall system performance, are highly uncertain. Overall, operational planning for the new logistical services becomes much more complex since the provider must solve dynamic, stochastic vehicle routing problems (Soeffker et al., 2022).
- **Opportunities for demand management**: When customers are given the freedom to choose their desired service parameters, they demand a very high service level, i.e., as little deviation as possible by the provider (Amorim et al., 2024). This poses another challenge for providers since it drastically reduces the opportunities for demand consolidation, which the traditional standardized services naturally allow (Ulmer, 2020). To navigate the trade-off between service level and demand consolidation, providers can apply demand management. The basic idea behind it is to actively control the availability or the pricing of fulfillment options that customers can choose from upon submitting a request. Thereby, customers' choices are influenced according to the providers objectives, and the system performance improves. The

methodological foundation of these approaches lies in traditional revenue management and dynamic pricing, which originated in the airline industry (Strauss et al., 2018).

As a result of these changes, providers in many different fields of application face a common optimization problem at the operational planning level: They must decide on which offers are made to requesting customers and (simultaneously) decide on a feasible route plan for the fulfillment of collected orders. Usually, both decisions are optimized with the objective of maximizing profit, i.e., revenue net of routing cost. The resulting novel optimization problems are called integrated demand management and vehicle routing problems (i-DMVRPs). While dynamic vehicle routing problems as well as revenue management and dynamic pricing problems alone are already highly complex, considering them in an integrated fashion entails additional complexity since both types of decisions are heavily intertwined. Therefore, integrated demand management and vehicle routing has emerged as a distinct research area, and its methodology combines elements from both dynamic vehicle routing as well as revenue management and dynamic pricing.

In the following, an initial understanding of the basic structure of i-DMVRPs is provided. To this end, we first consider a single customer's booking process before widening the view to consider the structure of the entire planning horizon.



Fig. 1 Prototypical booking process

Fig. 1 shows a prototypical, simplified example of a single customers' booking process and the associated planning tasks for the provider. The booking process of a single customer can be subdivided into four steps. The first step (request arrival) and the third step (order confirmation) require the customer to enter data into their smartphone or web application. The second step (demand management) and the fourth step (vehicle routing) require the provider to make planning decisions. We consider all steps in more detail below:

1. **Request arrival**: The customer partly specifies the desired logistical service. The point in time at which the request is submitted is termed *time of request*. Usually, a delivery *location* or drop-off location is required along with the *desired time* of service fulfillment. In some

applications, additional parameters, such as a pick-up location or a more detailed specification of the service type, are relevant. Depending on the relation between time of request and desired time, we can distinguish three types of requests: *advance requests* (time of request at least one day prior to desired time), *same-day requests* (time of request on the same day as desired time), and *ad-hoc requests* (time of request matches desired time). Upon entering the data, the customer requests an offer by the provider, which triggers Step 2.

- 2. Demand management: Having received the request, the provider must decide on the offer they make in response. Optimizing this decision represents the first subproblem, termed demand management subproblem, of the provider's overall operational planning problem. In its most basic form, the demand management decision corresponds to the provider deciding on whether the desired fulfillment option is offered, i.e., whether the request is accepted as it is or rejected. As an extension, the provider can also decide on whether alternative fulfillment options are offered and/or vary the price of each fulfillment option. In any case, the result of the demand management decision, which is termed the *offer set*, is then presented to the customer. This initiates Step 3.
- 3. **Order confirmation**: Following provider-side request acceptance, the customer either confirms their order or abandons the booking process if they are not satisfied with the provider's offer. If the offer set comprises multiple fulfillment options, the customer chooses exactly one as part of the order confirmation. A successful order confirmation triggers Step 4.
- 4. Vehicle routing: In the final step, it is again up to the provider to decide on how to fulfill the collected order, i.e., how the route plan is adapted to include all pending orders and the newly received one. To determine the updated route plan, the provider must solve a second subproblem, namely the vehicle routing subproblem.



#### Fig. 2 Planning process

Now, we consider the entire planning process. Fig. 2 shows a time period of five days. Typically, the provider only offers service during a certain period of a day, which is called a service horizon. We now focus on the three different service horizons depicted in different line styles (on Day 3, Day 4, and Day 5). As the matching line style indicates, each of these service horizons is associated with a corresponding booking horizon during which customers can submit requests for

service at this particular service horizon. The temporal relation between a pair of corresponding booking horizon and service horizon can be twofold:

- Overlapping horizons: In this case, both horizons may cover the exact same time period. Then, customers can only submit same-day requests and ad-hoc requests. Alternatively, as depicted in Fig. 2, the booking horizon additionally extends over multiple days prior to the day of the service horizon, also allowing customers to place advance requests. E.g., a service to be fulfilled during the service horizon of Day 3 (solid line) can be requested from Day 1 until the end of the service horizon on Day 3. Analogously, requests for the service horizon of Day 4 (dashed line) can be submitted on Day 2, Day 3, and Day 4 itself until the respective service horizon ends.
- **Disjoint horizons**: In case of disjoint horizons, the booking horizon also starts multiple days before the day of the service horizon but ends at a cut-off time before the start of the corresponding service horizon. In this setting, customers can only submit advance requests. E.g., requests for a service to be fulfilled during the service horizon of Day 3 (solid line) can only be submitted until the evening of Day 2. Likewise, the latest point in time for submitting a request for Day 4 (dashed line) lies in the evening hours of Day 3.

Dividing the planning horizon into several service horizons has a crucial advantage: Providers can then plan each pair of a service horizon and its corresponding booking horizon independently. Hence, a problem instance of an i-DMVRP is equivalent to one pair of finite horizons. As long as only few customers are willing to substitute between multiple days, this assumption is reasonably weak. If this is not the case, an infinite planning horizon results.

To illustrate the high relevance and broad variety of business models involving i-DMVRPs, we now briefly discuss three of their most prominent applications:

• Attended home delivery (AHD): Unlike standard parcels, certain deliveries require the customer to be present at home to receive them. Typical examples include bulky goods such as large household appliances and furniture, perishable goods like groceries, or pharmaceuticals that must be received personally by the customer by law. In the AHD business model, the delivery service is coupled with the purchase of these goods. Customer requests in AHD are characterized by a shopping basket and a delivery location. The offer set, which the provider presents in response to a request, consists of a menu of delivery time slots. The resulting vehicle routing subproblem is a vehicle routing problem with time windows. In AHD research, booking horizon and service horizon are often assumed to be disjoint. Then, the vehicle routing subproblem is static, and can be solved in the time period between cut-off time and start of the service horizon. In practice, however, more and more providers relax this assumption, resulting in overlapping horizons (Waßmuth et al., 2023).

In 2023, AHD providers in the e-grocery market globally incurred an estimated 490 billion USD revenue (Statista, 2024a). In Germany alone, e-grocery has become a steadily growing multi-billion Euro market, which is projected to grow even further in the coming years (Hofer

et al., 2024). Research on i-DMVRPs in AHD can be traced back to the seminal work by Campbell and Savelsbergh (2005). Regarding demand management, both availability control (e.g., Lang et al., 2021 and Mackert, 2019) and dynamic pricing (e.g., Koch and Klein, 2020 and Yang et al., 2016) are well-investigated. For an overview of this stream of literature, we refer the interested reader to the surveys by Cordeau et al. (2024), Snoeck et al. (2020), and Waßmuth et al. (2023).

- Same-day delivery (SDD): SDD refers to the same-day delivery of parcels, groceries, and meals. While traditional SDD services exclusively transported urgent, high value items in the business-to-business market, the business-to-consumer market has grown rapidly in recent years (Allen et al., 2018). In today's SDD business-to-consumer market, the delivered goods are predominantly food. They can be categorized into recurring groceries, emergency groceries, recipe boxes, and prepared meals (Buldeo Rai et al., 2023). While the request placement step is similar to AHD, the offer set comprises delivery deadlines instead of time windows. There is a clear trend toward shorter deadlines, which has even led to the establishment of own sub-types of SDD. These include instant delivery in less than two hours (Dablanc et al., 2017), and quick commerce in less than twenty minutes (Buldeo Rai et al., 2023). Due to booking horizon and service horizon being identical, both demand management and vehicle routing decisions must be made dynamically. The vehicle routing subproblem can be cast as a multi-trip vehicle routing problem with release and due times (Klein and Steinhardt, 2023). SDD has reached considerable market size. For the year 2023, global (German) revenue in the quick commerce segment is estimated to equal 144 billion USD (810 million EUR) (Statista, 2024c), and estimates for prepared meal delivery are even higher with 394 billion USD (7 billion EUR) (Statista, 2024b). The literature stream on i-DMVRPs in SDD is initiated by Azi et al. (2012) who apply accept/reject demand management. More recent studies also investigate dynamic pricing (Ulmer, 2020, Klein and Steinhardt, 2023). A comprehensive survey of the research in this field can be found in Li et al. (2024).
- Mobility-on-demand (MOD): i-DMVRPs also arise in local passenger transportation. To subsume these business models, the term MOD (e.g., Atasoy et al., 2015), or alternatively, demand-responsive transport (e.g., Brake et al., 2004) is used in the literature. Originating from phone-based dial-a-bus systems in the 1970s, a wide variety of services has emerged (Currie and Fournier, 2020). Enoch et al. (2004) categorize them into the following subtypes: Interchange MOD services operate as feeders for scheduled public transport. Network MOD services, which are potentially destination-specific, complement scheduled public transport, e.g., to link underserved areas or extend public transport into the night. Substitute MOD services fully replace scheduled public transport in an entire region. Another important distinction can be made based on whether the ride is exclusive for the requesting customer as in traditional taxis (e.g., Alonso-Mora et al., 2017) or combined with rides of unrelated customers (shared MOD), which is known as ridepooling (e.g., Zwick et al., 2022). Further, there

are substantial differences between urban services and rural services. However, all business models have in common that customers can request rides from a desired origin to a desired destination at a desired time for one or multiple passengers. Potentially, customers can choose between different products that differ, e.g., in terms of the maximum added ride time. After submitting their request, customers can usually choose between different ride options. Since booking horizon and service horizon typically overlap, the provider must solve a dynamic dial-a-ride problem (Cordeau and Laporte, 2007) or, in case of exclusive rides, a matching problem (Alonso-Mora et al., 2017).

The MOD market size is difficult to estimate due to the plethora of (small) service providers and because most services are run in a business-to-government setting as a form of public transport. According to an estimate by Foljanty (2024), the size of the global business-togovernment market alone was 1.15 billion USD in 2023 with Germany being among the three leading national markets. In the academic literature, Atasoy et al. (2015) are the first to consider an i-DMVRP in the context of MOD. While they apply availability control, there is a rich body of literature on dynamic pricing (e.g., Arian et al., 2022, Qiu et al., 2018, or Sharif Azadeh et al., 2022). Recent surveys of this literature stream are presented by Rammohan et al. (2024), Vansteenwegen et al. (2022), and Zwick et al. (2022).

The subject of the cumulative dissertation at hand is the analysis of the family of i-DMVRPs from a general point of view and the detailed investigation of one specific business model, namely rural shared MOD (SMOD) services. Throughout the dissertation, there is a strong focus on the demand management subproblem, and the applied methodology regarding modeling and solving the considered problems is mainly drawn from operations research.

The remainder of the introduction serves two purposes: On the one hand, it introduces the individual contributions of articles A1-A6. On the other hand, it explicitly highlights the interconnections between the articles. Section 1 provides more detail about the planning problems, models, and solution approaches considered in the academic literature based on Article A1. Sketching the contributions of Article A2 and Article A3, Section 2 outlines how the analytical and numerical analysis of opportunity cost yields domain knowledge about the general structure of i-DMVRPs and facilitates the selection and development of solution algorithms. Finally, Section 3 introduces the i-DMVRP arising in rural SMOD. Further, it presents the main contributions of the remaining three articles: Article A4 identifies patterns in the rural demand structure that can be exploited by demand management. Article A5 analyzes the impact of different availability control approaches on the sustainability of a rural SMOD system. Article A6 develops a multi-objective, sustainable dynamic pricing approach tailored to rural SMOD systems.

# 1 Problem Definitions, Models, and Solution Approaches for i-DMVPRs

In Article A1, we present the first cross-application survey of the i-DMVRP literature. The motivation for conducting such a study is the heterogeneity of research on i-DMVRPs due to authors having diverse backgrounds in terms of modeling techniques (e.g., mixed-integer programming vs. Markov decision process modeling), solution approaches (e.g., sampling-based vs. learningbased), and application areas (e.g., AHD, SDD, or MOD). This heterogeneity is in stark contrast to the common problem structure underlying all i-DMVRPs, which provides ample opportunities for transferring knowledge gained for a specific i-DMVRP to other i-DMVRPs. To leverage these opportunities, a unified, generalized methodological framework is necessary, which Article A1 provides. In the following, the organization of the article is briefly outlined and the relevance of the results for the subsequent articles of this cumulative dissertation is addressed.

In the article, we define three criteria that a problem must meet to be considered an i-DMVRP: First, it must feature a stochastic and dynamic booking process. Second, both types of decisions (demand management and vehicle routing) must be explicitly optimized aiming at some form of profitability (Article A5) or service level maximization, which may be derived from sustainability objectives (Article A6). Third, the provider must have full control over fleet operations. Considering all publications that meet these criteria, we apply morphological analysis to derive a generalized problem definition.

Further, we formulate a high-level Markov decision process (MDP) model that generalizes the many application-specific models for i-DMVRPs and join other researchers (Ulmer et al., 2020) in advocating for MDPs to become the standard modeling approach in i-DMVRP research. This is not only due to their suitability for precisely defining an i-DMVRP mathematically, but also since they allow analytical and numerical analyses of the problem structure (articles A2 and A3). Regarding solution approaches, we provide a structured overview of solution concepts and corresponding algorithms. At the top level, we distinguish static deterministic approximation and decomposition-based approximation, with the latter being by far the more popular solution concept. Its core idea is to further decompose the demand management subproblem into 1) feasibility check, 2) opportunity cost approximation, and 3) determining the actual demand management decision. This observation sets the basis for articles A2 and A3, which aim at guiding the selection and design of opportunity cost approximation algorithms. In articles A5 and A6, the solution concept is also decomposition-based approximation.

Finally, we characterize all relevant publications drawing on the developed framework and identify crucial implications for future research, some of which we take up in articles A2-A6. Among them is the finding that anticipation not only improves the performance regarding the primary objective but can also reduce disparities in the offer quality between different types of customer requests (see also Article A4 and Article A6). Further, we stress the importance of innovation in choice modeling and fulfillment option design, to which we contribute in Article A5 and Article A4, respectively. Another recommendation of Article A1 that the dissertation at hand follows, is transferring i-DVMRP research into practice. E.g., we present an explainability technique that contributes to a better understanding of decomposition-based solution approaches in Article A3. Articles A4, A5, and A6 are based on a collaboration with our industry partner FLEXIBUS and include computational analyses based on a large real-world data set. Finally, we stress the potential of demand management to improve the sustainability of business models, for which articles A5 and A6 provide extensive evidence.

## 2 Definition and Analysis of Opportunity Cost in i-DMVRPs

The decomposition-based solution concept introduced in Section 1 is also widely established in other application areas of demand management, such as traditional revenue management in the airline industry or car rental industry. Since it is well-known that opportunity cost approximation has a decisive impact on the overall performance of the solution approach (Klein et al., 2018), properties of opportunity cost are analyzed with the aim of developing improved solution approaches that exploit them (e.g., Adelman, 2007 or Koch, 2017). However, verifying whether properties such as monotonicity or non-negativity hold in a specific problem is not straightforward. In i-DMVRPs, it is even more challenging because the definition of opportunity cost is fundamentally different compared to traditional revenue management, where opportunity cost is "[...] the expected loss in future revenue from using the capacity now rather than reserving it for future use." (Talluri and Van Ryzin, 2004, p. 33). Given non-negligible and non-attributable variable routing cost, the profit impact of a demand management decision no longer equals displaced revenue but also includes the change in routing cost due to serving the additional order.

Hence, leveraging structural knowledge about opportunity cost in i-DMVRPs requires a rigorous analysis from a formal definition to proving the general validity of mathematical properties and deriving novel types of approximation approaches that exploit them. In Article A2, we provide this analysis and thereby close the corresponding research gap.

In the first step, drawing on a newly introduced model element, the interim state, we can isolate the impact of the demand management decision from the integrated vehicle routing decision. Based on that, we define opportunity cost as the difference in future profit resulting from collecting a potential order compared to not collecting it.

In the second step, we prove that opportunity cost in i-DMVRPs generally has four properties:

- 1. Decomposability into two components (displacement cost and marginal cost-to-serve)
- 2. Potential component-wise negativity
- 3. Overall non-negativity
- 4. State value monotonicity

In the third step, we present three approximation approaches that are designed to exploit the decomposability property and show that they yield promising results in a stylized numerical experiment involving many different problem settings.

Article A3 builds on the theoretical foundation developed in articles A1 and A2. Again, it aims to guide the selection and development of decomposition-based solution approaches for i-DMVRPs. By design, these approaches consist of several algorithmic elements, i.e., feasibility check, opportunity cost approximation, and solution methods for the demand management sub-problem and the vehicle routing subproblem. In Article A1, we point out that research should shed more light on how each of these elements influences the overall solution quality. This question refers to the issue of algorithmic explainability, which has not yet gained much attention in the operations research community (Goerigk and Hartisch, 2023) but is addressed more extensively in reinforcement learning (Milani et al., 2024).

Focusing again on opportunity cost approximation, we develop an explainability technique that consists of two building blocks: The first building block quantifies the chain of influencing factors between a systematic opportunity cost approximation error and the resulting performance loss. The second building block follows the concept of reward decomposition (Juozapaitis et al., 2019) and allows assessing the importance of each of the two opportunity cost components (Article A2) for approximation accuracy.

Applying the technique to the dataset introduced in Article A2, we characterize fundamental types of approximation errors, i.e., errors that are likely to occur in any i-DMVRP. Further, we discuss algorithmic approaches to mitigate them. Referring to the existing academic literature, we show that indications of the identified fundamental types of errors can be found in many computational experiments for specific i-DMVRPs and there seems to be implicit knowledge about them. With our work, we can transform it into explicit knowledge that other researchers and practitioners can take advantage of. Examples of Article A3's practical value can be found in articles A5 and A6. E.g., we find the myopic approaches applied therein are subject to overestimation errors, which can thwart stakeholders' support for dynamic pricing since early requests are systematically charged higher prices (Article A6). Still, the myopic approaches perform remarkably well in terms of objective value, which can also be better explained by the findings in Article A3.

# 3 Sustainability-oriented Demand Management for Rural SMOD Services

Following the generic, rather theoretical analyses of i-DMVRPs in articles A1-A3, the remainder of the dissertation at hand (articles A4-A6) focuses on rural shared mobility-on-demand (SMOD) as one specific application area. The particular relevance of rural SMOD stems from their ability to break the vicious cycle that scheduled public transport suffers from (Bar-Yosef et al., 2013). Due to low, dispersed demand, scheduled public transport only reaches a low service quality, e.g., in terms of service frequency. This leads to a low model split and under-utilized services being

cut back further. SMOD services, in contrast, offer favorable characteristics for providing areawide mobility coverage in low-demand areas (Mounce et al., 2020). This is because they combine demand orientation with (low-degree) demand consolidation (Fig. 3). As these theoretical considerations suggest, replacing or complementing scheduled services by SMOD services has been found to be beneficial by several studies (e.g., Mortazavi et al., 2024, Sieber et al., 2020, Viergutz and Schmidt, 2019).



Fig. 3 Key characteristics of SMOD services

Rural SMOD services differ from most i-DMVRP applications in four characteristics:

- Sustainability as the all-inclusive objective: In AHD, SDD, and many urban (S)MOD systems, providers are purely private companies whose primary objective is to maximize profit. In rural SMOD, municipal authorities either contract the company providing service or even own it (Lu et al., 2024). Therefore, they are deeply involved in planning service provision (Wang et al., 2015). The underlying root cause is that public transport in rural areas is financially unprofitable, and governments must step in to ensure basic mobility provision (Mounce et al., 2020). In the past, their decision-making was also largely economically driven, which limited the differences to purely privately operated services. With increasing efforts to combat climate change, however, the paradigm has shifted toward sustainability (Poltimäe et al., 2022), which is defined by three pillars: social sustainability, environmental sustainability, and economic sustainability (Purvis et al., 2019).
- **Regulated pricing**: Rural SMOD systems are part of the public transport system. Thus, their pricing must be closely aligned with scheduled public transport, which is often enforced by regulation (Schasché et al., 2022). This means that prices are predominantly static, even though dynamic pricing is conceivable within suitable regulation (VDV, 2023). This is different from other i-DMVRP applications, where providers are in full control of pricing.
- Long booking horizon overlapping with service horizon: To allow for as much planning reliability as possible, given the limited alternative modes of transport, rural SMOD providers typically extend the booking horizon well before the start of the service horizon. E.g., our

industry partner FLEXIBUS allows requesting service up to two weeks in advance, while other providers even extend the booking horizon to cover a full month before the start of service horizon (Chandakas, 2020). At the same time, customers should also be able to place requests during the service day since not all trips can be planned in advance. In contrast to AHD, where only advance requests are possible, and SDD and urban (S)MOD, where the booking horizon is much shorter, this results in a special type of i-DMVRP with a long booking horizon that overlaps with the service horizon.

• Low demand volume: Compared to urban SMOD providers, rural providers face much lower demand, and thus, operate a much smaller fleet. The difference can amount to several orders of magnitude. E.g., the urban provider MOIA transports about 7000 passengers per day with 280 vehicles in Hamburg (MOIA, 2024), while our industry partner FLEXIBUS transports only about 70 passengers per day with 2-3 vehicles in the service area Krumbach. For operational planning, this has two implications: On the one hand, ridepooling is particularly challenging since compatible orders are scarce. On the other hand, the instance size of the i-DMVRP is small, which means that more accurate solution approaches can be applied.

In view of these decisive differences compared to other application areas and their high societal significance, the i-DMVRPs resulting from rural SMOD services deserve thorough investigation. In particular, given the profit improvements achieved in other application areas, analyzing how demand management can improve the performance regarding sustainability objectives is of high practical relevance. As the literature reviews in articles A4-A6 reveal, much of the existing research focuses on urban SMOD services, which makes the research gap even larger. Each of the articles A4-A6 investigates the application of demand management to rural SMOD services with a different focus. The remainder of this section introduces the contribution of each article.

#### Article A4

Article A4 descriptively analyzes demand data with the aim of discovering structural patterns that have implications for demand management. Thereby, it provides essential groundwork for articles A5 and A6, which develop demand management approaches tailored to rural SMOD services. Based on the data set, we compile empirical evidence for demand patterns that arise as a direct consequence of the extended booking horizon, which is one of the unique characteristics of rural SMOD services. For each of the observed patterns, we discuss its implications for operational planning, which may be challenges or opportunities. Further, we elaborate criteria that suitable demand management approaches should meet to account for these patterns and improve the system performance. In more detail, we find the following four patterns:

• **Reservation behavior**: First, customers strategically reserve certain types of rides early in the booking horizon. While this means that the provider gives up some control over the booking process, it may even be desirable for providers and customers that such reservations are possible. If the provider wants to curtail them, we suggest applying anticipatory demand management. The approaches developed in article A5 and A6 meet this criterium.

- **Round trips**: Second, customers commonly request round trips in the following way: They request the outward ride in advance and the return ride ad-hoc. While this is advantageous for the provider since they can theoretically reject the return ride, it reduces the planning reliability for customers, which may lead to lost demand. Operationally, this issue can also be addressed by anticipatory demand management, as in articles A5 and A6. Strategically, providers may design a dedicated product for round trips with a larger time window for the return ride.
- **Cancellations**: Third, the long booking horizon leads to orders that are booked early and cancelled late, which we call ghost demand. It is problematic since it distorts the information basis for demand management decisions and blocks fleet resources. We find that cancellation probabilities can be accurately predicted by supervised learning models, which allows the explicit consideration of cancellations in demand management decision-making. Since neither we nor other authors have investigated this topic further, it is still open for future research.
- **Time flexibility**: Fourth, we find that customers have time flexibility in both directions, i.e., the provider has the opportunity to shift their pick-up time earlier or later by offering alternative times to the desired time. This suggests offering multiple alternative rides instead of purely accepting/rejecting the desired ride to improve performance. The demand management approaches presented in articles A5 and A6 follow this finding.

#### Article A5

In Article A5, we consider a strategic decision problem that arises when newly establishing a rural SMOD system or when a switch from first-come-first-served decision-making to active demand management is planned. In this situation, providers and municipal authorities face the strategic question which type of demand management to apply. To provide decision support, we propose a methodology based on a suitable model of the corresponding i-DMVRP. The model is solved with practical solution algorithms representing different classes of demand control policies. These classes differ in terms of the applied mechanisms (rejections or time shifts), criteria for decision-making (feasibility or profitability), and use of information (myopic or anticipatory). Based on a computational study with real-world data, we analyze how the proposed classes of demand control policies differ regarding their performance. For this study, we limit the scope to availability control. Further, we assume that the provider maximizes profit, which yields insights into whether social sustainability and environmental sustainability improve "automatically" when the provider follows the traditional economic planning paradigm.

At the heart of our proposed methodology is a semi-perfect information model, which is derived from an MDP model of the i-DMVRP that the provider faces at the operational planning level. Since the semi-perfect information model is intended to be applied at the strategic planning level, the realization of the i-DMVRP's stochastic information must be accurately simulated. Regarding the request arrivals, this is straightforward because the provider can track historical request arrivals and use them as a basis for the simulation. However, the customer choice behavior in the order confirmation step can hardly be statistically estimated. To avoid a distortion of the results by an inaccurate customer choice model, we propose to model it deterministically and to perform a systematic sensitivity analysis to account for the uncertainty about true customer choice behavior. The studied solution algorithms follow the decomposition-based solution concept introduced in Section 1. For anticipatory demand management, we propose a sampling-based look-ahead algorithm that is transferred from algorithms for related AHD problems (Koch and Klein, 2020, Köhler et al., 2024, and Yang et al., 2016).

In the computational study of Article A5, we find that availability control substantially improves performance in terms of economic sustainability and also environmental sustainability. However, this comes at the cost of social sustainability. Anticipating the future evolution of the booking process brings further, albeit much smaller benefits. Sensitivity analyses reveal another conflict of objectives between economical sustainability and environmental sustainability. The observed benefits of demand management together with the observed conflicts of objectives motivate the development of demand management approaches that explicitly consider and balance multiple sustainability objectives. One such approach is proposed in Article A6.

#### Article A6

Article A6 is motivated by one of the main findings in Article A5: Purely setting economic incentives for the provider to achieve sustainability improvements potentially causes an underrepresentation of the other two pillars of sustainability. Similar difficulties can be observed if the authority imposes certain minimum performance requirements for individual sustainability objectives (Anzenhofer et al., 2025). In Article A6, we investigate the alternative idea of enforcing the application of an explicitly sustainability-oriented demand management approach. Next to a potentially more balanced performance in terms of sustainability, such an approach also enables the application of dynamic pricing. Assuming a profit-maximizing provider as in Article A5 and the vast majority of the academic literature, unregulated dynamic pricing would lead to monopoly markups, i.e., a general increase of the price level solely to exploit customers' willingness-to-pay (Hörcher and Graham, 2020).

The sustainable dynamic pricing approach we propose in Article A6 is based on a thorough investigation of the relevant objectives by means of multi-attribute decision analysis (Keeney and Raiffa, 1993). Therewith, we identify five objectives:

- **Social sustainability**: From the social perspective, 1) the maximization of basic mobility provision and 2) the avoidance of monopoly markups are relevant.
- Environmental sustainability: In terms of environmental sustainability, the provider should 3) aim at a maximization of modal shift from motorized individual transport and 4) aim at the minimization of emissions per passenger km of the SMOD service.
- Economic sustainability: Finally, economic sustainability can be cast as equal to 5) minimizing subsidy requirements.

Based on the structuring of objectives, we formulate a constrained MDP model, which builds on the model presented in Article A5. The objective function only accounts for the primary objective of maximizing served demand. Thereby, we incorporate the social objective of maximizing basic mobility provision and the environmental objective of maximizing modal shift. To consider the remaining secondary objectives, we include a constraint that imposes a lower bound on the price of each ride. This bound is set equal to the marginal cost of the ride, i.e., variable routing cost and external cost resulting from emissions. Thereby, we implement the well-established concept of marginal cost pricing (Hörcher and Tirachini, 2021). Based on the price signal, it allows the customer to decide whether a ride is worth its societal cost, thereby steering demand (Eliasson, 2021). Making accurate dynamic pricing decisions requires an evaluation of rides, similar to opportunity cost approximation in profit-based i-DMVRPs, to approximate displaced demand and marginal cost. To this end, we also develop a post-decision rollout algorithm (Bertsekas et al., 1997) in Article A6. It shares many similarities with the sampling-based look-ahead algorithm presented in Article A5 but crucially allows for an assessment of demand displacement and the impact of future demand management decisions.

The computational study based on real-world data shows that sustainable dynamic pricing is advantageous both for the provider and for customers compared to static pricing, which is the statusquo in practice, and compared to profit-oriented dynamic pricing, which is almost exclusively proposed in the academic literature. Again, we find that anticipation is beneficial, not only in terms of sustainability but also since it allows for a more constant offer quality across the booking horizon due to the avoidance of opportunity cost overestimation errors (Article A3). Based on the results, we also discuss the application of our dynamic pricing approach in practice. In particular, we recommend an accompanying communication strategy that emphasizes the advantages for individual customers.

# References

- Adelman, D. (2007). Dynamic bid prices in revenue management. *Operations Research*, 55(4), 647-661. https://doi.org/10.1287/opre.1060.0368.
- Allen, J., Piecyk, M., & Piotrowska, M. (2018). An analysis of the same-day delivery market and operations in the UK. Retrieved from http://www.ftc2050.com/reports/ftc2050-2018-Sameday\_delivery\_market\_UK.pdf. Last accessed: 02/12/2024.
- Alonso-Mora, J., Samaranayake, S., Wallar, A., Frazzoli, E., & Rus, D. (2017). On-demand highcapacity ride-sharing via dynamic trip-vehicle assignment. *Proceedings of the National Academy of Sciences of the United States of America*, 114(3), 462-467. https://doi.org/10.1073/pnas.1611675114.
- Amorim, P., DeHoratius, N., Eng-Larsson, F., & Martins, S. (2024). Customer preferences for delivery service attributes in attended home delivery. *Management Science*, 70(11), 7559-7578. https://doi.org/10.1287/mnsc.2020.01274.
- Anzenhofer, F., Schmidbaur, S., Klein, R., & Steinhardt, C. (2025). The potential of governmental regulation on shared mobility-on-demand systems. *Transportation Research Part A: Policy* and Practice, 192, 104360. https://doi.org/10.1016/j.tra.2024.104360.
- 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.
- Atasoy, B., Ikeda, T., Song, X., & Ben-Akiva, M. E. (2015). The concept and impact analysis of a flexible mobility on demand system. *Transportation Research Part C: Emerging Technol*ogies, 56, 373-392. https://doi.org/10.1016/j.trc.2015.04.009.
- Azi, N., Gendreau, M., & Potvin, J.-Y. (2012). A dynamic vehicle routing problem with multiple delivery routes. *Annals of Operations Research*, 199(1), 103-112. https://doi.org/10.1007/s10479-011-0991-3.
- Bar-Yosef, A., Martens, K., & Benenson, I. (2013). A model of the vicious cycle of a bus line. *Transportation Research Part B: Methodological*, 54, 37-50. https://doi.org/10.1016/j.trb.2013.03.010.
- Bertsekas, D. P., Tsitsiklis, J. N., & Wu, C. (1997). Rollout algorithms for combinatorial optimization. *Journal of Heuristics*, *3*, 245-262. https://doi.org/10.1023/A:1009635226865.
- Brake, J., Nelson, J. D., & Wright, S. (2004). Demand responsive transport: Towards the emergence of a new market segment. *Journal of Transport Geography*, 12(4), 323-337. https://doi.org/10.1016/j.jtrangeo.2004.08.011.
- Buldeo Rai, H., Mariquivoi, J., Schorung, M., & Dablanc, L. (2023). Dark stores in the City of Light: Geographical and transportation impacts of 'quick commerce' in Paris. *Research in Transportation Economics*, 100, 101333. https://doi.org/10.1016/j.retrec.2023.101333.
- Campbell, A. M., & Savelsbergh, M. W. P. (2005). Decision support for consumer direct grocery initiatives. *Transportation Science*, *39*(3), 313-327. https://doi.org/10.1287/trsc.1040.0105.
- Chandakas, E. (2020). On demand forecasting of demand-responsive paratransit services with prior reservations. *Transportation Research Part C: Emerging Technologies*, *120*, 102817. https://doi.org/10.1016/j.trc.2020.102817.
- Cordeau, J. F., Iori, M., & Vezzali, D. (2024). An updated survey of attended home delivery and service problems with a focus on applications. *Annals of Operations Research*, *343*, 885-922. https://doi.org/10.1007/s10479-024-06241-9.
- Cordeau, J. F., & Laporte, G. (2007). The dial-a-ride problem: models and algorithms. *Annals of Operations Research*, 153, 29-46. https://doi.org/10.1007/s10479-007-0170-8.
- Currie, G., & Fournier, N. (2020). Why most DRT/Micro-Transits fail What the survivors tell us about progress. *Research in Transportation Economics*, *83*, 100895. https://doi.org/10.1016/j.retrec.2020.100895.

- Dablanc, L., Morganti, E., Arvidsson, N., Woxenius, J., Browne, M., & Saidi, N. (2017). The rise of on-demand 'Instant Deliveries' in European cities. *Supply Chain Forum: An International Journal*, 18(4), 203-217. https://doi.org/10.1080/16258312.2017.1375375.
- Eliasson, J. (2021). Efficient transport pricing Why, what, and when?. *Communications in Transportation Research*, *1*, 100006. https://doi.org/10.1016/j.commtr.2021.100006.
- Enoch, M., Potter, S., Parkhurst, G., & Smith, M. (2004). Intermode: Innovations in demand responsive transport. Working paper, Loughborough University.
- Foljanty (2024). On-demand transit market report 2023 Recap. Retrieved from https://www.linkedin.com/pulse/on-demand-transit-market-report-2023-recap-lukasfoljanty-ngyzc. Last accessed: 22/12/2024.
- Goerigk, M., & Hartisch, M. (2023). A framework for inherently interpretable optimization models. *European Journal of Operational Research*, 310(3), 1312-1324. https://doi.org/10.1016/j.ejor.2023.04.013.
- GSMA (2023). The mobile economy Europe 2023. Retrieved from https://www.gsma.com/solutions-and-impact/connectivity-for-good/mobile-economy/europe/. Last accessed: 02/12/2024.
- Hörcher, D., & Graham, D. (2020). Pricing and efficient public transport supply in a mobility as a service context. Working paper, Imperial College London.
- Hörcher, D., & Tirachini, A. (2021). A review of public transport economics. *Economics of Transportation*, 25, 100196. https://doi.org/10.1016/j.ecotra.2021.100196.
- Hofer, H., Eikelmann, S., Langkammer, J., Bartels, L., & Banken, J. (2024). Online grocery delivery in flux - An update on the state of eGrocery. Retrieved from https://www.strategyand.pwc.com/de/en/industries/consumer-markets/the-state-of-the-egrocery-market.html. Last accessed: 02/12/2024.
- Irnich, S. (2008). Solution of real-world postman problems. European Journal of Operational Research, 190(1), 52-67. https://doi.org/10.1016/j.ejor.2007.06.002.
- Juozapaitis, Z., Koul, A., Fern, A., Erwig, M., & Doshi-Velez, F. (2019). Explainable reinforcement learning via reward decomposition. In *IJCAI/ECAI Workshop on Explainable Artificial Intelligence*.
- Keeney, R. L., & Raiffa, H. (1993). Decisions with multiple objectives: Preferences and value tradeoffs. Cambridge: Cambridge University Press.
- Klein, R., Mackert, J., Neugebauer, M., & Steinhardt, C. (2018). A model-based approximation of opportunity cost for dynamic pricing in attended home delivery. *OR Spectrum*, 40(4), 969-996. https://doi.org/10.1007/s00291-017-0501-3.
- Klein, V., & Steinhardt, C. (2023). Dynamic demand management and online tour planning for same-day delivery. *European Journal of Operational Research*, 307(2), 860-886. https://doi.org/10.1016/j.ejor.2022.09.011.
- Koch, S. (2017). Least squares approximate policy iteration for learning bid prices in choicebased revenue management. *Computers & Operations Research*, 77, 240-253. https://doi.org/10.1016/j.cor.2016.07.015.
- Koch, S., & Klein, R. (2020). Route-based approximate dynamic programming for dynamic pricing in attended home delivery. *European Journal of Operational Research*, 287(2), 633-652. https://doi.org/10.1016/j.ejor.2020.04.002.
- Köhler, C., Campbell, A. M., & Ehmke, J. F. (2024). Data-driven customer acceptance for attended home delivery. OR Spectrum, 46(2), 295-330. https://doi.org/10.1007/s00291-023-00712-4.
- Lang, M. A. K., Cleophas, C., & Ehmke, J. F. (2021). Anticipative dynamic slotting for attended home deliveries. *Operations Research Forum*, 2(4), 70. https://doi.org/10.1007/s43069-021-00086-9.
- Lehrer, C., Wieneke, A., Vom Brocke, J. A. N., Jung, R., & Seidel, S. (2018). How big data analytics enables service innovation: materiality, affordance, and the individualization of

service. Journal of Management Information Systems, 35(2), 424-460. https://doi.org/10.1080/07421222.2018.1451953.

- Li, Y., Archetti, C., & Ljubic, I. (2024). Emerging optimization problems for distribution in sameday delivery. Working paper, ESSEC Business School. https://doi.org/10.48550/arXiv.2405.05620.
- Lu, C., Tiwari, S., Nassir, N., & Nagel, K. (2024). Efficient operation of demand-responsive transport (DRT) systems: Active requests rejection. *Procedia Computer Science*, 238, 81-90. https://doi.org/10.1016/j.procs.2024.06.001.
- Mackert, J. (2019). Choice-based dynamic time slot management in attended home delivery. *Computers* & *Industrial Engineering*, *129*, 333-345. https://doi.org/10.1016/j.cie.2019.01.048.
- Milani, S., Topin, N., Veloso, M., & Fang, F. (2024). Explainable reinforcement learning: A survey and comparative review. ACM Computing Surveys, 56(7), 1-36. https://doi.org/10.1145/3616864.
- MOIA (2024). MOIA auf einen Blick. Retrieved from https://www.moia.io/de-DE/news-center. Last accessed: 02/12/2024.
- Mortazavi, A., Ghasri, M., & Ray, T. (2024). Integrated demand responsive transport in lowdemand areas: A case study of Canberra, Australia. *Transportation Research Part D: Transport and Environment*, 127, 104036. https://doi.org/10.1016/j.trd.2023.104036.
- Mounce, R., Beecroft, M., & Nelson, J. D. (2020). On the role of frameworks and smart mobility in addressing the rural mobility problem. *Research in Transportation Economics*, *83*, 100956. https://doi.org/10.1016/j.retrec.2020.100956.
- Poltimäe, H., Rehema, M., Raun, J., & Poom, A. (2022). In search of sustainable and inclusive mobility solutions for rural areas. *European Transport Research Review*, 14(1), 13. https://doi.org/10.1186/s12544-022-00536-3.
- Purvis, B., Mao, Y., & Robinson, D. (2019). Three pillars of sustainability: In search of conceptual origins. *Sustainability Science*, 14, 681-695. https://doi.org/10.1007/s11625-018-0627-5.
- Qiu, H., Li, R., & Zhao, J. (2018). Dynamic pricing in shared mobility on demand service. Working paper, Massachusetts Institute of Technology.
- Rammohan, S., Marathe, R. R., & Sudarsanam, N. (2024). Recent advancements in revenue management of taxi services: A systematic review and research agenda. *Management Review Quarterly*, 74(2), 1029-1055. https://doi.org/10.1007/s11301-023-00326-8.
- Schasché, S. E., Sposato, R. G., & Hampl, N. (2022). The dilemma of demand-responsive transport services in rural areas: Conflicting expectations and weak user acceptance. *Transport Policy*, 126, 43-54. https://doi.org/10.1016/j.tranpol.2022.06.015.
- Schöbel, A. (2012). Line planning in public transportation: models and methods. *OR Spectrum*, 34(3), 491-510. https://doi.org/10.1007/s00291-011-0251-6.
- Sharif Azadeh, S., Atasoy, B., Ben-Akiva, M. E., Bierlaire, M., & Maknoon, M. Y. (2022). Choice-driven dial-a-ride problem for demand responsive mobility service. *Transportation Research Part B: Methodological*, 161, 128-149. https://doi.org/10.1016/j.trb.2022.04.008.
- Sieber, L., Ruch, C., Hörl, S., Axhausen, K. W., & Frazzoli, E. (2020). Improved public transportation in rural areas with self-driving cars: A study on the operation of Swiss train lines. *Transportation Research Part A: Policy and Practice*, 134, 35-51. https://doi.org/10.1016/j.tra.2020.01.020.
- Snoeck, A., Merchán, D., & Winkenbach, M. (2020). Revenue management in last-mile delivery: State-of-the-art and future research directions. *Transportation Research Procedia*, 46, 109– 116. https://doi.org/10.1016/j.trpro.2020.03.170.
- Soeffker, N., Ulmer, M. W., & Mattfeld, D. C. (2022). Stochastic dynamic vehicle routing in the light of prescriptive analytics: A review. *European Journal of Operational Research*, 298(3), 801-820. https://doi.org/10.1016/j.ejor.2021.07.014.

- Strauss, A. K., Klein, R., & Steinhardt, C. (2018). A review of choice-based revenue management: Theory and methods. *European Journal of Operational Research*, 271(2), 375-387. https://doi.org/10.1016/j.ejor.2018.01.011.
- Statista (2024a). Grocery delivery. Retrieved from https://de.statista.com/outlook/dmo/online-food-delivery/grocery-delivery/retail-delivery/. Last accessed: 02/12/2024.
- Statista (2024b). Meal delivery. Retrieved from https://de.statista.com/outlook/dmo/online-food-delivery/meal-delivery/. Last accessed: 02/12/2024.
- Statista (2024c). Quick commerce. Retrieved from https://de.statista.com/outlook/dmo/online-food-delivery/grocery-delivery/quick-commerce/. Last accessed: 02/12/2024.
- Talluri, K. T., & Van Ryzin, G. J. (2004). The theory and practice of revenue management. New York, NY: Springer. https://doi.org/10.1007/b139000.
- Ulmer, M. W. (2020). Dynamic pricing and routing for same-day delivery. *Transportation Science*, 54(4), 1016-1033. https://doi.org/10.1287/trsc.2019.0958.
- Ulmer, M. W., Goodson, J. C., Mattfeld, D. C., & Thomas, B. W. (2020). On modeling stochastic dynamic vehicle routing problems. *EURO Journal on Transportation and Logistics*, 9(2), 100008. https://doi.org/10.1016/j.ejtl.2020.100008.
- Vansteenwegen, P., Melis, L., Aktaş, D., Montenegro, B. D. G., Vieira, F. S., & Sörensen, K. (2022). A survey on demand-responsive public bus systems. *Transportation Research Part C: Emerging Technologies*, 137, 103573. https://doi.org/10.1016/j.trc.2022.103573.
- VDV (2023). Linienbedarfsverkehr: zukunftsgerecht, integriert und nachfragegesteuert. Retrieved from https://www.vdv.de/positionspapier-linienbedarfsverkehr-maerz23.pdfx. Last accessed: 15/11/2024.
- Viergutz, K., & Schmidt, C. (2019). Demand responsive-vs. conventional public transportation: A MATSim study about the rural town of Colditz, Germany. *Procedia Computer Science*, 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. *Case Studies on Transport Policy*, 3(2), 129-136. https://doi.org/10.1016/j.cstp.2014.12.006.
- Waßmuth, K., Köhler, C., Agatz, N., & Fleischmann, M. (2023). Demand management for attended home delivery – A literature review. *European Journal of Operational Research*, 311(3), 801-815. https://doi.org/10.1016/j.ejor.2023.01.056.
- Yang, X., Strauss, A. K., Currie, C. S. M., & Eglese, R. (2016). Choice-based demand management and vehicle routing in e-fulfillment. *Transportation Science*, 50(2), 473-488. https://doi.org/10.1287/trsc.2014.0549.
- Zwick, F., Kuehnel, N., & Axhausen, K. W. (2022). Review on theoretical assessments and practical implementations of ride-pooling. In 22nd Swiss Transport Research Conference (STRC 2022). https://doi.org/10.3929/ethz-b-000548675.

II Articles

# Article A1: Recent Advances in Integrating Demand Management and Vehicle Routing: A Methodological Review

David Fleckenstein<sup>a</sup>, Robert Klein<sup>a</sup>, Claudius Steinhardt<sup>b</sup>

(<sup>a</sup>University of Augsburg, <sup>b</sup>University of the Bundeswehr Munich)

Published in a category A journal: European Journal of Operational Research, 306(2), 499-518 (2023). https://doi.org/10.1016/j.ejor.2022.04.032.

#### Abstract

In logistics and mobility services, new business models such as "attended home delivery", "sameday delivery", and "mobility-on-demand" have been successfully established over the last decade. They have in common that customers order online, while the services are provided offline. To make such online-to-offline services profitable, the efficient operation of a vehicle fleet is an essential prerequisite. Therefore, researchers began to explore approaches for integrating demand management and vehicle routing to support such operations, and a rapidly growing body of literature emerged. However, due to the diversity of existing business models, the analysis and comparison of decision problems and solution concepts are challenging, especially across applications, making the search for appropriate models and algorithms for new problem settings nontrivial.

Therefore, in this survey, we structure this innovative research area and review the existing literature from a methodological perspective. We present a generalized problem definition of integrated demand management and vehicle routing, derive a high-level formulation for the underlying sequential decision process, and present a corresponding mathematical model. We then describe and characterize solution concepts and algorithms from the literature in a structured way. We also present a tabular overview of the literature that connects applications and problem characteristics with solution concepts and allows researchers to quickly step through already studied combinations. Finally, we comment on the state-of-the-art from a cross-application perspective and discuss future research opportunities.

Key words: Routing, Demand Management, Attended Home Delivery, Same-Day Delivery, Mobility-on-Demand

# 1 Introduction

Over the last decade, many new applications for vehicle routing models and corresponding solution methods have emerged, which have attracted great interest in the research community and in public. Starting points for this development were the introduction of new technologies like drones and delivery robots (Boysen et al., 2021) and the establishment of new business models such as attended home delivery, same-day delivery, and mobility-on-demand (e.g., Agatz et al., 2013, Voccia et al., 2019, and Qin et al., 2020). These business models, often characterized by the term "online-to-offline", allow a service to be booked online that is delivered offline by operating vehicles. Today, with services like Instacart, Amazon PrimeNow, and Uber being commonplace, corresponding business models represent such an essential part of the modern on-demand lifestyle that popular news media like the BBC have covered even the underlying mathematics (Church, 2019).

In this context, demand management has become a popular, often necessary tool. Requests for online-to-offline services arrive over time, and customers have different preferences concerning different fulfillment options. Hence, providers can shape demand, i.e., the set of resulting orders and their characteristics, by offering targeted fulfillment options to specific customers to allow efficient routing operations. A variety of approaches were proposed for this purpose: In the case of attended home delivery (AHD) and field service operations (FSO), the variation of prices or time window availability is often in the focus of demand control (e.g., Strauss et al., 2021 or Avraham and Raviv, 2021). For same-day delivery (SDD) and mobility-on-demand (MOD) services, accepting or rejecting customer requests may be the approach of choice (e.g., Klapp et al., 2020 or Fielbaum et al., 2022). In general, actively controlling demand entails the following benefits for providers: First, control decisions balance demand in temporal and geographical terms to avoid spilled demand on the one hand and low utilization of fulfillment resources on the other hand. This increases the number of orders served by a given fleet and, hence, the overall profit. Second, for time periods or areas where such smoothing does not eliminate capacity shortage, demand control enables allocating available capacity to the most profitable customers (Agatz et al., 2013) and possibly earning additional revenues in the form of delivery fees. Thereby, the average profit per order increases. Third, effective demand control stimulates demand and opens new markets in the form of initially low-demand and, therefore, unprofitable delivery areas (Yang and Strauss, 2017). Fourth, demand control contributes to increasing routing efficiency (Klein et al., 2019). By controlling the fulfillment options sold, service providers can "generate" a favorable instance of the resulting routing problem.

In principle, many established approaches from the field of revenue management, like availability control and dynamic pricing, can be used for demand management purposes (see Strauss et al., 2018 and Klein et al., 2020 for recent surveys). Unfortunately, the integration of demand management and vehicle routing turns out to be quite complex. More precisely, demand is stochastic

and realizes over time, which leads to a sequential decision problem. Providers must decide on fulfillment options for incoming requests without exactly knowing the number of future customers and their preferences. Depending on the orders made, different vehicle routing costs may result, and future revenues may even be displaced, e.g., if an accepted request prevents future orders due to capacity or service constraints. Anticipating these intertemporal effects requires solving vehicle routing problems, which, in general, are NP-hard. Furthermore, to meet customers' expectations, providers must make decisions in real-time (e.g., Poggi et al., 2014).

This complexity led to various new approaches to integrate demand management and vehicle routing, with the center of the respective contributions often depending on the authors' methodological backgrounds (e.g., integer programming or stochastic dynamic programming). However, analyzing the literature shows that the structure of the specific control problems considered is very similar. This observation even holds across application areas. As a consequence, demand management approaches, solution concepts, and algorithms applied in different areas are strongly related. Despite that, the relationships are usually not discussed beyond the areas' borders.

Motivated by these observations, the key contributions of this survey paper are as follows:

- 1) To foster a structured comparison of different real-world applications, we present a generalized definition of integrated demand management and vehicle routing problems. To analyze the characteristics of specific decision problems, we identify four components of the underlying sequential decision process: request capture, demand management, order confirmation, and vehicle routing. Using morphological analysis, we characterize each component regarding several dimensions. We summarize this analysis in a comprehensive morphological box and illustrate the results by describing possible realizations for existing applications in AHD, FSO, SDD, and MOD.
- 2) As a synthesis of specific modeling approaches existing in the literature, we formulate a high-level mathematical model of the generalized sequential decision problem. As tractable solution concepts for decision problems falling under this generalized formulation, we discuss static deterministic approximations as well as decomposition-based approximations. In particular, we investigate the tasks resulting from decomposition-based approximations, i.e., feasibility check, cost estimation, demand control, and routing control, and present corresponding solution approaches often based on specific auxiliary models.
- 3) We present an overview of the literature "at a glance" in two comprehensive tables, linking decision problems and solution concepts to applications. These tables allow researchers to check for suitable approaches without analyzing all possible related fields when they want to apply demand management in their area of interest. Furthermore, they can quickly verify whether certain combinations of specific decision problems and solution concepts have already been examined.
- 4) Complementary to the high-level overview of solution concepts, we discuss selected contributions to algorithms used as part of solution approaches for static deterministic

approximations and decomposition-based approximations in more detail. For the latter class, we highlight the algorithms that are suitable for addressing several tasks in combination.

5) Finally, we identify seven different topics around which we discuss the current state of research to deliver cross-application insights, and which represent fruitful starting points for future research.

The scope and the purpose of our work substantially differ from existing surveys. Agatz et al. (2013) focus more on optimizing demand management decisions and less on the associated routing problems. Besides this, they exclusively consider AHD problems. The latter also holds for the survey by Snoeck et al. (2020), who extensively outline possible extensions of AHD-specific problem settings and their implications. Yan et al. (2020) exclusively deal with matching and dynamic pricing in MOD. The recent survey by Soeffker et al. (2022) considers dynamic vehicle routing in general, with SDD being one of many application areas.

To allow for the necessary focus, we establish the following criteria for selecting the publications for this survey: First, we only include works investigating stochastic and dynamic booking processes. Second, we only consider settings where fulfillment operations must be optimized explicitly by integrating demand management and vehicle routing methods based on profitability or service quality. Hence, we exclude dynamic vehicle routing settings, where providers control service availability purely for ensuring the feasibility of routes and refer the interested reader to surveys by, e.g., Pillac et al. (2013), Psaraftis et al. (2016), and Ulmer et al. (2020). Finally, we assume full information and control regarding the resources needed to fulfill services. Consequently, we do not cover problems involving stochastic vehicle availability or platform-based service provision based on two-sided markets, which arise in the context of sharing-based or crowdsourced fulfillment systems (e.g., Afèche et al., 2023, Banerjee et al., 2016, and Taylor, 2018). Furthermore, we leave out special cases for readability.

Our survey is structured as follows: In Section 2, we first state the problem of integrating demand management and vehicle routing along a generic process formulation. Subsequently, we discuss the characteristics of this process for several areas of application. We then provide an exact, high-level mathematical model formulation for the resulting sequential decision problem in Section 3. In Section 4, we analyze different solution concepts based on tractable approximations of the exact model from Section 3. Section 4 concludes with a summary of all results up to this point in the form of comprehensive tables of the existing literature. Section 5 comprises a more detailed discussion on solution algorithms and may be skipped by readers only looking for the high-level overview provided in the preceding sections. Section 6 is devoted to key insights and take-aways and includes the discussion of promising research opportunities.

# 2 Generalized Problem Definition

This section first investigates a sequential decision process for integrating demand management and vehicle routing from an application-oriented perspective. We identify four essential components that are part of this process and present dimensions that characterize each component as well as possible realizations of each dimension in Section 2.1. Subsequently, we discuss prototypical applications in Section 2.2. The purpose is to show how different realizations of the dimensions relate to real-world implementations.

## 2.1 Sequential Decision Process

Providers that offer online-to-offline logistical services regularly face stochastic and dynamic decision problems that arise over time on an operational level. Such problems can be described as sequential decision processes, which cast the overall problem as a sequence of states (Powell, 2019). In each state, the provider must collect and evaluate (stochastic) information concerning customers, logistical resources, i.e., vehicles, and, possibly, the environment (Soeffker et al., 2022). Depending on the information's evaluation, they must also make different types of decisions.

To analyze the problem characteristics, we decompose the resulting decision process into four components for each state. Two of the components include *interactions* with customers, the remaining two deal with the provider's *decisions*. Different types of events may trigger these decisions. Fig. 1 shows the components and their relationships. We explain them in the following and introduce dimensions by which we characterize different realizations of the components as part of a morphological analysis. This technique allows us to systematically describe the entire spectrum of decision problems by reducing the problems to these key dimensions with a set of possible realizations.





**Request capture**: The arrival of a customer during a sales period, called the booking horizon, triggers this component. The provider can sell different *types of services*: pure transportation (e.g., a ride), transportation in combination with selling goods (e.g., groceries), or transportation in combination with selling ancillaries (e.g., installation). We refer to the latter two as coupled goods and coupled services, respectively. The customer makes a request by specifying parameters

of the service wanted, e.g., using a web application or via a call-center. These parameters can be origin and destination, time and mode of transport, and coupled goods or services. The provider must capture these parameters as input for their decisions.

**Demand management**: This component follows request capturing and must control demand with respect to the provider's *objective*. It tries to exploit that usually several feasible options for service fulfillment exist. Then, it aims at selling the available capacity in a way that maximizes a measure of profit. The profit comprises several components that represent revenues and costs. On the revenue-side, the fees for the logistical service itself and the revenues/profits of coupled goods or services may be relevant. On the cost-side, the unit costs of the coupled goods or services, possible discounts, and the transportation costs must be considered. Depending on the application, also the number of orders, i.e., accepted requests, may serve as an objective. Regardless of the objective function, the provider must ensure that the logistical services sold can be fulfilled subject to operational constraints. The implementation of demand management can be characterized along the following dimensions (Agatz et al., 2013):

- Concerning the *time of decision*, static and dynamic controls can be distinguished. Static controls determine all decisions before the start of the booking horizon based on exogenous information. They do not adjust them depending on endogenous information concerning customers but check for feasibility. As an example, an AHD provider may publish a static price list for their delivery time slots, which is valid for multiple weeks. During each booking process, any customer will be able to place an order at the price of the published delivery fee as long as the provider can feasibly fulfill the order. By contrast, dynamic controls make decisions based on the information becoming available during the booking and service horizon. Beyond the current request's parameters, such information includes existing orders, the vehicles' locations, and loads. In this case, an AHD provider would, e.g., offer individual delivery fees determined at the time of each customer request arrival based on the delivery location and the shopping basket value.
- To influence the customers' choices favorably, the provider can apply two *control types*, namely availability control or price-based control. In availability control, the provider makes decisions on which feasible fulfillment options to offer to the customer, e.g., when prices are fixed. In price-based control, they set fees for the different options. The set of fulfillment options along with their prices form an offer set, from which a customer can choose.
- Finally, the provider can decide on request *processing*, i.e., between real-time processing for single requests or batch processing. In the first case, the provider implements decisions immediately. In the second one, they postpone decisions until, e.g., a specific batch size or state is reached.

**Order confirmation**: After constructing the offer set, the order confirmation component represents a second interaction with the customers, which consists in presenting offer sets to customers and, potentially, closing a deal. If the provider generally offers only a single *fulfillment option*,

customers will either buy or not. If they provide an assortment of multiple options, e.g., different time windows for the transportation, customers will choose an option, which is potentially the nopurchase option, according to some individual preference function, e.g., by maximizing their utilities. If a sale takes place, the corresponding option becomes an order.



Fig. 2 Booking and service horizon

Vehicle routing: The vehicle routing component is executed before or during the period of service fulfillment, called the service horizon. Its task consists in determining feasible and costminimizing route plans for the given orders. Booking and service horizons can either be disjoint, overlapping, or infinite as illustrated in Fig. 2. In the first case, the provider collects orders until a cutoff time, which lies before the beginning of the service horizon. Here, the provider can postpone definitive vehicle routing decisions until the end of the booking horizon. However, sometimes they may perform tentative route planning as an input for demand management decisions. If the horizons overlap or are infinite, the provider needs to finalize routing decisions before the end of the booking horizon. Here, several events may trigger a decision for a given state. First of all, a new order may have been accepted. Other events include that a vehicle has become idle or must act, e.g., leave the depot, to fulfill operational constraints. Also, it can be reasonable to move a vehicle to another position to be better prepared for future requests. In the latter cases, the vehicle routing component is executed without a customer arrival. Depending on the transportation service sold, the provider must solve different types of vehicle routing problems (e.g., Toth and Vigo, 2014). For example, delivery or pickup problems may occur. Also, point-to-point problems may arise. Finally, routing decisions may be subject to different types of *constraints*. These may refer to the fleet size or composition, the vehicles' capacity, or service guarantees like delivery within a specific time window.

The morphological box in Table 1 summarizes the result of the morphological analysis, i.e., it describes the different components based on the dimensions and their potential realizations introduced above. Besides providing a compact summary, it also serves as a tool for further analyses. Specific decision problems, including novel ones, can be derived by selecting a certain realization for each dimension and combining them. In turn, existing decision problems can be classified according to their realizations for each dimension. In the survey at hand, we present the latter type of analysis for prototypical decision problems (Table 2) and decision problems considered in the existing literature (Table 4).

Process component	Dimension	Realization		
Request capture	Service type	Transportation (TR)	Coupled goods (CG)	Coupled services (CS)
Demand management	Objective	Profit (PR)	Revenue (RE)	Number of orders (NO)
	Time of decision	Static	Dynamic	
	Control type	Availability (AV)	Price-based (PB)	
	Processing	Real-time (RT)	Batch (BA)	
Order confirmation	Fulfillment options	Single (SI)	Multiple (MU)	
Vehicle routing	Booking/service horizon	Disjoint (DJ)	Overlapping (OL)	Infinite (IF)
	Routing problem	Delivery (DE)	Pick-up (PU)	Point-to-point (PP)
	Constraints	Fleet	Vehicles	Service guarantees

Table 1 Components and dimensions of the sequential decision process

#### 2.2 Applications

This section discusses prototypical applications for which integrating demand management and vehicle routing has already been established or is currently evolving. We deliberately do not explicitly refer to specific companies' existing applications because the underlying business models are adapted fast and refined continuously. However, in all cases, corresponding services exist and can easily be found by simple internet search. In Table 2, we describe the prototypical applications based on the morphological box (Table 1) developed in Section 2.1. Table 4 in Section 4.3 will characterize the related specific problems considered in the existing academic literature.

The most prominent application for AHD is e-groceries (Agatz et al., 2013). Here, transportation is combined with the sales of groceries. Most commonly, the providers try to maximize profit after fulfillment. This profit is determined by the profit per order, which considers the profit of the shopping basket plus the delivery fee, minus the cost of transportation. In the early days of AHD, the usual way to control demand was to define combinations of delivery areas and time windows. For these combinations, the provider computed static prices and the maximal number of customers to be served prior to the booking horizon which led to a form of availability control. Thus, it was possible to provide customers with feedback on fulfillment options after filling their shopping basket in real-time. Until recently, booking and service horizon have usually been disjoint. Customers had to place their orders until the evening before the delivery day. For all orders accepted, the provider must solve a capacitated vehicle routing problem with time windows. Please note that modern approaches do not only set prices dynamically but also offer overlapping time windows of different lengths.

SDD is also used for selling groceries (Archetti and Bertazzi, 2021). New market entrants currently try to establish services that deliver a restricted assortment of food products within very short deadlines. Established players like large grocery and wholesale retailers are experimenting with combining SDD and next-day delivery. However, the concept was initially introduced for courier and express services, the reason why we discuss a corresponding application here (Ghiani et al., 2009). Such services offer pure transportation for, e.g., pharmaceutical drugs or spare parts. Since the provider's capacity is usually fixed on a given day, they maximize the total revenue as a proxy for profit. Depending on the transport's origin and destination and the delivery deadline, the provider dynamically calculates a fee, i.e., sets a price. Again, the provider must process a captured request in real-time. New orders arrive while executing others, i.e., the booking and service horizon overlap. Hence, the provider must deal with a dynamic point-to-point (pickup and delivery) problem with deadlines.

Process AHD SDD MOD FSO Dimension component Request Coupled goods Service type Transportation Transportation Coupled services capture Objective Profit Number of orders Number of orders Revenue Time Static Dynamic Dynamic Dynamic Demand management Control type Availability Price-based Availability Availability Processing Real-time Real-time Real-time Batch Order Fulfillment Single Multiple Single Multiple confirmation options Booking/ Disjoint Overlapping Overlapping Infinite service horizon Vehicle Routing Point-to-point Point-to-point Delivery Delivery routing problem Waiting and Time windows, Delivery dead-Time windows, Constraints travel time worker skills vehicle capacity lines

Table 2 Sample applications

An increasingly popular form of public transport is MOD (Hazan et al., 2019). The transportation service is provided using mini-buses and taxis in a shared-ride mode. Public providers may aim at maximizing the number of orders, i.e., rides, performed. Customers can specify the origin and destination and the earliest pick-up or latest arrival time. The fee depends on the origin and destination and is commonly based on published tariffs, such that only the availability is subject to dynamic control. Hence, based on their request and the capacity utilization, customers are either offered a ride or are rejected in real-time. In the first case, a single option is provided which comes with a travel time, a possible waiting time, and the number of passengers on the ride. The customers can then accept the option or reject it. With the switch from call center- to application-based reservation systems, providers have allowed to make reservations on the day of travel lead-ing to overlapping booking and service horizons. Again, a point-to-point (dial-a-ride) transportation problem results whose constraints must consider the vehicles' capacities and ride-specific aspects like waiting and travel times.

FSO represents an emerging application of integrated demand management and vehicle routing (Chen et al., 2016). In a business-to-consumer context, customers receive some furniture,

electronics, or home appliances and may require a coupled service like installation for the items delivered. In a business-to-business context, on-site maintenance and repair may represent possible use cases. In the first case, which we consider here, it is common that the customer can select several options from a menu of delivery dates with corresponding time windows, i.e., the provider deliberately restricts the availability of options by availability control. When determining the corresponding offer sets, the provider usually tries to maximize the number of installations. Some days ahead of delivery, the provider informs about which of the customer's options they have chosen for installation. Since lead times for the products can depend heavily on the different products, the problem on hand has no finite horizon. New orders for products with a short lead time can arrive and be ready for installation while waiting for the completion of orders with longer ones. Like for AHD, the provider must solve a capacitated vehicle routing problem with time windows. However, additional constraints like worker skills come into play. Often, corresponding routing problems are identified as technician or field service routing problems.

## 3 Mathematical Model Formulation

In this section, we discuss the formalization of the generalized problem definition described in Section 2 by means of mathematical modeling. Since the problem at hand is stochastic and dynamic, an accurate formalization requires a dynamic control model, which is subject of Section 3.1. An integral element of this formalization is also the modeling of the customers' choice behavior, provided that they are given a choice between fulfillment options as part of the order confirmation component. Therefore, we elaborate on these customer choice models separately in Section 3.2.

## 3.1 Dynamic Control Model

Mathematically, Markov decision processes (MDPs) provide the foundation for describing most decision problems in demand management and vehicle routing. However, in contrast to, e.g., deterministic vehicle routing, it is not standard in the literature to present a corresponding MDP model, which is an observation already made by Ulmer et al. (2020) for stochastic, dynamic vehicle routing. Reasons may be that the notation is quickly becoming complex and awkward to handle. Moreover, solution approaches are generally approximative and do not rely directly on an exact dynamic control model. Further, the variety of problems leads to rather specific models from a notational point of view (e.g., Al-Kanj et al., 2020, Ulmer et al., 2019, Xu et al., 2018, or Yang et al., 2016). Therefore, in the following, we synthesize the models from existing works and provide a generalized, high-level model formulation. We structure the discussion along the model's primary building blocks using the language and notation common for MDPs (e.g., Powell, 2019). For similarly generic models, we refer to Klein et al. (2020), who propose a route-based modeling framework for dynamic routing.

In the model, demand is represented as a set of potential customers  $\mathcal{I} = \{1, ..., I\}$ . Each customer  $i \in \mathcal{I}$  comes with a location and has different preferences for the services offered. To serve the customers, the provider has vehicles  $h \in \mathcal{H} = \{1, ..., H\}$  available. The vehicles may have several restrictions concerning their capacity, which may refer to the maximal feasible load, the maximal travel distance, or the maximal travel time due to working shifts. Based on these assumptions, we describe the building blocks of MDP models. For each possible variant of modeling a certain building block, we provide exemplary references. Please note that the notation chosen makes several deliberate simplifications for the sake of readability. For example, numbers of customers *I*, in general, are stochastic. Furthermore, we omit indices where possible, and following Al-Kanj et al. (2020), we indicate unambiguous state-dependencies by an index *k*.

**Decision epochs**: The booking horizon and the service horizon encompass  $k \in \mathcal{K} = \{0, ..., K\}$  decision epochs, whose number can be stochastic. Decision epochs represent points in time at which the provider must make a demand management decision, a routing decision, or both. Three types of events can trigger a decision epoch, with the latter two only being relevant for problems with overlapping horizons. The first one is the arrival of a customer request (Ulmer, 2020a). Secondly, routing-related events may require decisions, e.g., if a vehicle becomes available after completing an order (Ulmer et al., 2018). Thirdly, a new decision epoch can be defined to occur after a certain amount of time in which vehicles were idle or orders were not assigned for fulfillment (Chen et al., 2019).

**States**: Tuples  $S_k = (S_k^{cust}, S_k^{peh})$  describe the system's state at the beginning of a decision epoch k and contain all information necessary to make a decision. The vectors  $S_k^{cust}$  and  $S_k^{peh}$  describe the customers' and vehicles' statuses. For customers, this status may indicate which customers are currently requesting service. Additionally, in case the provider receives orders, information on the orders' parameters (Koch and Klein, 2020) and, for problems with overlapping horizons, the fulfillment status is stored (Chen et al., 2022). For vehicles, the status may refer to the current location (Qiu et al., 2018), the time of arrival at the next customer (Chen et al., 2019) or at the depot (Voccia et al., 2019), or a route plan (Ulmer and Thomas, 2020). Note that information on vehicles is not required for problems with disjoint horizons because final routing is not necessary before the end of the booking horizon.

**Decisions**: Depending on the state in decision epoch k, the provider must either make a demand management decision and, potentially, a corresponding vehicle routing decision, or a stand-alone routing decision. When booking horizon and service horizon are disjoint, demand management decisions suffice. The decisions are summarized by variables  $x_k = (x_k^{dem}, x_k^{rout})$  that describe the actions taken and are defined as follows:

Vehicle routing decisions x<sub>k</sub><sup>rout</sup>: If the provider makes a routing decision x<sub>k</sub><sup>rout</sup> for state S<sub>k</sub>, they select a feasible route plan φ<sub>k</sub> = {ρ<sub>h</sub>: h ∈ H}, i.e., determine a route ρ<sub>h</sub> for each vehicle h ∈ H (Ulmer, 2020a). A route plan is called feasible if it does not violate any operational

restriction. In this context, the term route plan has a fairly broad meaning, i.e.,  $x_k^{rout}$  may only state which order to serve next for each vehicle (e.g., Xu et al., 2018). The set of all feasible route plans in state  $S_k$  is denoted by  $\Phi_k$ . In case the booking horizon and service horizon are disjoint, a single routing decision is made at decision epoch K + 1 (Klein et al., 2018), i.e., at the end of the booking horizon.

Demand management decisions x<sub>k</sub><sup>dem</sup>: A demand management decision x<sub>k</sub><sup>dem</sup> determines which offer the provider makes for providing a service requested by customer *i* at decision epoch *k*. The feasible fulfillment options available are given by O<sub>k</sub> = {1, ..., O<sub>k</sub>}. An option o ∈ O<sub>k</sub> is called feasible if a feasible route plan φ<sub>k+1</sub> exists when the request turns into an order due to the sale of o. When applying availability control, the provider determines an offer set Θ<sub>k</sub> ⊆ O<sub>k</sub> (Avraham and Raviv, 2021). Analogously, when using price-based control, the provider sets prices (service fees) p<sub>oi</sub> for all options o ∈ O<sub>k</sub> (Prokhorchuk et al., 2019).

**Transitions**: Transitions between states  $S_k$  and  $S_{k+1}$  may occur for several reasons: If customer *i* decides (stochastically) to buy an option *o*, the request becomes an order and  $S_k^{cust}$  is updated accordingly. The same holds if customers are served as the provider (partially) executes route plan  $\phi_k$ . In this case, the vehicles' status  $S_k^{veh}$  is also updated (Voccia et al., 2019). Mathematically, the transition can be described by a state equation  $S_{k+1} = S^M(S_k, x_k, W_{k+1})$ .  $W_{k+1}$  represents random variables affecting the transition from epoch *k* to *k* + 1. In our case, these include, e.g., the choice of customer *i*, the preferences and locations of incoming customers (Mackert, 2019), or stochastic travel times (Xu et al., 2018).

**Rewards**: If the provider sells an option *o* to a customer *i*, they obtain a reward  $R_k(S_k, x_k) = r_{oi}$ . Usually,  $r_{oi}$  represents the revenue per order or the profit per order possibly depending on a charged price (service fee)  $p_{oi}$  (Strauss et al., 2021). If the objective is to maximize the number of customers served, the reward is set to  $r_{oi} = 1$  (Ulmer et al., 2019). Fulfillment costs can be modeled as negative rewards that are incurred once the respective routing decisions become definitive and the route plan is (partly) executed (Klapp et al., 2018). For disjoint horizon problems, the terminal reward  $R_{K+1}$  summarizes all fulfillment cost (Yang et al., 2016).

**Policy**: A policy  $X^{\pi}(S_k)$  is a rule or function that determines a decision  $x_k$  for a state  $S_k$ . Here, it refers to vehicle routing and demand management decisions, which are often intertwined. For example, when deciding on an offer set, the provider may have to simultaneously make routing decisions anticipating the possible sale.

**Objective function**: In general, since the problems are stochastic, the objective consists of maximizing expected rewards (including terminal cost  $R_{K+1}$ ):

 $J(X^{\pi}) = \mathbb{E}\left\{\sum_{k=0}^{K} R_k(S_k, X^{\pi}(S_k)) + R_{K+1}\right\}$ 

In infinite state problems, we can discount rewards and define the objective as the limit of the expression above, when  $K \to \infty$  (Holler et al., 2019).

Value function: To evaluate possible decisions in state  $S_k$ , we define the value function  $V_k(S_k)$  the provider wants to maximize. It represents the objective function value at the end of the booking and service horizon that can be expected at decision epoch k by the corresponding Bellman equation:

$$V_k(S_k) = \max_{x_k} \mathbb{E} \{ R_k(S_k, x_k) + V_{k+1} (S^M(S_k, x_k, W_{k+1})) \}$$

Thus,  $J(X^{\pi}) = V_0(S_0)$  holds if  $X^{\pi}$  is an optimal policy. The correct computation of the value function requires optimal demand management decisions for future requests and optimal routing decisions for existing and future orders. Alternatively, it is possible to formulate a Bellman equation based on state-action values (Kullman et al., 2022).

#### 3.2 Customer Choice Modeling

In case the order confirmation component allows customers to select a fulfillment option from an offer set, any dynamic control model must include a customer choice model. Otherwise, if there is no such interaction during order confirmation, choice modeling can be omitted. More precisely, a choice model predicts a purchase probability  $P_o(\Theta_k)$  for each option  $o \in \Theta_k$  with respect to the offer set  $\Theta_k$  and, possibly, prices  $p_{oi}$ . For this purpose, parametric, non-parametric, and multi-stage models exist (Strauss et al., 2018 and Berbeglia et al., 2022).

In the context of vehicle routing applications, parametric models rooted in random utility theory dominate. Following this theory, each customer *i* evaluates the set of offered alternatives with respect to an individual utility function before deciding on either buying one option  $o \in \Theta_k$  or leaving the market (e.g., Train, 2009). In general, we assume that the resulting utility for an option  $o \in \Theta_k$  has a deterministic and a random part. Customers decide on the alternative that maximizes their utility. If  $|\Theta_k| > 1$ , customers may substitute across all  $o \in \Theta_k$ , in case their preferred one is not available (e.g., Kök and Fisher, 2007). In the literature, the existence of such substitution behavior is widely acknowledged (e.g., Ulmer, 2020a, Yan et al., 2020, or Yang et al., 2016). Thus, the resulting purchase decision is stochastic and depends on the characteristics of all options  $o \in \Theta_k$  including, if applicable, their prices  $p_{oi}$ .

The purpose of choice modeling is to obtain purchase probabilities for each  $o \in \Theta_k$ , which serve as input parameters for demand control. To this end, the specification of a utility function is necessary for random utility models. The deterministic part is usually expressed as a linear function of a vector of attributes that influence the purchase probabilities. In last-mile logistics, these include the associated time slot (e.g., Yang et al., 2016) and the delivery deadline (Prokhorchuk et al., 2019). Similarly, for passenger transportation, attributes encompass travel time (Atasoy et al., 2015 and Qiu et al., 2018) as well as origin, destination, and time of day (Al-Kanj et al., 2020). Also, the price  $p_{oi}$  represents an attribute if fees are charged.

Different choice models are obtained depending on the assumptions made on the distribution of the random utility part. Thereby, it is crucial to consider that model selection and model
specification significantly impact the quality of demand management decisions and the complexity of demand control (Berbeglia et al., 2022). The estimation of the utility function's parameters from historical data is also an optimization problem and can be of varying complexity.

With respect to our domain, authors use the following random utility models:

- Multinomial logit (MNL) model: This is the most prominent model. It assumes that the entire customer population can be described by a common utility function. Furthermore, it assumes that the random utility components are independent and identically distributed random variables following a Gumbel distribution. If  $O_k = 1$ , the MNL reduces to a binary logit model (Al-Kanj et al., 2020). In comparison to other random utility models, the MNL has advantages in terms of computational complexity (Berbeglia et al., 2022). However, it is not sufficiently accurate in many applications, even with a nearly perfect specification: First, it does not capture latent customer preferences. Second, the model suffers from the IIA property (independence from irrelevant alternatives) and therefore only allows for proportional substitution behavior (Train, 2009).
- Generalized attraction model: Compared to the MNL model, it captures customer dissatisfaction and thus reduces purchase probabilities for all offered products if the cardinality of an offer set is low (Gallego and Topaloglu, 2019).
- Finite mixture MNL model: This model assumes that demand is composed of homogeneous segments whose choice behavior can be described by standard MNL models (Strauss et al., 2018). If the segment affiliations of the arriving customers are unknown, the integration of the model into demand control significantly increases its complexity (Koch and Klein, 2020). The same holds for the parameter estimation problem. Otherwise, the segment-specific MNL models are independent, and there is no increase in complexity (e.g., Lang et al., 2021b).
- Nested logit model: The nested logit (NL) model is appropriate if we can aggregate alternatives into nests in a way such that the IIA holds within each nest but not across nests. Each nest represents a set of substitutes. The model by Wang et al. (2021) accounts for alternate pick-up and drop-off points customers can choose. Köhler et al. (2019) and Strauss et al. (2021) use the NL model to reflect demand interdependencies and non-negligible disproportional substitution behavior due to offering overlapping time windows of different lengths. Because of the higher complexity of demand management decisions, Strauss et al. (2021) approximate the NL model by a standard MNL model.

Lastly, some authors propose parametric models that are specifically designed for pricing control and are not rooted in random utility theory (Campbell and Savelsbergh, 2006, Chen et al., 2019, Haliem et al., 2021, Klein and Steinhardt, 2023, Ulmer, 2020a, and Vinsensius et al., 2020).

## 4 Solution Concepts

In this section, we discuss solution concepts for dealing with decision problems that fall under the generalized problem definition as presented in Section 2. Due to the problems' complexity, directly solving corresponding dynamic control models (Section 3) to optimality is computationally intractable. As the state space is very large even for small instances, it is not possible to evaluate, e.g., the Bellman equation for each potential state. Moreover, in each state, the determination of demand management decisions and vehicle routing decisions can represent challenging optimization problems of their own.

Instead, the existing literature follows two basic solution concepts, both based on approximations. In Section 4.1, we first describe decomposition-based approximations. Section 4.2 is devoted to static deterministic approximations. In Section 4.3, we merge the results of our analyses of problem characteristics and solution concepts in the form of a tabular overview. Thus, we only provide exemplary references in all the following subsections and refer the reader to Table 4 and Table 5 for the extensive classification of all works.

#### 4.1 Decomposition-based Approximation

In the academic literature, most authors resort to a decomposition-based approximation. For this purpose, they identify major tasks in the overall decision process to be addressed by the provider. Then, they formalize the tasks and solve corresponding subproblems or combinations of them sequentially. Different types of solution approaches exist: Sometimes, the authors explicitly formulate auxiliary or simplified mathematical models for the problems that are then tackled using a general-purpose solver or some special-purpose algorithm. In other cases, they only describe the problems verbally, propose a conceptual model to, e.g., deal with stochasticity or interdependencies among problems, and again, provide suitable algorithms. We define the tasks in Section 4.1.1 and describe the corresponding solution approaches in Sections 4.1.2-4.1.5. Solution algorithms are subject of Section 5.1.

#### 4.1.1 Task Definitions

In Section 2.1, we have identified two components that require the provider to make decisions: demand management and vehicle routing. When analyzing corresponding research papers, it turns out that authors consider up to three different tasks to support demand management decisions  $x_k^{dem}$ . Fig. 3 shows the sequence of these tasks and the input data they provide for the succeeding task. Routing control can be viewed as a fourth task associated with the vehicle routing component.



Fig. 3 Tasks of demand management component

**Feasibility check**: First, the provider must determine the set  $O_k$  of feasible options with respect to existing orders in state  $S_k$ . The exact type of the corresponding vehicle routing problem

depends on the application. In case the vehicle routing problem has a feasible solution, this implies that  $o \in O_k$ .

**Cost estimation**: Second, the provider must compute the value difference, i.e., the costs,  $\Delta V(S_{k+1}|o) = V_{k+1}(S_{k+1}) - V_{k+1}(S_{k+1}|o)$  for each feasible option  $o \in O_k$  in case the provider sells option o to customer i due to demand management decisions  $x_k^{dem}$  compared to not selling it. Hence, the result of the feasibility check is an input for cost estimation. The impact of selling option o is twofold: First, it can lead to the displacement of demand arriving later, in case not enough capacity will be left. Hence, a sale influences future rewards via the *displacement cost* well known from revenue management (Talluri and van Ryzin, 2004a). Second, due to deliveries, it also impacts the costs-side because the usage of some resources causes non-negligible (future) transport costs that are not attributable to requests ex-ante. These costs are captured by the term *marginal delivery cost* or *marginal cost-to-serve* (e.g., Yang and Strauss, 2017). However, due to the "curses of dimensionality" (Powell, 2011), i.e., the large number of possible states and actions, cost values  $\Delta V(S_{k+1}|o)$  can usually only be approximated by an estimate  $\Delta \tilde{V}(S_{k+1}|o)$ .

**Demand control**: Based on a cost estimate for each feasible option, the provider must make a demand management decision  $x_k^{dem}$ :

- When using price-based control, the provider again only offers an option o ∈ O<sub>k</sub> if r<sub>oi</sub> ≥ ΔṼ(S<sub>k+1</sub>|o) where r<sub>oi</sub> includes the price p<sub>oi</sub>. Hence, the ΔṼ(S<sub>k+1</sub>|o) represents a lower bound for the reward r<sub>oi</sub>, from which a lower bound for the price (service fee or discount) p<sub>oi</sub> can be derived. Based on this information, the provider can optimize prices to influence demand.

**Routing control**: The final task results from the vehicle routing component and consists in making routing decisions  $x_k^{rout}$ . Again, the feasibility check provides a crucial input to ensure that routing decisions do not violate the operational constraints.

As we show in the following sections, there exist individual solution approaches for each task. Yet, as the tasks build upon each other, the corresponding subproblems are often related. For example, explicit route planning approaches can be applied to feasibility check, cost estimation, and routing control. Therefore, one could argue that solution approaches exist that solve tasks in combination. However, for the sake of clarity, we discuss the approaches for each task individually (Sections 4.1.2-4.1.5). Table 3 provides an overview of the fundamental solution approaches for each task.

Task	Solution approach			
Feasibility check	Route-based (RO)	Capacity-based (CA)		
Cost estimation	Myopic (MY)	Sampling-based (SA)	Deterministic linear program (DL)	Predictive (PR)
Demand control	Accept/reject (AR)	Assortment optimization (AO)	Discrete pricing (DP)	Continuous pricing (CP)
Routing control	Full route plan (FP)	Single route (SR)	Leg-oriented (LO)	

Table 3 Overview of task-specific solution approaches

#### 4.1.2 Feasibility Check

As stated before, the provider can check the feasibility of a potential order as a separate task. In this case, we can distinguish two types of checks:

**Route-based check**: This type solves some auxiliary model that explicitly considers the constraint satisfaction version of a vehicle routing problem for each fulfillment option o being a candidate for  $\mathcal{O}_k$  (e.g., Brailsford et al., 1999 and Berbeglia et al., 2011 or Elting and Ehmke, 2021 in the context of point-to-point transportation). The models are deterministic because the already existing orders and the option o are known for a state  $S_k$ . In case a solution exists for the resulting instance, o is included in  $\mathcal{O}_k$ .

**Capacity-based check**: These checks determine capacity limits for the number of feasible orders depending on criteria like the location or the time of delivery (e.g., Lang et al., 2021a) and thereby approximate the constraint satisfaction problem. During the booking horizon, an option o is considered feasible, i.e., included in  $O_k$ , if the number of similar orders with respect to the criteria is below the capacity limit. Capacity-based feasibility checks are generally more suitable for disjoint-horizon problems because no routing decisions are required during the booking horizon and, thus, route-based planning is not essential.

## 4.1.3 Cost Estimation

The literature distinguishes between myopic cost estimation and anticipative cost estimation depending on the use of information.

Myopic estimation solely incorporates information about orders that have already been received (Haferkamp and Ehmke, 2022) and does not require any (probabilistic) information about future demand. Therefore, it only aims at marginal cost-to-serve and does not capture a decision's impact on future rewards, i.e., neglects displacement cost. However, the reduced data requirements compared to anticipative estimation can also be a significant advantage in practice if data on future demand are sparse, unreliable, or even not available at all. Usually, myopic estimation relies on a formulation of a static routing problem, so that marginal cost-to-serve is estimated as the increase in total routing cost caused by adding another order to the respective problem instance.

If information about future demand is available, anticipative estimation is applicable. It addresses two aspects to improve the estimate. First, it can achieve a more accurate estimate of marginal cost-to-serve compared to myopic estimation. For example, this cost may be overestimated in myopic estimates if not considering consolidation opportunities with future orders. Second, anticipation enables an approximation of displacement costs in the first place. Not surprisingly, empirically, many studies demonstrated that anticipative estimation yields better results compared to myopic estimation (Section 6). However, the extent to which this potential can be realized in practice depends on the quality of available data regarding future demand.

Depending on which techniques are used to deal with uncertainty, i.e., characteristics of future requests including the customers' preferences, we distinguish three subclasses of anticipative approaches, namely *sampling-based approaches*, *deterministic linear programming approaches*, and *predictive approaches*. In the following, we characterize these subclasses:

- **Sampling-based**: To obtain a more precise estimation of marginal cost-to-serve, several authors propose the inclusion of sampled future orders into a single (tentative) route plan or a pool of tentative route plans, i.e., a static routing problem. If the corresponding problem allows displacements of sampled orders, its solution also yields an estimate of displacement cost. The idea behind this type of approaches, known as scenario-based planning, is to anticipate how the instance of the routing problem will be structured at the time a potential order is fulfilled. The resulting gain of accuracy is particularly high in the early phase of each booking horizon (Yang et al., 2016). The concept goes back to Bent and van Hentenryck (2004) and Ichoua et al. (2006), who apply it to pure dynamic vehicle routing problems. While scenario-based planning considers the future evolution of the decision process from a hindsight perspective, sampling is also possible by dynamically simulating the evolution of the decision process from the current state onward over a limited horizon (Soeffker et al., 2022). This is the principle of rollout approaches, which provide an estimate of both cost components as future decisions are simulated sequentially according to a base policy (e.g., Ulmer, 2020b).
- Deterministic linear programming: Originally developed in revenue management (Gallego and Topaloglu, 2019), several publications show that deterministic linear programming techniques are transferable to the field of vehicle routing. They define corresponding auxiliary models, which provide two types of information: On the one hand, the objective function value approximates a certain state value, and hence, the model can be solved twice to calculate a cost estimate  $\Delta \tilde{V}(S_{k+1}|o)$  (e.g., Klein et al., 2018). On the other hand, the solution yields information that may also serve directly as an input for demand control. Such models are related to sampling-based approaches in that they also assume expected future demand to be deterministic and include it as an input in aggregated or disaggregated form. The goal is to use this information to predict the expected evolution of the remaining booking and service horizon depending on the demand management and routing decisions. To model customer choice behavior, the inclusion of choice models (Section 3.2) is also possible.

- **Predictive**: A considerable number of authors use predictive models borrowed from the field of statistical learning (Powell, 2019). We can distinguish three types of solution approaches depending on the values to be predicted:
  - The first type approximates the *state value* function  $\tilde{V}_{k+1}(S_{k+1}|o)$  for each resulting state  $S_{k+1}$  and option o to calculate the cost  $\Delta \tilde{V}(S_{k+1}|o)$  as a value difference  $\tilde{V}_{k+1}(S_{k+1}) \tilde{V}_{k+1}(S_{k+1}|o)$  (e.g., Lang et al., 2021a).
  - The second one provides a *direct cost approximation*  $\Delta \tilde{V}(S_{k+1}|o)$  (e.g., Qiu et al., 2018).
  - Finally, the third one predicts *state-action values* by Q-learning based on approximating the value of a demand management decision in a particular state. Since maximizing the state-action value in a state  $S_k$  directly leads to an optimal solution for demand control, an explicit cost calculation is no longer required (e.g., Chen et al., 2023).

Any type of prediction can generally be encoded using three types of approximations (Powell, 2011). All of these have in common that values are computed dependent on a set of preselected features representing the state in an aggregated form. Besides the decision epoch k this may include order characteristics as well as route-based features of tentative routes like the vehicles' idle times. The approximations are:

- Lookup tables: They store an estimate for all possible resulting combinations of feature values, which is updated each time one of the corresponding states occurs throughout the learning process (e.g., Ulmer et al., 2018).
- Parametric approximations: They represent the prediction by an expression of a particular functional form dependent on a set of parameters and the feature values. Most often, a *linear* function, i.e., the weighted sum of all feature values, is chosen (e.g., Yang and Strauss, 2017). However, *piecewise-linear* or *non-linear* specifications are also possible (e.g., Ni et al., 2021 and Lebedev et al., 2020).
- Non-parametric approximations: In contrast to parametric ones, these approximations
  do not assume that the relationship between the estimate and the feature values is of a
  particular functional form. Therefore, they can adapt more flexibly to the actual functional relationship, which is likely non-linear. Examples are kernel regression and (deep)
  neural networks (e.g., Dumouchelle et al., 2021).

## 4.1.4 Demand Control

The demand control task yields the demand management decisions  $x_k^{dem}$  that are made in response to an arriving request in stage  $S_k$ . For optimizing the demand management decision, potentially based on customer choice behavior, three types of control are proposed in the literature: *accept/reject* control, *assortment optimization*, and *pricing* control.

Accept/reject: If the order confirmation step does not involve any stochastic customer choice decision, demand control boils down to an accept or reject decision for each request. The resulting subproblem can be cast in two ways, both derived from traditional demand management

applications (Talluri and van Ryzin, 2004a). First, the provider can subdivide the set of possible requests into subsets according to certain parameters and assign a booking limit to each subset, i.e., an upper bound on the number of orders (Giallombardo et al., 2022). In this case, a request is accepted if this does not cause the corresponding limit to be exceeded. Second, the cost estimate (Section 4.1.3) can serve as a bid price, i.e., as the minimum profit of a request for it to be accepted. This type of control is also applicable for batched request processing (Ulmer et al., 2018).

Assortment optimization: Under the assumption of substitution behavior and multiple fulfillment options, the demand control task is called an assortment optimization problem (see Gallego and Topaloglu, 2019 for an in-depth introduction). Due to the decision space growing exponentially with  $O_k$ , i.e., the number of fulfillment options, it becomes combinatorial. Given  $O_k$  as well as  $r_{oi}$ ,  $\Delta \tilde{V}(S_{k+1}|o)$ , and the offer set-dependent purchase probabilities  $P_o(\Theta_k)$  for all  $o \in \Theta_k$  and  $\Theta_k \subseteq O_k$  provided by the choice model, the objective is to maximize the expected profit after fulfillment:

$$\Theta_{k}^{*} = \operatorname*{argmax}_{\Theta_{k} \subseteq \mathcal{O}_{k}} \left\{ \sum_{o \in \Theta_{k}} P_{o}(\Theta_{k}) \cdot \left( r_{oi} - \Delta \tilde{V}(S_{k+1}|o) \right) \right\}$$

If necessary, certain structural properties of the offer set can be specified by adding constraints. Additionally, problem structure and problem complexity depend on the choice model (Section 3.2).

**Pricing**: The basic principle of price-based control is to offer each feasible option  $o \in \Theta_k = O_k$ at some dynamic price  $p_{oi}$ , i.e., determine a price vector  $\mathbf{p}_i = (p_{oi})_{O_k \times 1}$ . Thus, rewards  $r_{oi}(p_{oi})$ depend on the respective price  $p_{oi}$ . In general, pricing optimization requires the same types of input data as assortment optimization, and the problem structure again depends on the choice model defining the purchase probabilities  $P_o(p_{oi}, \Theta_k)$ . The decision space, i.e., the feasible price vectors, can be similarly vast even if restrictions are imposed. In case the price is only subject to an upper or a lower bound or is entirely unrestricted, a *continuous pricing* problem results, which is modeled as follows (e.g., Yang et al., 2016):

$$\boldsymbol{p}_{i}^{*} = \operatorname*{argmax}_{\boldsymbol{p}_{i}} \left\{ \sum_{o \in \Theta_{k}} P_{o}(p_{oi}, \Theta_{k}) \cdot \left( r_{oi}(p_{oi}) - \Delta \tilde{V}(S_{k+1}|o) \right) \right\}$$

Specifying a set of feasible price points leads to a *discrete pricing* problem, which is a special case of the assortment optimization problem described above. Alternatively, auxiliary models based on quadratic programming (Campbell and Savelsbergh, 2006 and Vinsensius et al., 2020) and predictive models (Chen et al., 2019 and Al-Kanj et al., 2020) are proposed in the academic literature. Finally, note that discounts can also be modeled by allowing  $p_{oi} < 0$ .

### 4.1.5 Routing Control

Routing control is inherently related to the tasks of the demand management component discussed in 4.1.2-4.1.4. Its goal is to optimize the route plan for serving the set of previously received orders augmented by the newly received one and to potentially make additional routing decisions based on expected demand. In contrast to checking feasibility, the objective is to not only determine a feasible route plan but a cost-minimal one. Depending on the control problem at hand, three types of routing control are possible that differ in what portion of the route plan is determined.

**Full route plan**: For disjoint horizon problems, routing control is in fact static and deterministic as definitive routing decisions are made after the booking horizon. Therefore, the provider makes a single decision on the *full route plan* under certainty by solving a static vehicle routing problem (Toth and Vigo, 2014). Note that, additionally, tentative route planning is part of some solution approaches for feasibility checking, cost estimation, and demand control of disjoint problems but we do not categorize it as routing control.

**Single route**: Conversely, for overlapping horizons, some fulfillment planning decisions must be made during the booking horizon and cannot be postponed until its end. Routing control decisions can then be made by repeatedly fixing *complete routes for single vehicles* over time, e.g., when the capacity limit of a vehicle is reached. For this purpose, corresponding routing problems may include tentative decisions for other vehicles. Consequently, most problems consider a tentative route plan beyond the route to be optimized (e.g., Klein and Steinhardt, 2023). This is particularly suitable for deliveries from a central depot as, once a set of orders is loaded onto a vehicle, the route usually cannot be changed any more.

**Leg-oriented**: Overlapping horizons also allow only fixing a certain part of a route, i.e., the next leg or the next few legs for each vehicle. A leg may correspond to serving an order, moving empty to another location or a charging station, or even idling until the next decision epoch. This type of routing control is often applied to point-to-point transportation problems. In this context, ful-fillment planning at each decision epoch only needs to cover a short time span in the case of tight waiting time limits and the absence of pre-bookings (e.g., Kullman et al., 2022). Decisions on relocations and deliberate waiting times of the vehicles, i.e., anticipative routing decisions based on expected demand, can be incorporated, e.g., by means of predictive modeling (e.g., Holler et al., 2019). We refer the interested reader to the works of Berbeglia et al. (2010), Ulmer (2017), Soeffker et al. (2022), and Pillac et al. (2013) for an in-depth consideration of these aspects.

#### 4.2 Static Deterministic Approximation

Integrated demand management and vehicle routing problems can also be cast as static deterministic problems assuming given deterministic customer requests and customer preferences. Only a subset of requests must be accepted as orders. If multiple fulfillment options are defined, it may also be part of the optimization which option should be sold to each customer. Hence, for a fleet of several vehicles, profitable capacitated tour problems or team orienteering problems result (Vansteenwegen and Gunawan, 2019). Therefore, they can be formulated as mixed-integer programs (MIPs). As is the case for dynamic control models, their structure depends on the problem setting. Depending on their use, we distinguish two types of static deterministic approximations for the dynamic control model:

**Offline static control**: Here, we assume perfect information on incoming customer requests and customer preferences. This assumption reflects an ex-post perspective at the end of the booking horizon. Solution approaches based on offline static control auxiliary models yield static controls, which determine definitive demand management decisions before the start of the booking horizon (e.g., Agatz et al., 2011, Klein et al., 2019, and Mackert et al., 2019). Another motivation for explicitly considering such models results from the fact that their solutions serve as a bound for any policy's performance for the corresponding dynamic problem (e.g., Hosni et al., 2014).

**Online static control**: The underlying idea of this approach is to derive both demand management and vehicle routing decisions from a static snapshot of the original dynamic control problem at a specific decision epoch. Consequently, perfect information is only available about existing orders and newly arrived requests. Online static control is applicable for both real-time request processing and batched request processing (e.g., Erdmann et al., 2021). Expected future orders can, e.g., be integrated by simulating customer arrivals or using aggregated expectations, which results in anticipative auxiliary models. Note that in addition, constraints must ensure all previously made decisions. Exemplary formulations of auxiliary models can be found in Klapp et al. (2020), Voccia et al. (2019), and Wang et al. (2021).

### 4.3 Tabular Overview

This section provides an overview of the literature on modeling and dynamically solving integrated demand management and vehicle routing problems that we consider to be in scope for this survey. To this end, we use the morphological analysis of the problem characteristics from Section 2.1 to classify the individual publications (see Table 1 for the possible realizations of all dimensions). Table 4 comprehensively merges the results of this analysis (Columns 3-10) with the application (Column 2), the selected customer choice model (Column 11), and the basic solution concept of the respective work (Columns 12 and 13). Please note that in addition to the applications considered in Section 2.2, we use the entry "GEN" for publications that consider a generic problem setting and do not specify an application. Also, Column 10 sketches the constraint structure of the respective problem in more detail than given in Table 1. The following entries are possible: single vehicle fleet (SV), heterogeneous fleet (HF), multiple trips per vehicle (MT), maximum route duration (RD), order pickup range (PR), physical vehicle capacity (PC), time windows (TW), delivery deadlines (DD), maximum waiting time (WT), maximum ride time (RT), and battery charging level (CL). Since we focus on dynamic decision making, all publications listed in Table 4 propose dynamic controls, and we omit the dimension "time of decision". Column 11 specifies whether the authors apply a multinomial logit model (ML), a generalized attraction model (GA), a finite mixture MNL model (FM), a nested logit model (NL), or a pricingspecific parametric model (PM). To characterize the solution concept, Column 12 states whether the authors apply a decomposition-based approximation ( $\checkmark$ ) or a static deterministic one (X). Additionally, Column 13 indicates whether the approach is anticipative ( $\checkmark$ ) or myopic (X). For the works applying decomposition-based approximation, we summarize the task-specific solution approach that the authors selected in Table 5. We use the classification scheme given in Table 3. In case predictive cost estimation is applied, we additionally state whether it provides a state value estimate (SV), a direct cost estimate (DC), or a state-action value estimate (AV).

Authors	Appli- cation	Service type	Objective	Control type	Pro- cessing	Fulfillment options	Booking/ser- vice horizon	Routing problem	Constraints	Choice model	Decom- position	Antici- pation
Al-Kanj et al. (2020)	MOD	TR	PR	PB	BA	SI	OL	РР	PC, CL, PR	ML	$\checkmark$	$\checkmark$
Alonso-Mora et al. (2017)	MOD	TR	NO	AV	BA	SI	OL	PP	PC, WT, RT	_	Х	$\checkmark$
Angelelli et al. (2021)	GEN	TR	PR	AV	RT	SI	DJ	PU	SV	-	$\checkmark$	$\checkmark$
Archetti et al. (2021)	GEN	TR	PR	AV	RT	SI	OL	DE	PC, HF, TW, MT	-	$\checkmark$	Х
Atasoy et al. (2015)	MOD	TR	PR	AV	RT	MU	OL	PP	PC, HF, WT, RT	ML	$\checkmark$	Х
Avraham and Raviv (2021)	FSO	CS	NO	AV	RT	MU	IF	DE	TW	ML	$\checkmark$	$\checkmark$
Azi et al. (2012)	SDD	TR	PR	AV	RT	SI	OL	DE	TW, MT, RD	_	$\checkmark$	$\checkmark$
Bertsimas et al. (2019)	MOD	TR	PR	AV	BA	SI	OL	PP	PC, TW	-	Х	Х
Campbell and Savelsbergh (2005)	AHD	CG	PR	AV	RT	MU	DJ	DE	PC, TW	_	$\checkmark$	$\checkmark$
Campbell and Savelsbergh (2006)	AHD	CG	PR	PB	RT	MU	DJ	DE	PC, TW	PM	$\checkmark$	Х
Chen et al. (2019)	MOD	TR	RE	PB	RT	SI	OL	PP	PC, PR	PM	$\checkmark$	$\checkmark$
Chen et al. (2023)	SDD	TR	NO	AV	RT	SI	OL	DE	DD, MT	_	$\checkmark$	$\checkmark$
Chen et al. (2022)	SDD	TR	NO	AV	RT	SI	OL	DE	PC, HF, DD, MT	—	$\checkmark$	$\checkmark$
Côté et al. (2021)	SDD	TR	PR	AV	BA	SI	OL	DE	TW, MT	—	Х	$\checkmark$
Dayarian et al. (2020)	SDD	TR	NO	AV	BA	SI	OL	DE	SV, PC, HF, DD	_	Х	Х
Dumouchelle et al. (2021)	GEN	TR	PR	AV	RT	SI	DJ	PU	PC	_	$\checkmark$	$\checkmark$
Erdmann et al. (2021)	MOD	TR	PR	AV	RT, BA	SI	OL	PP	PC, TW, WT	_	Х	Х
Fielbaum et al. (2022)	MOD	TR	PR	AV	BA	SI	OL	PP	PC, WT, RT	_	Х	$\checkmark$
Giallombardo et al. (2022)	GEN	TR	PR	AV	RT	SI	DJ	PU	PC	_	$\checkmark$	$\checkmark$
Haferkamp and Ehmke (2022)	MOD	TR	NO	AV	RT	SI	OL	PP	WT, RT	—	Х	$\checkmark$
Haliem et al. (2021)	MOD	TR	PR	PB	BA	SI	OL	PP	PC	PM	$\checkmark$	$\checkmark$
Holler et al. (2019)	MOD	TR	RE	AV	BA	SI	OL	PP	PC, PR, WT	—	$\checkmark$	$\checkmark$
Hosni et al. (2014)	MOD	TR	PR	AV	RT	SI	OL	PP	PC, HF, WT, RT	—	$\checkmark$	Х
Jahanshahi et al. (2022)	SDD	TR	NO	AV	RT	SI	OL	PP	PC, DD	—	$\checkmark$	$\checkmark$
Klapp et al. (2018)	SDD	TR	PR	AV	BA	SI	OL	DE	SV, MT, RD	_	Х	$\checkmark$
Klapp et al. (2020)	SDD	TR	PR	AV	RT	SI	OL	DE	SV, MT, RD	—	Х	$\checkmark$
Klein et al. (2018)	AHD	CG	PR	PB	RT	MU	DJ	DE	PC, TW	ML	$\checkmark$	$\checkmark$

 Table 4 General overview (continued on next page)

Authors	Appli- cation	Service type	Objective	Control type	Pro- cessing	Fulfillment options	Booking/ser- vice horizon	Routing problem	Constraints	Choice model	Decom- position	Antici- pation
Klein and Steinhardt (2023)	SDD	CG	PR	PB	RT	MU	OL	DE	DD, MT	PM	$\checkmark$	$\checkmark$
Koch and Klein (2020)	AHD	CG	PR	PB	RT	MU	DJ	DE	TW	FL	$\checkmark$	$\checkmark$
Köhler et al. (2019)	AHD	CG	NO	PB	RT	MU	DJ	DE	TW	NL	$\checkmark$	Х
Köhler et al. (2020)	AHD	CG	NO	AV	RT	MU	DJ	DE	TW	_	$\checkmark$	Х
Kullman et al. (2022)	MOD	TR	PR	AV	RT	SI	OL	PP	PC, WT, CL	-	$\checkmark$	$\checkmark$
La Rocca and Cordeau (2019)	MOD	TR	RE	AV	BA	SI	IF	PP	PC, WT, CL	_	Х	Х
Lang et al. (2021a)	AHD	CG	RE	AV	RT	MU	DJ	DE	TW	FL	$\checkmark$	$\checkmark$
Lang et al. (2021b)	AHD	CG	RE	AV	RT	MU	DJ	DE	TW	FL	$\checkmark$	$\checkmark$
Lebedev et al. (2020)	AHD	CG	PR	PB	RT	MU	DJ	DE	TW	ML	$\checkmark$	$\checkmark$
Lebedev et al. (2022)	AHD	CG	PR	PB	RT	MU	DJ	DE	TW	ML	$\checkmark$	$\checkmark$
Lotfi and Abdelghany (2022)	MOD	TR	PR	AV	BA	MU	OL	PP	PC, TW	_	$\checkmark$	Х
Mackert (2019)	AHD	CG	PR	AV	RT	MU	DJ	DE	PC, TW	GA	$\checkmark$	$\checkmark$
Ni et al. (2021)	MOD	TR	PR	PB	BA	MU	OL	PP	CL	PM	$\checkmark$	$\checkmark$
Prokhorchuk et al. (2019)	SDD	TR	RE	PB	RT	MU	OL	DE	DD, MT	ML	$\checkmark$	$\checkmark$
Qiu et al. (2018)	MOD	TR	PR	PB	RT	MU	OL	PP	PC, HF, RT, PR	ML	$\checkmark$	$\checkmark$
Strauss et al. (2021)	AHD	CG	PR	PB	RT	MU	DJ	DE	PC, TW	NL	$\checkmark$	$\checkmark$
Ulmer (2020a)	SDD	CG	RE	PB	RT	MU	OL	DE	DD, MT	PM	$\checkmark$	$\checkmark$
Ulmer (2020b)	FSO	CS	NO	AV	BA	SI	OL, IF	DE	SV	_	$\checkmark$	$\checkmark$
Ulmer et al. (2018)	GEN	TR	NO	AV	BA	SI	OL	PU	SV	-	$\checkmark$	$\checkmark$
Ulmer et al. (2019)	GEN	TR	NO	AV	BA	SI	OL	PU	SV	-	$\checkmark$	$\checkmark$
Ulmer and Thomas (2020)	GEN	TR	RE	AV	RT	SI	DJ	DE	SV, PC	_	$\checkmark$	$\checkmark$
Vinsensius et al. (2020)	AHD	CG	PR	PB	RT	MU	DJ	DE	PC, HF, RD, MT, TW	PM	$\checkmark$	$\checkmark$
Voccia et al. (2019)	SDD	TR	NO	AV	BA	SI	OL	DE	TW, MT	_	Х	$\checkmark$
Wang et al. (2021)	MOD	TR	PR	PB	BA	MU	OL	PP	PC, TW	NL	Х	Х
Xu et al. (2018)	MOD	TR	RE	AV	BA	SI	OL	PP	PC	_	$\checkmark$	$\checkmark$
Yang and Strauss (2017)	AHD	CG	PR	PB	RT	MU	DJ	DE	PC, TW	ML	$\checkmark$	$\checkmark$
Yang et al. (2016)	AHD	CG	PR	PB	RT	MU	DJ	DE	PC, TW	ML	$\checkmark$	$\checkmark$
Zhang et al. (2023)	GEN	TR	NO	AV	RT	SI	OL	DE	MT	_	$\checkmark$	$\checkmark$

Authors	Feasibility check	Cost estimation	Demand control	Routing control
Al-Kanj et al. (2020)	RO	PR, SV	DP	LO
Angelelli et al. (2021)	RO	SA	AR	FP
Archetti et al. (2021)	RO	MY	AR	SR
Atasoy et al. (2015)	RO	MY	AO	SR
Avraham and Raviv (2021)	RO	PR, DC	AO	FP
Azi et al. (2012)	RO	SA	AR	SR
Campbell and Savelsbergh (2005)	RO	SA	AR	FP
Campbell and Savelsbergh (2006)	RO	MY	DP	FP
Chen et al. (2019)	RO	PR, SV	DP	LO
Chen et al. (2023)	RO	PR, AV	AR	SR
Chen et al. (2022)	RO	PR, AV	AR	SR
Dumouchelle et al. (2021)	CA	PR, AV	AR	FP
Giallombardo et al. (2022)	CA	DL	AR	FP
Haliem et al. (2021)	RO	MY	СР	LO
Holler et al. (2019)	RO	PR, AV	AR	LO
Hosni et al. (2014)	RO	MY	AR	LO
Jahanshahi et al. (2022)	RO	PR, AV	AR	LO
Klein et al. (2018)	RO	DL	CP	FP
Klein and Steinhardt (2023)	RO	SA	DP	SR
Koch and Klein (2020)	RO	PR, SV	DP	FP
Köhler et al. (2019)	RO	MY	DP	FP
Köhler et al. (2020)	RO	MY	AO	FP
Kullman et al. (2022)	RO	PR, AV	AR	LO
Lang et al. (2021a)	CA	PR, SV	AO	FP
Lang et al. (2021b)	CA	PR, SV	AO	FP
Lebedev et al. (2020)	CA	PR, SV	СР	FP
Lebedev et al. (2022)	CA	PR, SV	СР	FP
Lotfi and Abdelghany (2022)	RO	MY	AR	SR
Mackert (2019)	RO	DL	AO	FP
Ni et al. (2021)	RO	PR, SV	СР	LO
Prokhorchuk et al. (2019)	RO	PR, SV	СР	SR
Qiu et al. (2018)	RO	PR, DC	СР	LO
Strauss et al. (2021)	CA	DL	DP	FP
Ulmer (2020a)	RO	PR, SV	СР	SR
Ulmer (2020b)	RO	SA, PR, SV	AR	LO
Ulmer et al. (2018)	RO	PR, SV	AR	LO
Ulmer et al. (2019)	RO	SA, PR, SV	AR	LO
Ulmer and Thomas (2020)	RO	PR, SV	AR	FP
Vinsensius et al. (2020)	CA	PR, SV	DP	FP
Xu et al. (2018)	RO	PR, SV	AR	LO
Yang and Strauss (2017)	CA	PR, SV	СР	FP
Yang et al. (2016)	RO	SA	СР	FP
Zhang et al. (2023)	RO	DL	AR	SR

<b>Table</b> 5 Overview of decomposition-based approaches	Table	5 Ov	verview	of deco	mposition	-based	approaches
---	-------	------	---------	---------	-----------	--------	------------

## 5 Solution Algorithms

In this section, we provide a more detailed analysis of the specific algorithms used as part of the solution concepts from Section 4. Hence, this section is intended particularly for readers who would like to dive deeper into the literature. We discuss algorithms for both classes of solution concepts in Sections 5.1 and 5.2, respectively.

## 5.1 Algorithms for Decomposition-based Approximation

In the following, we discuss algorithms for the tasks individually in Sections 5.1.1-5.1.4. We structure our discussion along the types of solution approaches characterized in Section 4.1. An essential observation is that authors rarely fully decompose the problem, i.e., they often propose

a particular algorithm to tackle more than one task. Therefore, at the end of each section, we highlight the algorithms suitable for solving a combination of the current and preceding tasks.

## 5.1.1 Feasibility Check

The complexity of this task ranges from almost trivial (e.g., if the fleet consists of vehicles with a physical capacity of one) to NP-hard for time-window-constrained problems (Savelsbergh, 1985). Consequently, *exact* as well as *heuristic* algorithms are applied. Heuristic algorithms are usually considerably faster compared to exact ones. Thus, as feasibility checks are required simultaneously for all potential options in real-time, the former prevail in the literature. However, they may return false-positive or false-negative results, i.e., incorrectly categorize an option as feasible or infeasible, respectively. The consequence of a false-positive statement and a resulting order of the corresponding option is that the provider cannot serve the respective customer or other customers due to insufficient capacity. This could cause a loss of customer goodwill (e.g., Wang et al., 2011) or require expensive short-term capacity enhancement measures (e.g., Vinsensius et al., 2020). By way of contrast, false-negative statements might lead to lost sales if a feasible and profitable option is not offered.

Algorithms for route-based checks: Most publications apply route-based feasibility checks, drawing on the extensive set of existing methods for solving classical static vehicle routing problems:

- Heuristics: In heuristic algorithms, at least one route plan φ ∈ Φ<sub>k</sub> serving all orders accepted so far is maintained or generated online at each decision epoch. If the heuristic finds that augmenting φ to a plan φ' for an option o is feasible, the check returns a positive result. Most approaches use an insertion heuristic to this end (Solomon, 1987). Insertion heuristics offer high flexibility regarding the extensiveness of the search for a feasible position and are adaptable to many generalizations of the vehicle routing problem (Campbell and Savelsbergh, 2004). For that reason, they are applied to almost any problem setting. The following works present interesting contributions regarding this method: Campbell and Savelsbergh (2005) generate a pool of tentative route plans using a randomized insertion procedure and evaluate all potential insertion positions for a particular fulfillment option. Yang et al. (2016) additionally maintain a tentative route plan from the previous decision epoch. Azi et al. (2012) allow splitting routes if there is no feasible insertion position in the original routes of a single route plan, given some maximum route length constraint. Prokhorchuk et al. (2019) check for infeasible and undoubtedly unprofitable options to reduce the computational effort for the downstream tasks.
- **Exact algorithms**: As opposed to heuristics, exact algorithms thoroughly search a static routing problem's solution space. Thus, they do not return false results but at the cost of higher time consumption. In the surveyed literature, authors only consider total enumeration and apply it to less complex problems. For example, they examine problem settings that only

involve vehicles with a physical capacity of one (e.g., Chen et al., 2019). Qiu et al. (2018) show that total enumeration is also applicable for vehicle capacities in the lower one-digit range.

Algorithms for capacity-based checks: Since capacity-based feasibility checks approximate route-based auxiliary models, they are heuristic by design. The corresponding algorithms differ in how capacity limits are determined offline. Lebedev et al. (2020), Yang and Strauss (2017), and Strauss et al. (2021) draw on routing approximation techniques by Daganzo (1987). Lang et al. (2021a) apply the iterated local search algorithm by Souffriau et al. (2013) to solve sampled instances of the offline static control problem (Section 4.2). Lang et al. (2021b) also solve an offline problem with forecasted orders but assume that the provider must serve all orders.

### 5.1.2 Cost Estimation

As outlined in Section 4.1.3, any cost estimation is approximate due the task's high complexity. This section discusses the algorithms presented in the literature, which again may yield cost estimates of varying quality.

Algorithms for myopic estimation: Applying myopic estimation yields an estimate of *tentative* marginal cost-to-serve, i.e., the increase in total delivery cost caused by additionally serving a potential order o. The term *tentative* expresses that they only refer to the orders accepted so far. For this estimate to be exact, routing costs of the optimal route plans  $\phi'^*$  and  $\phi^*$  with and without the potential order need to be determined, which is often very time-consuming. Therefore, only Hosni et al. (2014) apply a standard mixed-integer programming solver (MIP solver) to search for the minimum-cost update, however, separately for each vehicle and thus heuristically. The other algorithms rely on insertion heuristics.

Campbell and Savelsbergh (2006) approximate the tentative marginal cost-to-serve by the insertion cost of a potential order concerning a pool of tentative route plans. Atasoy et al. (2015) develop a similar procedure that differentiates between different vehicle types but is based only on a single current route plan. Köhler et al. (2019) and Köhler et al. (2020) observe that the insertion cost decreases depending on the routing flexibility for a given set of orders. Hence, they use measures for the routing flexibility as a cost estimate.

It is important to note that even exact tentative marginal cost-to-serve are an approximation of the true marginal cost-to-serve. The latter can be computed at the end of the booking horizon, being the cost difference between optimal route plans with and without the potential order. In the following, we denote this true hindsight cost as *ex-post marginal cost-to-serve*. This distinction is required because a tentative route plan can structurally differ from the final route plan to a large extent (Yang et al., 2016).

Hence, we have a chain of three potential sources of inaccuracy for myopic estimation: First, heuristic algorithms only approximate the exact tentative marginal cost-to-serve. Second, even the exact tentative marginal cost-to-serve only approximate the ex-post marginal cost-to-serve.

Third, as explained in Section 4.1.1, the ex-post marginal cost-to-serve is just one cost component and must be complemented by the exact displacement cost to obtain a perfectly accurate cost estimate. Within the class of myopic approaches, an algorithmic improvement can just tackle the first source of inaccuracy as the other two are of a structural nature.

It is only through anticipation that a *refinement of the marginal cost-to-serve estimate* beyond the exact tentative value and toward the ex-post value and the *estimation of displacement cost* becomes possible. However, not all approaches take advantage of both opportunities, as explained in the following.

Algorithms for sampling-based estimation: Solution algorithms for sampling-based cost estimation are related to those for myopic estimation in that they are also essentially routing heuristics. However, the inclusion of sampled orders necessitates adaptions.

Azi et al. (2012) calculate the average insertion cost of a potential order into a pool of route plans, each initialized with sampled orders. They permanently insert new orders into the sampled route plans and reoptimize them using an adaptive large neighborhood search heuristic. Yang et al. (2016) compute a weighted combination of the average insertion cost regarding two pools of route plans: One contains route plans of all received orders. The other consists of historic final route plans and, hence, entirely contains sampled orders. The tentative insertion cost is expected to gain accuracy throughout the booking horizon, so its weight is gradually increased. Displacement of sampled customers is not possible in either approach.

In contrast, the following three algorithms also estimate displacement cost. Campbell and Savelsbergh (2005) construct a single route plan from scratch using a profit-based insertion heuristic for each potential order. In the first phase, they insert all existing orders. In the second one, they include the potential order together with a set of sampled ones. Thereby, they adjust the sampled orders' revenues by the probabilities of their arrival. The resulting objective function values of the solution with and without the potential order are used to determine a cost estimate, including displacement cost. Angelelli et al. (2021) follow the same ideas but draw on a different routing heuristic (Chao et al., 1996). Klein and Steinhardt (2023) propose a method to refine cost estimates derived from scenario-sampling through the explicit integration of future demand control decisions and the resulting customer choice behavior. Ulmer (2020b) presents a rollout algorithm. It uses a pre-trained state value approximation and an insertion heuristic to simulate demand control and routing control, respectively.

Algorithms for deterministic linear programming: Deterministic linear programming models are usually solved through MIP solvers. To achieve tractable formulations, authors propose several techniques. Such formulations require an approximation of final routing cost based on known and expected orders. Since expected orders depend on future demand management decisions and, potentially, on customer choice behavior, they must also include these aspects.

Klein et al. (2018) solve a model leaning on the choice-based deterministic linear program (e.g., Liu and van Ryzin, 2008). For estimating routing cost, they combine insertion-based tentative route planning with a seed-based routing approximation developed by Fisher and Jaikumar (1981). To account for expected demand management and the resulting purchase decisions, they define a set of potential price lists and pre-compute choice probabilities. Mackert (2019) uses the same routing approach to adapt the sales-based deterministic linear program by Gallego et al. (2015), which endogenizes a choice model in the form of linearized constraints. The same is true for the formulation used by Strauss et al. (2021). However, they apply the approximation developed by Daganzo (1987) to estimate the final routing cost. Zhang et al. (2023) solve a multiple-knapsack problem approximating both future demand management and routing decisions. Giallombardo et al. (2022) geographically aggregate requests to allow for explicit route planning. If the request arrival rate is prohibitively high for real-time decisions, Klein et al. (2018) and Giallombardo et al. (2022) propose solving their auxiliary model at larger time intervals and re-using current cost estimates until an update is available.

Algorithms for predictive estimation: For predictive approaches, algorithms solve the estimation problem of the statistical model, i.e., they train the model based on historical or simulated booking data (e.g., Powell, 2019). This training involves several steps, such as feature value calculation, model updates, and exploration. For each of these steps, a wide range of methods from the field of statistical learning can be applied in various combinations to the control problem considered in this survey. Therefore, we refrain from discussing the individual methods and their composition in detail and only give an overview of the most important contributions.

State value approximations: Lang et al. (2021b) apply a backward dynamic programming algorithm to compute a lookup table. Ulmer et al. (2018) and Al-Kanj et al. (2020) propose dynamically refining the partitioning of lookup tables during the offline learning process. Ulmer et al. (2019) amend this approach by an online rollout component. The parametric models by Prokhorchuk et al. (2019) and Koch and Klein (2020) entail features derived from route plans. Like sampling-based approaches, the latter include sampled orders into the route plan, which they gradually remove during the booking horizon. Both works use linear regression for policy updates. Koch and Klein (2020) find that side constraints incorporating the value function's structural properties improve the learning performance. The algorithms of Yang and Strauss (2017) and Vinsensius et al. (2020) do not require any tentative route planning. Both update the parameters using a stochastic gradient step immediately after each value observation but differ in the way of calculating the final delivery cost: Yang and Strauss (2017) use a routing approximation by Daganzo (1987), Vinsensius et al. (2020) construct each final route plan with a minimum-regret insertion heuristic (Pisinger and Ropke, 2007). For non-linear statistical models, Lebedev et al. (2020) and Lebedev et al. (2022) show that policy updates are not prohibitively complex. The same is true for non-parametric statistical

models, i.e., neural networks, for which special policy update methods exist depending on the model specification (Chen et al., 2019 and Lang et al., 2021a).

- **Direct cost approximations**: To directly learn a non-linear cost function, Avraham and Raviv (2021) conduct an iterative local search within a gradient descend framework and use simulation to evaluate a parameter set's quality. Qiu et al. (2018) employ a covariance matrix adaption evolution strategy, i.e., a numerical optimization method, to learn the parameters of a linear function.
- State-action value approximations: Instead of a value function or a cost function, Q-learning is based on learning a state-action value function. Combining Q-learning with a deep neural network representation of the state-action value function is called Deep Q-learning. It is, e.g., applied in the following two works: Chen et al. (2023) train the network such that it learns a policy which balances acceptance rates over sub-areas. Kullman et al. (2022) estimate a separate Q-value for each vehicle and mimic centralized control during training by a reward-sharing mechanism. Holler et al. (2019) propose a proximal policy optimization method that also relies on a neural network representation of the policy. Jahanshahi et al. (2022) train a Double Deep Q-Network with prioritized experience replay. Finally, Dumouchelle et al. (2021) train a neural network combining Monte Carlo tree search with the SARSA algorithm.

*Combination with other tasks*: All algorithms for myopic cost estimation simultaneously provide a cost estimate and a statement on the feasibility for each potential order.

Yang et al. (2016) and Klein and Steinhardt (2023) simultaneously check feasibility when applying their routing heuristics to determine sampling-based cost estimates. Since some predictive cost estimation algorithms require tentative route planning to calculate feature values, such as the free time budget, combining them with a route-based feasibility check (e.g., Ulmer et al., 2018) is natural.

Integrated capacity-based feasibility checks are, on the one hand, possible via the routing approximations used as part of the deterministic linear programming approach by Strauss et al. (2021) as well as the predictive approaches by Lang et al. (2021a), Lebedev et al. (2020), and Yang and Strauss (2017). On the other hand, the cost estimate can incorporate the likelihood that a potential order leads to an infeasible route plan. If the likelihood is high, the aim is to set the value of the cost estimate sufficiently high to prevent offering the respective fulfillment option. Vinsensius et al. (2020) and Dumouchelle et al. (2021) propose such algorithms.

#### 5.1.3 Demand Control

In this section, we examine algorithms for the demand control task. The complexity of this task depends on both the type of solution approach, according to which we structure the following discussion, and the choice model providing purchase probabilities.

Algorithms for accept/reject control: Accept/reject decisions based on both booking limits and bid prices require minimal computational effort. For booking limits, it is sufficient to check

whether a potential order causes the respective limit to be exceeded (Giallombardo et al., 2022). Controlling demand based on bid prices requires checking whether a potential order's profit is larger than or at least equal to the cost estimate. If not, the request is rejected (Hosni et al., 2014). However, some algorithms allow such requests to be reconsidered in subsequent decision epochs until they expire (Holler et al., 2019). Maximizing state-action values (Kullman et al., 2022) or solving a matching problem via the Kuhn-Munkres algorithm (Xu et al., 2018) are also suitable for accept/reject control.

Algorithms for assortment optimization: Under the assumption of an MNL choice model, an optimal offer set exists among the nested-by-revenue ones (Talluri and van Ryzin, 2004b). Lang et al. (2021a) and Lang et al. (2021b) take advantage of this property, which does no longer hold in case of side constraints. The application considered by Atasoy et al. (2015) requires such constraints to guarantee that at most one option of different classes of fulfillment options is offered. However, the total unimodularity of this constraint allows formulating the problem as a linear program (see Davis et al., 2013 and Bechler et al., 2021 for an overview of such linearization techniques). Similarly, Mackert (2019) uses a linearized formulation of the assortment optimization problem arising under the assumption of a generalized attraction choice model. For problem settings where  $|\Theta_k|$  is low, Avraham and Raviv (2021) find that total enumeration is an efficient method to solve assortment optimization problems given that all options with  $r_{oi} < \Delta \tilde{V}(S_{k+1}|o)$  can be excluded.

Algorithms for pricing: Discrete pricing problems can be modeled as assortment optimization problems, such that algorithms described in the previous paragraph are applicable. Like Atasoy et al. (2015), Strauss et al. (2021) solve an MNL-based pricing problem with unimodular constraints using a MIP solver. The constraints guarantee that less convenient options are priced lower than more convenient ones. Koch and Klein (2020) tackle the discrete pricing problem under a finite-mixture MNL model through a greedy construction heuristic.

Yang et al. (2016) are the first to describe the continuous pricing problem resulting from applying the MNL model in the context of demand management for a vehicle routing application. Drawing on Dong et al. (2009), they show that the problem is non-linear but concave, so they can apply any numerical optimization method.

While all pricing policies discussed so far involve discrete choice models, the literature describes some other variants. Campbell and Savelsbergh (2006) propose a two-step algorithm. First, they perform a rule-based selection of feasible options to be offered at a discount. By solving the piecewise linear approximation of a quadratic program, they determine the value of all discounts. Vinsensius et al. (2020) apply a similar algorithm and solve the resulting quadratic program directly in closed form. Ulmer (2020a) proposes a rule-based policy that makes offers at a static base price or a price equal to the cost estimate if the latter exceeds the base price. Haliem et al. (2021) use a similar method. Köhler et al. (2019) present another rule-based algorithm analogous to the assortment optimization method by Köhler et al. (2020). Al-Kanj et al. (2020) and Chen et

al. (2019) show that machine learning methods are also suitable for solving pricing problems heuristically.

*Combination with other tasks*: As booking limits generally reflect the available logistical capacity, their use for the demand control task involves a capacity-based feasibility check. Concerning the other demand control approaches, existing works exclusively tackle demand control separate from other tasks.

## 5.1.4 Routing Control

Algorithms for determining vehicle routing decisions for control problems with integrated demand management have much in common with pure vehicle routing algorithms (Soeffker et al., 2022). Due to the constraint structure depending on operational restrictions, they are also highly specific to the problem setting of individual applications. As we generally focus on demand management, we only provide a high-level overview.

Algorithms for full route plan approaches: In problem settings with disjoint booking and service horizons, a static vehicle routing problem arises after the cutoff time. Thus, any route planning heuristic suitable for the respective model can be applied.

Algorithms for single route approaches: In the case of overlapping horizons, routing control may rely on fixing complete routes. Here, it is possible to extend the feasibility check to not only search for a feasible update for tentative route planning but a cost-minimal one. Azi et al. (2012) were the first to propose such an algorithm. They insert every new order into a valid route plan containing all received orders and reoptimize it by adaptive large neighborhood search upon each insertion. Archetti et al. (2021) periodically perform a local search, Lotfi and Abdelghany (2022) apply a greedy heuristic. Atasoy et al. (2015) consider a problem setting where each vehicle can be used for different transportation modes. Thus, they divide each route into blocks within which the mode remains the same. If possible, they insert new orders into an existing block. Otherwise, they create a new block solving a shortest path problem.

Algorithms for leg-oriented approaches: Alternatively, the provider can decide on the next legs of vehicles. The methods by Ulmer et al. (2018) and Ulmer et al. (2019) require a decision whether to wait at the current location or to proceed toward the next location according to the updated route plan at each decision epoch. For applications with point-to-point transportation, stand-alone algorithms can determine empty relocations as shown by Chen et al. (2019), who use a random walk process. In contrast, Ni et al. (2021) apply a MIP solver to determine all routing decisions including relocations.

*Combination with other tasks*: For the routing control task, there are many combination opportunities with preceding tasks. Many algorithms for feasibility check and cost estimation already yield route plans as a "side-product." Hence, these plans can be used directly (e.g., Klein and Steinhardt, 2023) or optimized further by the heuristics described above. Xu et al. (2018) and Qiu et al. (2018) show that algorithms for demand control can also yield routing decisions. Decisions on relocations can also be made in conjunction with demand control. State(-action) value-based accept/reject methods offer one way to integrate these tasks (Al-Kanj et al., 2020, Holler et al., 2019, Jahanshahi et al., 2022, and Kullman et al., 2022). Haliem et al. (2021) estimate dedicated state-action values for relocations, which also serve as an input for pricing decisions.

## 5.2 Algorithms for Static Deterministic Approximation

In contrast to solution concepts based on decomposition, which are often inspired by traditional demand management applications, this class of concepts rather originate from pure dynamic vehicle routing (Berbeglia et al., 2010) and, hence, are only suitable in case of overlapping horizons. In each decision epoch, solving an auxiliary online static control model (Section 4.2) simultaneously provides a demand control and routing control decision. This results in another important characteristic compared to decomposition-based approximations: the lack of an explicit cost estimate. However, the notion of myopic and anticipative decision making is transferable since online static control models may also include information on future demand. In the literature, static deterministic approximation is only applied, with one exception, for accept/reject control. Therefore, the complexity of the periodic optimization problem is mainly determined by the vehicle routing component and the use of anticipation. Consequently, we consider algorithms for myopic and anticipative approaches separately in Sections 5.2.1 and 5.2.2.

## 5.2.1 Algorithms for Myopic Approaches

La Rocca and Cordeau (2019) present the only exact solution algorithm within the class of myopic approaches. They apply a MIP solver to a linear assignment problem with dummy vehicles, which leads to a set of new orders with vehicle assignments. The route plan is then complemented with charging and relocation decisions for unassigned vehicles by separate rule-based policies dependent on the current system state.

Other authors rely on heuristics: Erdmann et al. (2021) propose a greedy matching heuristic to determine order-vehicle assignments. Bertsimas et al. (2019) solve an auxiliary network flow model using a MIP solver. They use the solution from the previous decision epoch as a warm start and a backbone algorithm for preprocessing to reduce the computational effort.

The auxiliary bi-level programming model used by Wang et al. (2021) is very complex as it incorporates choice-based pricing control and thus needs to be solved by a specialized heuristic search algorithm. Dayarian et al. (2020) use a two-stage heuristic that first creates a potentially infeasible route plan serving all received orders and potential orders using a large neighborhood search with a worst-removal destroy operator. Second, potential orders are removed following a greedy scheme until reaching feasibility.

#### 5.2.2 Algorithms for Anticipative Approaches

Some works solving anticipative auxiliary models also consider problems that allow a thorough search of the solution space. For vehicle capacities in the lower one-digit range, Alonso-Mora et

al. (2017) show that total enumeration is applicable. They construct a shareability graph, first proposed by Santi et al. (2014), to identify the set of all feasible routes and solve a matching problem to determine which of these to assign to vehicles. To allow for anticipation and relocations, a set of sampled requests with reduced rejection penalty costs is added to the batch of newly arrived ones. Fielbaum et al. (2022) propose two extensions for this algorithm: First, they modify arc costs according to the expected demand at the vehicle's destination. Second, they refine the sampling procedure for future orders through an online method for estimating demand distributions for sampling that does not require historical data. Klapp et al. (2018) and Klapp et al. (2020) consider single-vehicle, multi-trip problems for which the application of a MIP solver is also practical. Both works develop policies based on a-priori plans, which are computed by solving the offline static control problem based on expected customer arrivals associated with rejection penalties. The a-priori plan is then updated at each decision epoch by solving the online static control problem or a simplified version of it. Following the authors mentioned above, Klapp et al. (2020) state that it is beneficial to warm-start the solver with data from the previous decision epoch. However, they also present a metaheuristic tailored to the problem's structure to reduce computation time further.

This leads us to more complex static control problems where metaheuristics are, in fact, the only practical solution approach. Haferkamp and Ehmke (2022) apply a large neighborhood search with three classical removal operators and regret-insertion. Voccia et al. (2019) generate scenarios by sampling future requests and solve a relaxation of the online static control problem for each scenario instance by a variable neighborhood search. They then apply a consensus function (Bent and van Hentenryck, 2004) to the set of resulting scenario plans. This function identifies which part of each idle vehicle's route can accommodate new orders in each scenario plan, selects the best plan, and with it the subset of requests to accept. The chosen plan is then repaired for feasibility by removing potential orders. Also based on scenario-sampling, Côté et al. (2021) first evaluate whether it is beneficial to delay the start of all planned routes. If not, they first ensure that each request is either planned to be served by a vehicle departing in the current decision epoch, a later decision epoch, or is rejected consistently in all scenarios before applying the consensus function. For route planning, they use an adaptive large neighborhood search.

## 6 Conclusion and Research Opportunities

In this survey, we reviewed the methodological advances regarding the integration of demand management and vehicle routing. This research area, whose origins can be situated around the mid-2000s, encompasses a wide range of applications. Therefore, we first developed a generalized definition and a high-level mathematical model of the underlying sequential decision process, and then used this as a basis for analyzing and classifying the literature concerning the decision problems, solution concepts, and algorithms presented. Based on this analysis, we can now discuss important insights and challenges from a cross-application perspective. In particular, we draw conclusions regarding the current state of research and, simultaneously, point toward future research opportunities. We structure the elaboration along the following seven topics:

Generic model formulations: Establishing some form of a common modeling language is undoubtedly beneficial to describe problem settings in a standardized and concise manner and to be able to relate these settings to each other on a formal level. To this end, it seems most natural to formalize the various settings in terms of corresponding Markov decision processes to fully capture the dynamic and stochastic nature of the underlying control problems. Since many existing works already include such models, we advocate that these become a standard for future publications and introduced a generic, high-level formulation representing a possible starting point for modeling specific control problems in any area of application. One example in this context is the model by Yang et al. (2016) for dynamic pricing in AHD, on which several authors have based their models afterward. A particular challenge arises because vehicle routing dynamics and the reactions of customers to demand management must be modeled. Klein et al. (2020) discuss examples of modeling integrations of demand management techniques and operational decision making from different fields of applications. An important step to improve the presentation of relevant control problems toward a more generic description is to establish and use common terminology that this review aims to contribute to.

Generic solution frameworks: Just like standardized modeling, a uniform description of solution concepts enables methodological transfers within and between the application-specific literature streams and thus a faster progress of research overall. We aimed to contribute toward such a unification by explicitly distinguishing decomposition-based approximations and static deterministic ones as well as the associated solution approaches (Section 4). We encourage authors of future works on decomposition-based approximation to be explicit about how they address each task, how they orchestrate their complete solution method, and how it could possibly be adapted to other problem settings. It is also promising to align solution approaches and model formulation more closely. Substantial efforts in this direction already exist in related fields, e.g., by Ulmer et al. (2020) introducing a route-based Markov decision process for dynamic vehicle routing problems.

Advancement of solution approaches: We also see opportunities for future research at the methodological level. For the feasibility check, machine learning methods suitable for solving binary classification problems could be a valuable extension of the existing body of methods for capacity-based checks. Recent work by Dumouchelle et al. (2021) and Van der Hagen et al. (2024) shows that this is a promising research avenue. The same observation accounts for constraint programming techniques for route-based feasibility checks. Recent advances in approximate dynamic programming could improve cost estimates (Ulmer et al., 2019). To derive more accurate features from route plans, the inclusion of sampled orders could be further investigated (Koch and Klein, 2020). The application of more accurate choice models, whose major drawback is that they cause an increase in complexity of the demand control task, could be facilitated by developing tailored assortment planning and pricing heuristics. Sampling methods that rely on online demand data could enable anticipation in the absence of a reliable source of historical data (Fielbaum et al., 2022).

**Performance assessment:** Due to the abovementioned heterogeneity of the problem settings and dependencies on instance characteristics, comparing the performance of complete solution approaches on a general level is difficult. However, there seems to be a universally valid insight repeatedly reported in different areas of application: Anticipative approaches consistently dominate myopic ones, particularly in problem instances characterized by a medium scarcity of fulfillment capacity (e.g., Azi et al., 2012 and Voccia et al., 2019). Especially the anticipation of displacement effects is found to have a significant impact by several authors (e.g., Klein et al., 2018) comparing their approaches with the method by Yang et al. (2016), which uses anticipation only to refine the estimate of marginal cost-to-serve. Another interesting finding is that anticipation reduces the systematic discrimination against customers based on their location (e.g., Prokhorchuk et al., 2019), an issue that Soeffker et al. (2017) raised first. We believe that researchers should put more emphasis on identifying the components of the overall solution procedure to which a certain increase in performance can be attributed. To a certain extent, this is examined, e.g., in the study by Haferkamp and Ehmke (2022). Regarding the performance, authors should also evaluate the robustness of anticipative approaches in case that the parameters used in choice models and demand distributions differ from the real-world (Srour et al., 2018). To allow a generalized empirical validation of these performance insights, Lang et al. (2021a) identify the development of a benchmarking tool as an essential task for future research. First promising steps in this direction are being taken (Bertsimas et al., 2019 and Lang and Cleophas, 2020).

**Suitability of demand control policies**: Whether providers should prefer availability control or price-based control policies cannot be answered equally clearly, which is why no approach has become dominant in the literature either. Several authors argue that persuasive control strategies, i.e., those using incentives, are superior to coercive ones restricting service availability because they are more likely to be endorsed by customers. Consequently, availability control and especially policies that might reject customers without offering alternative fulfillment options are seen critically (e.g., Asdemir et al., 2009). As Lee and Savelsbergh (2015) point out, the resulting dissatisfaction in MOD settings is amplified by the fact that rejected customers might have to switch to an alternative means of transportation at short notice. Offering a set of fulfillment options instead of only a single one bears the potential to reduce the rate of these provider-side rejections substantially. On the other hand, charging dynamic prices for a logistical service is an inherent competitive disadvantage (Lang et al., 2021b). It may even be restricted or forbidden due to regulation (Bruck et al., 2018). Other types of incentives can use discounts or vouchers (Agatz

et al., 2008) or highlight the environmental benefits of specific fulfillment options (Agatz et al., 2021a) to alleviate these issues.

Advancement of choice modeling and fulfillment options: In both availability control and price-based control, the path toward more customer-friendly controls leads to the growing importance of choice modeling and the design of fulfillment options. As illustrated in Section 3.2, accurately modeling customer choice behavior is widely recognized as a success factor for demand management in general. The results of Mackert et al. (2019) show that this is also the case for vehicle routing applications. In future research, instead of passively fitting choice models, choice behavior could be actively explored, especially if the available historical data are sparse or biased due to suboptimal demand management in the past (Bondoux et al., 2020). Furthermore, the integration of more advanced choice models like the exponomial (Alptekinoğlu and Semple, 2016) or the Markov chain model (Feldman and Topaloglu, 2017) is a promising topic for future research. Likewise, the development of suitable types of fulfillment options should depend on the application examined. Although this is a strategic planning task (Talluri and van Ryzin, 2004a), it often has methodological implications. We believe the potential for future research in this regard exists in all application areas. For instance, Strauss et al. (2021) apply the concept of flexible products in AHD. Atasoy et al. (2015) propose an MOD system that allows customers to choose the mode of transport. Avraham and Raviv (2021) suggest offering arriving customers time slots of several consecutive working days simultaneously.

**Transfer into practice**: More research also seems necessary, in our view, to address problems that arise when transferring existing methods into practice. These include concurrency issues (Avraham and Raviv, 2021) as well as the management of disruptions and failed fulfillments, which can be investigated, for example, by taking stochastic travel times into account (Prokhorchuk et al., 2019). Another issue lies in the scalability of solution approaches concerning large instances, as they usually occur in practice (Bertsimas et al., 2019).

With the survey at hand, we hope to promote the transfer of the large body of existing approaches to novel problem settings or even new applications. Interestingly, all three themes that Agatz et al. (2021b) identify as characteristics of impactful research in the field of transportation are present in the surveyed research area: multi-objective optimization, stochastic optimization, and the integration of stakeholder behavior. Therefore, we believe active demand management to be a key enabler of new, sustainable business models for smart mobility and transportation applications.

## Acknowledgment

We sincerely thank Arne Strauss for his valuable suggestions and the many fruitful discussions on an earlier version of this survey.

## References

- Afèche, P., Liu, Z., & Maglaras, C. (2023). Ride-hailing networks with strategic drivers: The impact of platform control capabilities on performance. *Manufacturing & Service Operations Management*, 25(5), 1890-1908. https://doi.org/10.1287/msom.2023.1221.
- Agatz, N. A. H., Campbell, A. M., Fleischmann, M., & Savelsbergh, M. W. P. (2011). Time slot management in attended home delivery. *Transportation Science*, 45(3), 435-449. https://doi.org/10.1287/trsc.1100.0346.
- Agatz, N. A. H., Campbell, A. M., Fleischmann, M., van Nunen, J., & Savelsbergh, M. W. P. (2013). Revenue management opportunities for internet retailers. *Journal of Revenue and Pricing Management*, 12(2), 128-138. https://doi.org/10.1057/rpm.2012.51.
- Agatz, N. A. H., Fan, Y., & Stam, D. (2021a). The impact of green labels on time slot choice and operational sustainability. *Production and Operations Management*, 30(7), 2285-2303. https://doi.org/10.1111/poms.13368.
- Agatz, N. A. H., Fleischmann, M., & Van Nunen, J. A. (2008). E-fulfillment and multi-channel distribution – A review. *European Journal of Operational Research*, 187(2), 339-356. https://doi.org/10.1016/j.ejor.2007.04.024.
- Agatz, N. A. H., Hewitt, M., & Thomas, B. W. (2021b). "Make no little plans": Impactful research to solve the next generation of transportation problems. *Networks*, 77(2), 269-286. https://doi.org/10.1002/net.22002.
- Al-Kanj, L., Nascimento, J., & Powell, W. B. (2020). Approximate dynamic programming for planning a ride-hailing system using autonomous fleets of electric vehicles. *European Jour*nal of Operational Research, 284(3), 1088-1106. https://doi.org/10.1016/j.ejor.2020.01.033.
- Alonso-Mora, J., Wallar, A., & Rus, D. (2017). Predictive routing for autonomous mobility-ondemand systems with ride-sharing. In 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems, 3583-3590. https://doi.org/10.1109/IROS.2017.8206203.
- Alptekinoğlu, A., & Semple, J. H. (2016). The exponomial choice model: A new alternative for assortment and price optimization. *Operations Research*, 64(1), 79-93. https://doi.org/10.1287/opre.2015.1459.
- Angelelli, E., Archetti, C., Filippi, C., & Vindigni, M. (2021). A dynamic and probabilistic orienteering problem. *Computers & Operations Research*, 136, 105454. https://doi.org/10.1016/j.cor.2021.105454.
- Archetti, C., & Bertazzi, L. (2021). Recent challenges in routing and inventory routing: E-commerce and last-mile delivery. *Networks*, 77(2), 255-268. https://doi.org/10.1002/net.21995.
- Archetti, C., Guerriero, F., & Macrina, G. (2021). The online vehicle routing problem with occasional drivers. *Computers & Operations Research*, 127, 105144. https://doi.org/10.1016/j.cor.2020.105144.
- Asdemir, K., Jacob, V. S., & Krishnan, R. (2009). Dynamic pricing of multiple home delivery options. *European Journal of Operational Research*, 196(1), 246-257. https://doi.org/10.1016/j.ejor.2008.03.005.
- Atasoy, B., Ikeda, T., Song, X., & Ben-Akiva, M. E. (2015). The concept and impact analysis of a flexible mobility on demand system. *Transportation Research Part C: Emerging Technol*ogies, 56, 373-392. https://doi.org/10.1016/j.trc.2015.04.009.
- Avraham, E., & Raviv, T. (2021). The steady-state mobile personnel booking problem. Transportation Research Part B: Methodological, 154, 266-288. https://doi.org/10.1016/j.trb.2021.10.008.
- Azi, N., Gendreau, M., & Potvin, J.-Y. (2012). A dynamic vehicle routing problem with multiple delivery routes. *Annals of Operations Research*, 199(1), 103-112. https://doi.org/10.1007/s10479-011-0991-3.
- Banerjee, S., Johari, R., & Riquelme, C. (2016). Dynamic pricing in ridesharing platforms. ACM SIGecom Exchanges, 15(1), 65-70. https://doi.org/10.1145/2994501.2994505.

- Bechler, G., Steinhardt, C., & Mackert, J. (2021). On the linear integration of attraction choice models in business optimization problems. SN Operations Research Forum, 2(1), 1-13. https://doi.org/10.1007/s43069-021-00056-1.
- Bent, R. W., & van Hentenryck, P. (2004). Scenario-based planning for partially dynamic vehicle routing with stochastic customers. *Operations Research*, 52(6), 977-987. https://doi.org/10.1287/opre.1040.0124.
- Berbeglia, G., Cordeau, J. F., & Laporte, G. (2010). Dynamic pickup and delivery problems. *European Journal of Operational Research*, 202(1), 8-15. https://doi.org/10.1016/j.ejor.2009.04.024.
- Berbeglia, G., Garassino, A., & Vulcano, G. (2022). A comparative empirical study of discrete choice models in retail operations. *Management Science*, 68(6), 4005-4023. https://doi.org/10.1287/mnsc.2021.4069.
- Berbeglia, G., Pesant, G., & Rousseau, L.-M. (2011). Checking the feasibility of dial-a-ride instances using constraint programming. *Transportation Science*, 45(3), 399-412. https://doi.org/10.1287/trsc.1100.0336.
- Bertsimas, D., Jaillet, P., & Martin, S. (2019). Online vehicle routing: The edge of optimization in large-scale applications. *Operations Research*, 67(1), 143-162. https://doi.org/10.1287/opre.2018.1763.
- Bondoux, N., Nguyen, A. Q., Fiig, T., & Acuna-Agost, R. (2020). Reinforcement learning applied to airline revenue management. *Journal of Revenue and Pricing Management*, 19(5), 1-17. https://doi.org/10.1057/s41272-020-00228-4.
- Boysen, N., Fedtke, S., & Schwerdfeger, S. (2021). Last-mile delivery concepts: A survey from an operational research perspective. *OR Spectrum*, 43(1), 1-58. https://doi.org/10.1007/s00291-020-00607-8.
- Brailsford, S. C., Potts, C. N., & Smith, B. M. (1999). Constraint satisfaction problems: Algorithms and applications. *European Journal of Operational Research*, 119(3), 557-581. https://doi.org/10.1016/S0377-2217(98)00364-6.
- Bruck, B. P., Cordeau, J.-F., & Iori, M. (2018). A practical time slot management and routing problem for attended home services. *Omega*, 81, 208-219. https://doi.org/10.1016/j.omega.2017.11.003.
- Campbell, A. M., & Savelsbergh, M. W. P. (2004). Efficient insertion heuristics for vehicle routing and scheduling problems. *Transportation Science*, 38(3), 369-378. https://doi.org/10.1287/trsc.1030.0046.
- Campbell, A. M., & Savelsbergh, M. W. P. (2005). Decision support for consumer direct grocery initiatives. *Transportation Science*, *39*(3), 313-327. https://doi.org/10.1287/trsc.1040.0105.
- Campbell, A. M., & Savelsbergh, M. W. P. (2006). Incentive schemes for attended home delivery services. *Transportation Science*, *40*(3), 327-341. https://doi.org/10.1287/trsc.1050.0136.
- Chao, I. M., Golden, B. L., & Wasil, E. A. (1996). A fast and effective heuristic for the orienteering problem. *European Journal of Operational Research*, 88(3), 475-489. https://doi.org/10.1016/0377-2217(95)00035-6.
- Chen, H., Jiao, Y., Qin, Z., Tang, X., Li, H., An, B., Hongtu, Z., & Ye, J. (2019). InBEDE: Integrating contextual bandit with TD learning for joint pricing and dispatch of ride-hailing platforms. In 2019 IEEE International Conference on Data Mining (ICDM), 61-70. https://doi.org/10.1109/ICDM.2019.00016.
- Chen, X., Thomas, B. W., & Hewitt, M. (2016). The technician routing problem with experiencebased service times. *Omega*, *61*, 49-61. https://doi.org/10.1016/j.omega.2015.07.006.
- Chen, X., Ulmer, M. W., & Thomas, B. W. (2022). Deep Q-learning for same-day delivery with vehicles and drones. *European Journal of Operational Research*, 298(3), 939-952. https://doi.org/10.1016/j.ejor.2021.06.021.

- Chen, X., Wang, T., Thomas, B. W., & Ulmer, M. W. (2023). Same-day delivery with fair customer service. *European Journal of Operational Research*, 308(2), 738-751. https://doi.org/10.1016/j.ejor.2022.12.009.
- Church, G. (2019). The maths problem that could bring the world to a halt. *BBC*. Retrieved from https://www.bbc.com/future/article/20190606-the-maths-problem-that-modern-life-depends-on. Last accessed: 22/12/2024.
- Côté, J. F., de Queiroz, T. A., Gallesi, F., & Iori, M. (2021). Dynamic optimization algorithms for same-day delivery problems. Working paper, Université Laval.
- Daganzo, C. F. (1987). Modeling distribution problems with time windows: Part I. *Transportation Science*, 21(3), 171-179. https://doi.org/10.1287/trsc.21.3.171.
- Davis, J., Gallego, G., & Topaloglu, H. (2013). Assortment planning under the multinomial logit model with totally unimodular constraint structures. Working paper, Cornell University.
- Dayarian, I., Savelsbergh, M. W. P., & Clarke, J.-P. (2020). Same-day delivery with drone resupply. *Transportation Science*, *54*(1), 229-249. https://doi.org/10.1287/trsc.2019.0944.
- Dong, L., Kouvelis, P., & Tian, Z. (2009). Dynamic pricing and inventory control of substitute products. *Manufacturing & Service Operations Management*, 11(2), 317-339. https://doi.org/10.1287/msom.1080.0221.
- Dumouchelle, J., Frejinger, E., & Lodi, A. (2021). Can machine learning help in solving cargo capacity management booking control problems?. Working paper, Polytechnique Montréal.
- Elting, S., & Ehmke, J. F. (2021). Potential of shared taxi services in rural areas A case study. *Transportation Research Procedia*, *52*, 661-668. https://doi.org/10.1016/j.trpro.2021.01.079.
- Erdmann, M., Dandl, F., & Bogenberger, K. (2021). Combining immediate customer responses and car-passenger reassignments in on-demand mobility services. *Transportation Research Part C: Emerging Technologies*, *126*, 103104. https://doi.org/10.1016/j.trc.2021.103104.
- Feldman, J. B., & Topaloglu, H. (2017). Revenue management under the Markov chain choice model. *Operations Research*, 65(5), 1322-1342. https://doi.org/10.1287/opre.2017.1628.
- Fielbaum, A., Kronmueller, M., & Alonso-Mora, J. (2022). Anticipatory routing methods for an on-demand ridepooling mobility system. *Transportation*, 49, 1921-1962. https://doi.org/10.1007/s11116-021-10232-1.
- Fisher, M. L., & Jaikumar, R. (1981). A generalized assignment heuristic for vehicle routing. *Networks*, 11(2), 109-124. https://doi.org/10.1002/net.3230110205.
- Gallego, G., Ratliff, R., & Shebalov, S. (2015). A general attraction model and sales-based linear program for network revenue management under customer choice. *Operations Research*, 63(1), 212-232. https://doi.org/10.1287/opre.2014.1328.
- Gallego, G., & Topaloglu, H. (2019). Revenue management and pricing analytics. New York, NY: Springer. https://doi.org/10.1007/978-1-4939-9606-3.
- Ghiani, G., Manni, E., Quaranta, A., & Triki, C. (2009). Anticipatory algorithms for same-day courier dispatching. *Transportation Research Part E: Logistics and Transportation Review*, 45(1), 96-106. https://doi.org/10.1016/j.tre.2008.08.003.
- Giallombardo, G., Guerriero, F., & Miglionico, G. (2022). Profit maximization via capacity control for distribution logistics problems. *Computers & Industrial Engineering*, 171, 108466. https://doi.org/10.1016/j.cie.2022.108466.
- Haferkamp, J., & Ehmke, J. F. (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.
- Haliem, M., Mani, G., Aggarwal, V., & Bhargava, B. (2021). A distributed model-free ride-sharing approach for joint matching, pricing, and dispatching using deep reinforcement learning. *IEEE Transactions on Intelligent Transportation Systems*, 22(12), 7931-7942. https://doi.org/10.1109/TITS.2021.3096537.

- Hazan, J., Lang, N., Wegscheider, A., & Fassenot, B. (2019). On-demand transit can unlock urban mobility. *Boston Consulting Group*. Retrieved from https://www.bcg.com/de-de/publications/2019/on-demand-transit-can-unlock-urban-mobility. Last accessed: 21/03/2022.
- Holler, J., Vuorio, R., Qin, Z., Tang, X., Jiao, Y., Jin, T., Singh, S., Wang, C., & Ye, J. (2019).
  Deep reinforcement learning for multi-driver vehicle dispatching and repositioning problem.
  In 2019 IEEE International Conference on Data Mining (ICDM), 1090-1095.
  https://doi.org/10.1109/ICDM.2019.00129.
- Hosni, H., Naoum-Sawaya, J., & Artail, H. (2014). The shared-taxi problem: Formulation and solution methods. *Transportation Research Part B: Methodological*, 70, 303-318. https://doi.org/10.1016/j.trb.2014.09.011.
- Ichoua, S., Gendreau, M., & Potvin, J.-Y. (2006). Exploiting knowledge about future demands for real-time vehicle dispatching. *Transportation Science*, 40(2), 211-225. https://doi.org/10.1287/trsc.1050.0114.
- Jahanshahi, H., Bozanta, A., Cevik, M., Kavuk, E. M., Tosun, A., Sonuc, S. B., Kosucu, B., & Başar, A. (2022). A deep reinforcement learning approach for the meal delivery problem. *Knowledge-Based Systems*, 243, 108489. https://doi.org/10.1016/j.knosys.2022.108489.
- Klapp, M. A., Erera, A. L., & Toriello, A. (2018). The dynamic dispatch waves problem for sameday delivery. *European Journal of Operational Research*, 271(2), 519-534. https://doi.org/10.1016/j.ejor.2018.05.032.
- Klapp, M. A., Erera, A. L., & Toriello, A. (2020). Request acceptance in same-day delivery. *Transportation Research Part E: Logistics and Transportation Review*, 143, 102083. https://doi.org/10.1016/j.tre.2020.102083.
- Klein, R., Koch, S., Steinhardt, C., & Strauss, A. K. (2020). A review of revenue management: Recent generalizations and advances in industry applications. *European Journal of Operational Research*, 284(2), 397-412. https://doi.org/10.1016/j.ejor.2019.06.034.
- Klein, R., Mackert, J., Neugebauer, M., & Steinhardt, C. (2018). A model-based approximation of opportunity cost for dynamic pricing in attended home delivery. *OR Spectrum*, 40(4), 969-996. https://doi.org/10.1007/s00291-017-0501-3.
- Klein, R., Neugebauer, M., Ratkovitch, D., & Steinhardt, C. (2019). Differentiated time slot pricing under routing considerations in attended home delivery. *Transportation Science*, 53(1), 236-255. https://doi.org/10.1287/trsc.2017.0738.
- Klein, V., & Steinhardt, C. (2023). Dynamic demand management and online tour planning for same-day delivery. *European Journal of Operational Research*, 307(2), 860-886. https://doi.org/10.1016/j.ejor.2022.09.011.
- Koch, S., & Klein, R. (2020). Route-based approximate dynamic programming for dynamic pricing in attended home delivery. *European Journal of Operational Research*, 287(2), 633-652. https://doi.org/10.1016/j.ejor.2020.04.002.
- Köhler, C., Ehmke, J. F., & Campbell, A. M. (2020). Flexible time window management for attended home deliveries. *Omega*, *91*, 102023. https://doi.org/10.1016/j.omega.2019.01.001.
- Köhler, C., Ehmke, J. F., Campbell, A. M., & Cleophas, C. (2019). Flexible dynamic time window pricing for attended home deliveries. Working paper, University of Magdeburg. https://doi.org/10.24352/UB.OVGU-2019-087.
- Kök, A. G., & Fisher, M. L. (2007). Demand estimation and assortment optimization under substitution: Methodology and application. *Operations Research*, 55(6), 1001-1021. https://doi.org/10.1287/opre.1070.0409.
- Kullman, N. D., Cousineau, M., Goodson, J. C., & Mendoza, J. E. (2022). Dynamic ride-hailing with electric vehicles. *Transportation Science*, 56(3), 775-794. https://doi.org/10.1287/trsc.2021.1042.
- Lang, M. A. K., & Cleophas, C. (2020). Establishing an extendable benchmarking framework for E-fulfillment. In *Proceedings of the 53rd Hawaii International Conference on System Sciences*, 1589-1598. https://doi.org/10.24251/HICSS.2020.195.

- Lang, M. A. K., Cleophas, C., & Ehmke, J. F. (2021a). Anticipative dynamic slotting for attended home deliveries. *Operations Research Forum*, 2(4), 70. https://doi.org/10.1007/s43069-021-00086-9.
- Lang, M. A. K., Cleophas, C., & Ehmke, J. F. (2021b). Multi-criteria decision making in dynamic slotting for attended home deliveries. *Omega*, 102, 102305. https://doi.org/10.1016/j.omega.2020.102305.
- La Rocca, C. R., & Cordeau, J.-F. (2019). Heuristics for electric taxi fleet management at Teo Taxi. *Information Systems and Operational Research*, 57(4), 642-666. https://doi.org/10.1080/03155986.2019.1607808.
- Lebedev, D., Margellos, K., & Goulart, P. (2020). Approximate dynamic programming for delivery time slot pricing: A sensitivity analysis. Working paper, University of Oxford. https://doi.org/10.48550/arXiv.2008.00780.
- Lebedev, D., Margellos, K., & Goulart, P. (2022). Convexity and feedback in approximate dynamic programming for delivery time slot pricing. *IEEE Transactions on Control Systems Technology*, 30(2), 893-900. https://doi.org/10.1109/TCST.2021.3093648.
- Lee, A., & Savelsbergh, M. W. P. (2015). Dynamic ridesharing: Is there a role for dedicated drivers?. *Transportation Research Part B: Methodological*, 81(2), 483-497. https://doi.org/10.1016/j.trb.2015.02.013.
- Liu, Q., & van Ryzin, G. (2008). On the choice-based linear programming model for network revenue management. *Manufacturing & Service Operations Management*, 10(2), 288-310. https://doi.org/10.1287/msom.1070.0169.
- Lotfi, S., & Abdelghany, K. (2022). Ride matching and vehicle routing for on-demand mobility services. *Journal of Heuristics*, 28(3), 235-258. https://doi.org/10.1007/s10732-022-09491-7.
- Mackert, J. (2019). Choice-based dynamic time slot management in attended home delivery. *Computers* & *Industrial Engineering*, *129*, 333-345. https://doi.org/10.1016/j.cie.2019.01.048.
- Mackert, J., Steinhardt, C., & Klein, R. (2019). Integrating customer choice in differentiated slotting for last-mile logistics. *Logistics Research*, *12*(1), 5. https://doi.org/10.23773/2019\_5.
- Ni, L., Sun, B., Wang, S., & Tsang, D. H. (2021). Dynamic pricing mechanism design for electric mobility-on-demand systems. *IEEE Transactions on Intelligent Transportation Systems*, 23(8), 11361-11375. https://doi.org/10.1109/TITS.2021.3103199.
- Pillac, V., Gendreau, M., Guéret, C., & Medaglia, A. L. (2013). A review of dynamic vehicle routing problems. *European Journal of Operational Research*, 225(1), 1-11. https://doi.org/10.1016/j.ejor.2012.08.015.
- Pisinger, D., & Ropke, S. (2007). A general heuristic for vehicle routing problems. *Computers & Operations Research*, 34(8), 2403-2435. https://doi.org/10.1016/j.cor.2005.09.012.
- Poggi, N., Carrera, D., Gavaldà, R., Ayguadé, E., & Torres, J. (2014). A methodology for the evaluation of high response time on E-commerce users and sales. *Information Systems Frontiers*, 16(5), 867-885. https://doi.org/10.1007/s10796-012-9387-4.
- Powell, W. B. (2011). Approximate dynamic programming: Solving the curses of dimensionality. 2nd ed. Hoboken, NJ: John Wiley & Sons. https://doi.org/10.1002/9781118029176.
- Powell, W. B. (2019). A unified framework for stochastic optimization. *European Journal of Operational Research*, 275(3), 795-821. https://doi.org/10.1016/j.ejor.2018.07.014.
- Prokhorchuk, A., Dauwels, J., & Jaillet, P. (2019). Stochastic dynamic pricing for same-day delivery routing. Working paper, Nanyang Technological University.
- Psaraftis, H. N., Wen, M., & Kontovas, C. A. (2016). Dynamic vehicle routing problems: Three decades and counting. *Networks*, 67(1), 3-31. https://doi.org/10.1002/net.21628.
- Qin, Z., Tang, X., Jiao, Y., Zhang, F., Xu, Z., Zhu, H., & Ye, J. (2020). Ride-hailing order dispatching at DiDi via reinforcement learning. *INFORMS Journal on Applied Analytics*, 50(5), 272-286. https://doi.org/10.1287/inte.2020.1047.

- Qiu, H., Li, R., & Zhao, J. (2018). Dynamic pricing in shared mobility on demand service. Working paper, Massachusetts Institute of Technology.
- Santi, P., Resta, G., Szell, M., Sobolevsky, S., Strogatz, S. H., & Ratti, C. (2014). Quantifying the benefits of vehicle pooling with shareability networks. *Proceedings of the National Academy* of Sciences of the United States of America, 111(37), 13290-13294. https://doi.org/10.1073/pnas.1403657111.
- Savelsbergh, M. W. P. (1985). Local search in routing problems with time windows. *Annals of Operations Research*, *4*, 285-305. https://doi.org/10.1007/BF02022044.
- Snoeck, A., Merchán, D., & Winkenbach, M. (2020). Revenue management in last-mile delivery: State-of-the-art and future research directions. *Transportation Research Procedia*, 46, 109-116. https://doi.org/10.1016/j.trpro.2020.03.170.
- Soeffker, N., Ulmer, M. W., & Mattfeld, D. C. (2017). On fairness aspects of customer acceptance mechanisms in dynamic vehicle routing. *Proceedings of Logistikmanagement*, 17-24.
- Soeffker, N., Ulmer, M. W., & Mattfeld, D. C. (2022). Stochastic dynamic vehicle routing in the light of prescriptive analytics: A review. *European Journal of Operational Research*, 298(3), 801-820. https://doi.org/10.1016/j.ejor.2021.07.014.
- Solomon, M. M. (1987). Algorithms for the vehicle routing and scheduling problems with time window constraints. *Operations Research*, 35(2), 254-265. https://doi.org/10.1287/opre.35.2.254.
- Souffriau, W., Vansteenwegen, P., Vanden Berghe, G., & Van Oudheusden, D. (2013). The multiconstraint team orienteering problem with multiple time windows. *Transportation Science*, 47(1), 53-63. https://doi.org/10.1287/trsc.1110.0377.
- Srour, F. J., Agatz, N., & Oppen, J. (2018). Strategies for handling temporal uncertainty in pickup and delivery problems with time windows. *Transportation Science*, 52(1), 3-19. https://doi.org/10.1287/trsc.2015.0658.
- Strauss, A. K., Gülpinar, N., & Zheng, Y. (2021). Dynamic pricing of flexible time slots for attended home delivery. *European Journal of Operational Research*, 294(3), 1022-1041. https://doi.org/10.1016/j.ejor.2020.03.007.
- Strauss, A. K., Klein, R., & Steinhardt, C. (2018). A review of choice-based revenue management: Theory and methods. *European Journal of Operational Research*, 271(2), 375-387. https://doi.org/10.1016/j.ejor.2018.01.011.
- Talluri, K. T., & Van Ryzin, G. J. (2004a). The theory and practice of revenue management. New York, NY: Springer. https://doi.org/10.1007/b139000.
- Talluri, K. T., & van Ryzin, G. (2004b). Revenue management under a general discrete choice model of consumer behavior. *Management Science*, 50(1), 15-33. https://doi.org/10.1287/mnsc.1030.0147.
- Taylor, T. A. (2018). On-demand service platforms. *Manufacturing & Service Operations Management*, 20(4), 704-720. https://doi.org/10.1287/msom.2017.0678.
- Toth, P., & Vigo, D. (Eds.). (2014). Vehicle routing: Problems, methods, and applications. Philadelphia, PA: Society for Industrial and Applied Mathematics. https://doi.org/10.1137/1.9781611973594.
- Train, K. E. (2009). Discrete choice methods with simulation. Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9780511805271.
- Ulmer, M. W. (2017). Approximate dynamic programming for dynamic vehicle routing. Cham: Springer. https://doi.org/10.1007/978-3-319-55511-9.
- Ulmer, M. W. (2020a). Dynamic pricing and routing for same-day delivery. *Transportation Science*, 54(4), 1016-1033. https://doi.org/10.1287/trsc.2019.0958.
- Ulmer, M. W. (2020b). Horizontal combinations of online and offline approximate dynamic programming for stochastic dynamic vehicle routing. *Central European Journal of Operations Research*, 28(1), 279-308. https://doi.org/10.1007/s10100-018-0588-x.

- Ulmer, M. W., Goodson, J. C., Mattfeld, D. C., & Hennig, M. (2019). Offline-online approximate dynamic programming for dynamic vehicle routing with stochastic requests. *Transportation Science*, 53(1), 185-202. https://doi.org/10.1287/trsc.2017.0767.
- Ulmer, M. W., Goodson, J. C., Mattfeld, D. C., & Thomas, B. W. (2020). On modeling stochastic dynamic vehicle routing problems. *EURO Journal on Transportation and Logistics*, 9(2), 100008. https://doi.org/10.1016/j.ejtl.2020.100008.
- Ulmer, M. W., Mattfeld, D. C., & Köster, F. (2018). Budgeting time for dynamic vehicle routing with stochastic customer requests. *Transportation Science*, 52(1), 20-37. https://doi.org/10.1287/trsc.2016.0719.
- Ulmer, M. W., & Thomas, B. W. (2020). Meso-parametric value function approximation for dynamic customer acceptances in delivery routing. *European Journal of Operational Research*, 285(1), 183-195. https://doi.org/10.1016/j.ejor.2019.04.029.
- Van der Hagen, L., Agatz, N., Spliet, R., Visser, T. R., & Kok, A. L. (2024). Machine learningbased feasibility checks for dynamic time slot management. *Transportation Science*, 58(1), 94-109. https://doi.org/10.1287/trsc.2022.1183.
- Vansteenwegen, P., & Gunawan, A. (2019). Orienteering problems: Models and algorithms for vehicle routing problems with profits. Cham: Springer. https://doi.org/10.1007/978-3-030-29746-6.
- Vinsensius, A., Wang, Y., Chew, E. K., & Lee, L. H. (2020). Dynamic incentive mechanism for delivery slot management in e-commerce attended home delivery. *Transportation Science*, 54(3), 567-587. https://doi.org/10.1287/trsc.2019.0953.
- Voccia, S. A., Campbell, A. M., & Thomas, B. W. (2019). The same-day delivery problem for online purchases. *Transportation Science*, 53(1), 167-184. https://doi.org/10.1287/trsc.2016.0732.
- Wang, Y.-S., Wu, S.-C., Lin, H.-H., & Wang, Y.-Y. (2011). The relationship of service failure severity, service recovery justice and perceived switching costs with customer loyalty in the context of e-tailing. *International Journal of Information Management*, 31(4), 350-359. https://doi.org/10.1016/j.ijinfomgt.2010.09.001.
- Wang, L., Zeng, L., Ma, W., & Guo, Y. (2021). Integrating passenger incentives to optimize routing for demand-responsive customized bus systems. *IEEE Access*, 9, 21507-21521. https://doi.org/10.1109/ACCESS.2021.3055855.
- Xu, Z., Li, Z., Guan, Q., Zhang, D., Li, Q., Nan, J., Liu, C., Bian, W., & Ye, J. (2018). Large-scale order dispatch in on-demand ride-hailing platforms: A learning and planning approach. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery* & *Data Mining*, 905-913. https://doi.org/10.1145/3219819.3219824.
- Yan, C., Zhu, H., Korolko, N., & Woodward, D. (2020). Dynamic pricing and matching in ridehailing platforms. *Naval Research Logistics*, 67(8), 705-724. https://doi.org/10.1002/nav.21872.
- Yang, X., & Strauss, A. K. (2017). An approximate dynamic programming approach to attended home delivery management. *European Journal of Operational Research*, 263(3), 935-945. https://doi.org/10.1016/j.ejor.2017.06.034.
- Yang, X., Strauss, A. K., Currie, C. S. M., & Eglese, R. (2016). Choice-based demand management and vehicle routing in e-fulfillment. *Transportation Science*, 50(2), 473-488. https://doi.org/10.1287/trsc.2014.0549.
- Zhang, J., Luo, K., Florio, A. M., & Van Woensel, T. (2023). Solving large-scale dynamic vehicle routing problems with stochastic requests. *European Journal of Operational Research*, 306(2), 596-614. https://doi.org/10.1016/j.ejor.2022.07.015.

# Article A2: On the Concept of Opportunity Cost in Integrated Demand Management and Vehicle Routing

David Fleckenstein<sup>a</sup>, Robert Klein<sup>a</sup>, Vienna Klein<sup>a</sup>, Claudius Steinhardt<sup>b</sup> (<sup>a</sup>University of Augsburg, <sup>b</sup>University of the Bundeswehr Munich) Published in a category A journal: Transportation Science, online first (2024). https://doi.org/10.1287/trsc.2024.0644.

#### Abstract

Integrated demand management and vehicle routing problems are characterized by a stream of customers arriving dynamically over a booking horizon and requesting logistical services, fulfilled by a given fleet of vehicles during a service horizon. Prominent examples are attended home delivery and same-day delivery problems, where customers commonly have heterogeneous preferences regarding service fulfillment and requests differ in profitability. Thus, demand management methods are applied to steer the booking process to maximize total profit considering the cost of the routing decisions for the resulting orders. To measure the requests' profitability for any demand management method, it is common to estimate their opportunity cost. In the context of integrated demand management and vehicle routing problems, this estimation differs substantially from the estimation in the well-examined demand management problems of traditional revenue management applications as, for example, found in the airline or car rental industry. This is because of the unique interrelation of demand control decisions and vehicle routing decisions as it inhibits a clear quantification and attribution of cost, and of displaced revenue, to certain customer requests. In this paper, we extend the theoretical foundation of opportunity cost in integrated demand management and vehicle routing problems. By defining and analyzing a generic Markov decision process model, we formally derive a definition of opportunity cost and prove opportunity cost properties on a general level. Hence, our findings are valid for a wide range of specific problems. Further, based on these theoretical findings, we propose approximation approaches that have not yet been applied in the existing literature, and evaluate their potential in a computational study. Thereby, we provide evidence that the theoretical results can be practically exploited in the development of solution algorithms.

Key words: Last-mile Logistics, Demand Management, Markov Decision Process, Opportunity Cost

## 1 Introduction

The widespread adoption of digital distribution channels both enables and forces more and more logistical service providers to manage booking processes actively, in order to maintain competitiveness. As a result, their operational planning is no longer limited to solving vehicle routing problems (VRPs). Instead, providers integrate demand management to steer the booking process and either make established business models more profitable or operate novel ones profitably in the first place. These demand management approaches can comprise *demand control decisions* on *prices* of fulfillment options, the *availability* of fulfillment options, or the *acceptance/rejection* of requests.

Generally, the resulting integrated demand management and vehicle routing problems (i-DMVRPs) share a common structure (Fleckenstein et al., 2023 and Waßmuth et al., 2023): A service provider offers logistical services characterized by origin and destination in combination with other parameters like, e.g., service fees, time commitments, or vehicle types. These services are sold throughout a *booking horizon*, with customer requests arriving dynamically. The provider specifies a set of *fulfillment options* to offer in response to an incoming customer request, consisting of only a single option or multiple options with fixed or varying fees. Subsequently, the customer makes a purchase choice, i.e., places an order, based on their individual preferences and the offered options. Fulfillment of all customer orders takes place throughout the service horizon by means of a fixed number of vehicles. Capacities of other resources, like driver working hours, may also be limited. The booking and service horizons can be disjoint, which is typical for attended home delivery problems, or overlapping, which is common for same-day delivery and mobility-on-demand problems. Given the capacity restrictions as well as other operational constraints, such as potentially guaranteed service levels, the provider's typical objective is to control demand and routing in a profit-maximizing way, that is, to maximize the difference between revenue and cost.

Because i-DMVRPs are stochastic and dynamic, they can be modeled as Markov decision processes (MDPs) and, theoretically, decisions can be optimized by solving the well-known Bellman equation. However, as in demand management problems from traditional revenue management applications, like the airline or car rental industry (Klein et al., 2020), solving the Bellman equation is intractable for industry-sized instances. Therefore, it is common to approach demand management problems by decomposing them into two subproblems (e.g., Gallego and Topaloglu, 2019, p. 25), with the aim of eliminating the Bellman equation's recursiveness (Klein et al., 2018). The respective subproblems are 1) Approximating the opportunity cost for every potential fulfillment option to measure its profitability considering the remaining booking process and 2) solving the actual demand control problem based on the opportunity cost approximation. For this general decomposition-based solution concept, the overall performance largely depends on the quality of the underlying opportunity cost approximation (Klein et al., 2018), which is typically calculated as the difference of the state value approximation for the two resulting post-decision states (selling the fulfillment option vs. not selling it). Hence, in revenue management, the analysis of opportunity cost and its properties has already become a standard tool (e.g., Adelman (2007), Gallego and Topaloglu (2019), p. 10, Talluri and Van Ryzin (2004), p. 92). However, the corresponding results cannot be transferred directly to i-DMVRPs, due to the mutual interdependencies between demand control decisions and vehicle routing decisions across the entire planning horizon (Agatz et al., 2013).

This observation is the motivation for our work, which aims to inform and accelerate the development of more accurate opportunity cost approximation approaches. To this end, we consider opportunity cost at a formal level, and hence, draw on MDP models of i-DMVRPs as our primary object of study. Investigating these models, we derive mathematical properties and prove that they are valid for the entire family of i-DMVRPs. In a further step, we present three opportunity cost approximation approaches, which exploit the central property, namely the decomposability of opportunity cost. In a computational study, we analyze the potential of these approaches for a variety of problem settings. Primarily, this study is intended as a proof of concept for how the theoretical knowledge about the concept of opportunity cost can drive the development of solution approaches. In addition, the performance evaluation of the presented approaches can serve as a starting point for future research because it hints which approaches have the greatest potential in a certain setting.

Overall, our work has the following contributions:

- We deepen the theoretical foundation of opportunity cost in the context of i-DMVRPs. We
  do so by elaborating on decisive differences between opportunity cost in traditional revenue
  management applications and in i-DMVRPs and also by introducing a formal opportunity
  cost definition that is specifically tailored to i-DMVRPs.
- 2. We contribute to the existing literature on modeling i-DMVRPs and strengthen the connection between models and solution approaches by introducing a generic MDP model for i-DMVRPs. Further, to the best of our knowledge, we are the first to show formally how to separate demand control decisions from vehicle routing decisions at each decision epoch. This allows investigating the profit impact of demand control decisions in isolation. Additionally, we present a valid model transformation to restore properties in case a certain i-DMVRP application does not naturally inherit them.
- 3. We introduce and prove four central properties of opportunity cost for our i-DMVRP model that can be exploited within opportunity cost approximation approaches. Those properties are decomposability into two components, potential component-wise negativity, overall nonnegativity, and state value monotonicity.
- 4. Based on our theoretical findings and focusing on the decomposability as the central property, we present three types of approximation approaches that exploit this property and have yet to be applied to i-DMVRPs: single component approximation, a rather naive hybrid reward

approximation, and a more sophisticated hybrid reward approximation. Thereby, we illustrate how the theoretical results lead to direct, practical advances in algorithm development.

The remainder of this paper is structured as follows: In Section 2, we review the literature on opportunity cost in the context of demand management problems in general and on i-DMVRPs in particular. In Section 3, we first model the generic i-DMVRP, for which we show how its opportunity cost differs from the traditional interpretation in revenue management. Then, we present a formal definition of opportunity cost for i-DMVRPs. In Section 4, we elaborate and prove four central opportunity cost properties, which hold for the generic i-DMVRP. In Section 5, we present the approximation approaches and discuss the computational results. In Section 6, we summarize our work and outline opportunities for future research.

## 2 Literature Review

In this section, we give an overview of the related literature. We divide this overview into two parts. First, we briefly sketch the evolution of the integration of demand management and vehicle routing as a distinct research area (Section 2.1). Second, we review the scientific contributions to the analysis of opportunity cost for both traditional revenue management problems and i-DMVRPs and position our own work relative to these publications (Section 2.2).

## 2.1 Integrating Demand Management and Vehicle Routing

Specific similarities between traditional revenue management applications and selling logistical services, such as fixed resources and heterogeneous demand, have prompted the establishment of vehicle routing as a new application for demand management (Agatz et al., 2013). The differentiation of i-DMVRPs from "pure" stochastic and dynamic vehicle routing problems (for a recent review, see Soeffker et al. (2022)) is nontrivial. In the remainder of this work, we follow the definition in Fleckenstein et al. (2023). According to that, the control of demand with respect to profitability instead of only feasibility is the distinguishing feature of i-DMVRPs.

Thus, although there is some earlier work in the field of stochastic and dynamic vehicle routing, the works by Campbell and Savelsbergh (2005) and Campbell and Savelsbergh (2006) can be viewed as the first contributions to integrating active demand management and vehicle routing. These publications initiate the literature stream on attended home delivery problems (e.g., Koch and Klein, 2020, Vinsensius et al., 2020, and Yang et al., 2016), for which Snoeck et al. (2020) and Waßmuth et al. (2023) provide in-depth reviews. The corresponding problems feature disjoint booking and service horizons, meaning that bookings for a specific service horizon, usually a working day, are only possible until a certain cutoff time, such that there is no temporal overlap between booking and fulfillment processes.

On the contrary, Azi et al. (2012) present the first work on steering booking processes in parallel to fulfillment operations, that is, problems with overlapping booking and service horizons.
Therewith, they start the literature stream on same-day delivery (e.g., Klein and Steinhardt, 2023, and Ulmer, 2020), which is reviewed thoroughly by Li et al. (2024).

Another stream of literature on i-DMVRPs is initiated by Atasoy et al. (2015) and Hosni et al. (2014) and considers (shared) passenger transportation problems summed up under the term mobility-on-demand (e.g., Al-Kanj et al., 2020, Arian et al., 2022, Bertsimas et al., 2019, and Kullman et al., 2022). The corresponding problems commonly feature overlapping booking and service horizons as well. For a more extensive, cross-application review of the literature on i-DMVRPs, we refer the interested reader to the recent survey by Fleckenstein et al. (2023).

# 2.2 Theoretical Analysis of Opportunity Cost

As explained in Section 1, instead of approaching demand management problems holistically, solution concepts usually rely on decomposition. The idea is to separate the opportunity cost approximation from optimizing the demand control decision. Both subproblems can be tackled with separate approaches. In traditional revenue management applications, properties of the state values and opportunity cost, such as monotonicity or nonnegativity, have been successfully exploited to improve both the performance of opportunity cost approximation approaches and approaches for optimizing the subsequent demand control decisions. For the first task, that is, the approximation of opportunity cost, approaches based on linear programming (Adelman, 2007) as well as on statistical learning (Koch, 2017) are known to perform better if constraints are imposed to ensure that the resulting approximation also exhibits existing properties of opportunity cost. For the second task, that is, solving the demand control problem, certain properties simplify the computation of the optimal demand control decisions as, for example, shown in Talluri and Van Ryzin (2004) (p. 38) for single-resource capacity control. When designing solution approaches for specific applications of demand management, it is necessary to prove whether such structural properties hold or do not hold for the specific demand management problem, that is, under the assumptions associated with the underlying application (see, e.g., Maddah et al. (2010) for an example from cruise ship revenue management or Quante et al. (2009) for a manufacturing problem).

Because of the unique problem structure of i-DMVRPs, we cannot directly transfer findings from analyses of demand management problems in traditional revenue management or from other application areas of demand management. In the academic literature on i-DMVRPs, the majority of authors follow the general decomposition-based solution approach described in Section 1 (Fleck-enstein et al., 2023). However, the development of solution approaches is mainly driven by structural reasoning based on characteristics of specific i-DMVRPs paired with the validation of the respective approaches in a computational study. First, there are works that design opportunity cost approximations with the aim of capturing displacement effects regarding potential future orders (Avraham and Raviv, 2021, Lang et al., 2021, Ulmer, 2020, and Prokhorchuk et al., 2019). Second, some authors suggest that there is another component of opportunity cost in i-DMVRPs to be considered besides displacement cost: E.g., Arian et al. (2022) define opportunity cost as a

difference in future profit, which includes fulfillment cost. According to Klein et al. (2018), opportunity cost quantifies the "[...] "consequences" concerning potential future requests and the resulting routing cost [...]" (p. 971). Koch and Klein (2020), Yang et al. (2016), and Campbell and Savelsbergh (2005) state that the lost revenue of potential future orders as well as final fulfillment cost have to be anticipated when approximating opportunity cost. Vinsensius et al. (2020) introduce the term "marginal fulfillment cost" of a potential order, and Akkerman et al. (2022) aim at approximating changes in transportation cost. Abdollahi et al. (2023), Strauss et al. (2021), Mackert (2019), and Yang and Strauss (2017) provide the most extensive discussion of both revenue-side and cost-side future effects of a demand control decision.

Despite the substantial progress regarding approximation approaches, which the aforementioned publications have contributed to, there is hardly any work on formalizing and generalizing the underlying considerations aside from the following three publications (see also Waßmuth et al., 2023): Lebedev et al. (2021) and Asdemir et al. (2009) conduct a structural analysis of a specific i-DMVRP, namely an attended home delivery dynamic pricing problem with disjoint horizons. Based on MDP formulations that only implicitly model vehicle routing, they derive properties of the value function and the optimal pricing policy. In contrast to Asdemir et al. (2009), Lebedev et al. (2021) also model the cost-side via a fulfillment cost approximation. In comparison, our study is more general in that it considers a generic i-DMVRP and explicit vehicle routing decisions. It also focuses on the concept of opportunity cost independent from a specific demand management approach such as dynamic pricing. While the former two works take up a demand management-oriented view, Ulmer et al. (2020) focus on dynamic vehicle routing problems, which do not necessarily include demand management. They propose a novel MDP modeling framework and show its benefits for informing the design of solution approaches. Our work also aims to establish connections between modeling and solving i-DMVRPs but, different from Ulmer et al. (2020), with a focus on the demand control subproblem rather than on the vehicle routing subproblem.

In summary, our work closes existing research gaps in two ways: First, we provide a theoretical foundation for the existing qualitative reasoning and computational results. Second, our analysis provides the basis for developing algorithmic approaches that have not been considered in existing works, which we demonstrate in a computational study.

# 3 Opportunity Cost in Integrated Demand Management and Vehicle Routing Problems

In this section, we adapt and discuss the concept of opportunity cost specifically for i-DMVRPs. We first introduce a generic problem definition, on which we base our discussion throughout the whole section. For didactical reasons, we consider a problem as generic as possible and show later on (Section 4.5) how our insights can be transferred to more specific i-DMVRPs. We model the prototypical problem as an MDP (Section 3.1) and show the relevance of opportunity cost for

solving the introduced MDP (Section 3.2). Then, we discuss the structural differences of opportunity cost in i-DMVRPs compared with those in traditional revenue management applications and formalize a unified definition of opportunity cost specifically tailored to i-DMVRPs (Section 3.3). For ease of readability, we provide a list of the notation used throughout this paper in Appendix A.

### 3.1 Generic Problem Definition and Modeling

We discuss the concept of opportunity cost in the light of i-DMVRPs for the following generic problem: In each stage t = 1, ..., T within a finite booking horizon, at most one *customer request* of type  $c \in C$  can arrive with a certain arrival rate  $\lambda_c^t$ . A customer request of type c is characterized by the locations of its origin and destination, stored by parameter  $l_c$ , and revenue  $r_c$ . Without loss of generality, we assume that multiple individual customer requests can arrive from the same location with the same revenue, such that the arrival rates  $\lambda_c^t$  are independent of whether an individual customer request of type c has already realized before or not. Because the customer requests arrive sequentially, we can distinguish individual customer requests by their request time  $\tau$ .

The provider offers each arriving customer an offer set, which is a subset of a set of fulfillment options, for example, different time windows for order delivery. Once a customer books definitively, their request turns into a confirmed *customer order* that requires a certain amount of fulfillment resources, such as driving time or physical space in a vehicle, depending on the characteristics of the corresponding fulfillment option stored in parameter *o*. Requests may arrive in parallel to fulfillment operations, that is, the booking and the finite service horizon overlap. During the service horizon, all confirmed customer orders are served, and the provider incurs the resulting fulfillment cost defined as variable overhead cost arising from the execution of planned routes.

Because an individual opportunity cost value is associated with each (potential) order, that is, with a certain fulfillment option, modeling only a single fulfillment option is sufficient to generally analyze the concept of opportunity cost. By omitting explicitly modeling multiple fulfillment options, we obtain a much simpler model because the provider's decision space for demand control reduces from all possible offer sets to an accept/reject decision per request. Hence, for the sake of simplicity, we consider a single-option model in the remainder of this work. However, we will explain how to generalize the results of our discussion for problems requiring the explicit modeling of multiple fulfillment options in Section 4.5. In the following, we state the corresponding MDP model:

Decision epoch – A decision epoch defines the beginning of a stage of the MDP. In the considered problem, such stages correspond to (constant) time steps t = 1, ..., T. These time steps are sufficiently small to ensure that at most one customer request arrives between two decision epochs.

Hence,  $\lambda_c^t$  can well approximate the probability for observing exactly one request per time interval (Subramanian et al., 1999, Lee and Hersh, 1993).

State – The state of the system consists of all information that is known so far and relevant for decision-making. In i-DMVRPs with overlapping booking and service horizons, information about orders' and vehicles' statuses stored in two separate components are part of the state definition. First, state  $s_t$  at decision epoch t stores all confirmed customer orders for which fulfillment has not yet started as tuples  $(l_c, \tau, o)$  in set  $C_t$ . The second component is the overall tour plan at decision epoch t, denoted by  $\phi_t$ . It contains the currently running tours  $\theta_t^v$  for every vehicle  $v \in V$ . Thus, we define the state as:  $s_t = (C_t, \phi_t)$ . It is important to note that we construct the MDP around the post-decision state because this is the "natural" formulation when making decisions for an observed arriving request (Powell, 2022, p. 490). Hence,  $s_t$  consistently refers to a post-decision state. An important consequence of this is that a decision in decision epoch t is based on the information stored in state  $s_{t-1}$ .

Action – In an MDP model, the action  $a_t$  taken at decision epoch t corresponds to the realization of a certain decision. In the considered problem, at most one customer request, denoted by its request type c, can arrive in any stage between t = 1 and t = T. In the case of a request arrival, the provider must integratively make a demand control decision  $g_t \in \mathcal{G}(s_{t-1}, c) \subseteq \{0,1\}$ , i.e., accept or reject the arriving request of type c, and a tour planning decision  $\phi_t(g_t) \in \Phi(s_{t-1}, c, g_t)$ depending on whether the arriving request must be served according to the demand control decision. A tour planning decision is a decision on an update of the tour plan stored in the system state, which can also be the decision to leave the tour plan unchanged. The set  $\Phi(s_{t-1}, c, g_t)$ defines all potential tour plans that are feasible given the preceding (post-decision) state  $s_{t-1}$  and the demand control decision  $g_t$  for the arriving customer request of type c. In general, the tour plan must allow for the duly fulfillment of all confirmed customer orders. However, the precise definition of any feasible tour plan, that is, of the action space for the tour planning decision, depends on the specific problem. If no customer request arrives, we set  $g_t = 0$  and, thus, only a tour planning decision  $\phi_t(0) \in \Phi(s_{t-1}, 0)$  not including a new request is required. In summary, the action  $a_t$  is formally defined for three distinct cases:

$$a_{t} = \begin{cases} (0, \phi_{t}(0)) & \text{for } t = 1, ..., T, \text{ if there is no customer request arrival} \\ (0, \phi_{t}(0)) & \text{for } t = 1, ..., T, \text{ if the current customer request is rejected} \\ (1, \phi_{t}(1)) & \text{for } t = 1, ..., T, \text{ if the current customer request is accepted.} \end{cases}$$
(1)

The corresponding action space  $\mathcal{A}_t(s_{t-1}, c)$  at a decision epoch t when being in state  $s_{t-1}$  and receiving a customer request of type c comprises the two, above introduced components, that is, the demand control component  $\mathcal{G}(s_{t-1}, c)$  and the tour planning component  $\Phi(s_{t-1}, c, g_t)$ . Thus, the action space is formally defined as  $\mathcal{A}_t(s_{t-1}, c) = \{(g_t, \phi_t(g_t)): g_t \in \mathcal{G}(s_{t-1}, c), \phi_t(g_t) \in \Phi(s_{t-1}, c, g_t)\}$ . Thereby, the action space for the demand control decision depends on tour planning, because accepting a request of type c given the relevant state  $s_{t-1}$  is only feasible if at least one feasible tour plan exists, that is, if  $\Phi(s_{t-1}, c, 1) \neq \emptyset$ . However, for better readability, we omit this dependency as well as the dependency of the demand control decision  $g_t$  on the request type c in the notation. In case there is no customer request arrival at decision epoch t, the respective action space is defined as  $\mathcal{A}_t(s_{t-1}) = \Phi(s_{t-1}, 0)$ .

*Rewards* – The demand-control-related rewards  $r_c$  are received with actions  $g_t = 1$  for t = 1, ..., T. They are positive and equal the revenue of the customer request type c that is accepted at t, respectively. Demand control actions  $g_t = 0$  for t = 1, ..., T entail no rewards. Because we assume that the triangle inequality holds, the reward accrued with a tour planning decision is either zero (in case no new fulfillment vehicle operation is triggered) or negative (otherwise). We call those rewards logistics-related rewards and denote them formally by  $r_{\phi_t(g_t)}$ . They equal the negative of all fulfillment cost that are newly triggered with a decision  $\phi_t(g_t)$ , that is, the variable overhead cost of all new fulfillment operations that are executed definitively.

Transition – As a consequence of actions and stochasticity, the MDP transitions from a given state  $s_{t-1}$  to a successor state  $s_t$ . The second state component changes due to the execution of fulfillment operations according to the tour planning decision  $\phi_t(g_t)$  in action  $a_t$ . In the absence of stochastic elements in the fulfillment operations, such as stochastic travel times, this transition is purely deterministic. Thus, for state  $s_t$ ,  $\phi_t$  is set to  $\phi_t(g_t)$  from  $a_t$ . The first state component  $C_{t-1}$ , that is, the set of confirmed customer orders for which fulfillment has not yet started, changes as follows: First, the stochasticity of the i-DMVRP influences the transition of the first state component in the form of the potential request arrival according to time-dependent arrival rates  $\lambda_c^t$ . If a customer request arrives and turns into a customer order, it is added. We denote this particular order by  $c_t = (l_c^t, \tau^t, o^t)$ . Second, the subset of orders  $\Psi(\phi_t)$  for which the fulfillment process has started according to the new tour plan  $\phi_t$  are removed. The transitions of the state components can be formalized as follows:

$$\phi_t = \phi_t(g_t) \tag{2}$$

$$C_t = \begin{cases} C_{t-1} \setminus \Psi(\phi_t), & \text{if there is no customer request arrival or if } g_t = 0\\ (C_{t-1} \cup \{c_t\}) \setminus \Psi(\phi_t), & \text{if } g_t = 1 \end{cases}$$
(3)

*Objective* – The objective of the generic i-DMVRP is maximizing the expected profit across all decision epochs starting in state  $s_0$ . Thus, we aim at determining a policy x, with  $a_t^x(s_{t-1}, c_t) = (g_t^x(s_{t-1}, c_t), \phi_t^x(g_t^x(s_{t-1}, c_t)))$  denoting the action selected by the policy x at decision epoch t, according to the following objective function:

$$\max_{x} \mathbb{E} \left( \sum_{t=1}^{T} \left( r_{c_{t}} \cdot g_{t}^{x}(s_{t-1}, c_{t}) + r_{\phi_{t}^{x}}(g_{t}^{x}(s_{t-1}, c_{t})) \right) \mid s_{0} \right).$$
(4)

### 3.2 State Values and Opportunity Cost

We now represent the previously introduced objective of the generic i-DMVRP (4) by the corresponding value function that equals the well-known Bellman equation. It captures the value of being in a given state and can be applied to find an optimal policy for the MDP model (Powell, 2022, p. 46). Specified for the generic model, the value function explicitly models the mutual temporal interdependencies of the two integrated decisions, that is, the demand control decision and the tour planning decision:

$$V_{t-1}(s_{t-1}) = \sum_{c \in C} \lambda_c^t \cdot \max_{g_t \in \mathcal{G}(s_{t-1},c)} \left( g_t r_c + \max_{\phi_t(g_t) \in \Phi(s_{t-1},c,g_t)} \left( r_{\phi_t(g_t)} + V_t \left( s_t \mid s_{t-1}, \phi_t(g_t) \right) \right) \right) + (1 - \sum_{c \in C} \lambda_c^t) \cdot \max_{\phi_t(0) \in \Phi(s_{t-1},0)} \left( r_{\phi_t(0)} + V_t \left( s_t \mid s_{t-1}, \phi_t(0) \right) \right),$$
(5)

with boundary condition:

$$V_{T}(s_{T}) = 0.$$
(6)
  
Fulfillment operations arrival type c of transition transitio

Fig. 1 Overview of the MDP model of the i-DMVRP booking and fulfillment process including the interim state In i-DMVRPs, an action can comprise two types of integrated decisions, namely, demand control and tour planning decisions. In this work, the effects of a demand control decision are of interest. Thus, it is the target to calculate opportunity cost from comparing state values that reflect such effects separated from potential effects of tour planning decisions at the same decision epoch. Therefore, we introduce a fictive state for each decision epoch t = 1, ..., T. We refer to it as the *interim state* and denote it as  $s'_t | s_{t-1}, c, g_t$ . Technically, c and  $g_t$  are captured in additional state dimensions of the interim state. The interim state describes the state that is reached if the provider accepts ( $g_t = 1$ ) a customer request of type c starting in state  $s_{t-1}$  or rejects it ( $g_t = 0$ ). In other words, the state is measured after the demand control decision but before the integrated tour planning decision. The idea behind it is comparable to the idea of the post-decision state introduced by Powell (2011) (p. 129) with the aim of isolating different effects of decisions and information on the state variable. However, the post-decision state separates the *deterministic effect* of a decision from the stochastic effect of the same decision in order to ease decision-making. The interim state, instead, separates the effects of two different decisions, that is, the effects of the demand control decision from the effects of a tour planning decision taken in the same decision epoch. Fig. 1 illustrates the interim state within the decision process and its components.

We denote the value of interim state  $s'_t | s_{t-1}, c, g_t$  by  $V'_t(s'_t | s_{t-1}, c, g_t)$ . Generally, it can be calculated as the sum of the succeeding post-decision state's value, that is, of state  $s_t | s_{t-1}, \phi^*_t(g_t)$ , and the logistics-related rewards of decision epoch t:

$$V_t'(s_t' \mid s_{t-1}, c, g_t) = \max_{\phi_t(g_t) \in \Phi(s_{t-1}, c, g_t)} \left( r_{\phi_t(g_t)} + V_t(s_t \mid s_{t-1}, \phi_t(g_t)) \right)$$

$$= r_{\phi_t^*(g_t)} + V_t(s_t \mid s_{t-1}, \phi_t^*(g_t)),$$
(7)

with  $\phi_t^*(g_t)$  denoting the optimal tour planning decision given demand control decision  $g_t$  at decision epoch t. Note that this simplification of notation will be used repeatedly throughout the remainder of the paper. Based on interim state values, we can formulate a simplified variant of the value function (5) isolating the demand control decision:

$$V_{t-1}(s_{t-1}) = \sum_{c \in C} \lambda_c^t \cdot \max_{g_t \in \mathcal{G}(s_{t-1},c)} (g_t \cdot r_c + V_t'(s_t' \mid s_{t-1}, c, g_t)) + (1 - \sum_{c \in C} \lambda_c^t) \cdot V_t'(s_t' \mid s_{t-1}, 0).$$
(8)

In the remainder of our discussion, we denote interim states  $s'_t | s_{t-1}, c, 1$  by  $s'_t(c)$  and interim states  $s'_t | s_{t-1}, c, 0$  (or  $s'_t | s_{t-1}, 0$  in case there is no customer request) by  $s'_t(0)$  for ease of presentation. Based on the interim state, the following definition formalizes the concept of opportunity cost for solving the demand control problem of our generic i-DMVRP:

**Definition 1.** The opportunity cost  $\Delta V_t(s_{t-1}, c)$  of accepting a customer request of type c in a certain state  $s_{t-1}$  is calculated as the difference of the values of the following two interim states: (1) the interim state following the rejection of customer request c and (2) the interim state following the acceptance of c. Thus, it is defined as:

$$\Delta V_t(s_{t-1}, c) = V'_t(s'_t(0)) - V'_t(s'_t(c)).$$
(9)

This opportunity cost is then used as input to solve the demand control problem, which can be illustrated by the following reformulation of the value function (8). This reformulation is typical in the revenue management literature (e.g., Strauss et al., 2018) and yields:

$$V_{t-1}(s_{t-1}) = \sum_{c \in C} \lambda_c^t \cdot \underbrace{\max_{g_t \in \mathcal{G}(s_{t-1},c)} \left( g_t \cdot \left( r_c - \Delta V_t(s_{t-1},c) \right) \right)}_{\text{Demand control subproblem}} + V_t'(s_t'(0)). \tag{10}$$

Because the provider only takes a demand control decision when a certain customer request arrives, the probability  $\sum_{c \in C} \lambda_c^t$  is not relevant for decision-making. Also, the second summand of equation (10), that is,  $V'_t(s'_t(0))$ , is not relevant as it is a constant and independent of the decision. Further, the provider knows  $r_c$ . Thus, given the opportunity cost of a customer request of type c,  $\Delta V_t(s_{t-1}, c)$ , it is possible to solve the demand control problem as a deterministic subproblem. This is why, for industry-sized problems, it is necessary to find accurate and efficient approximation approaches for opportunity cost, that is, for the value function (Strauss et al., 2018). This motivates a deeper understanding of opportunity cost and of its peculiarities and properties in i-DMVRPs. Knowing certain properties enables exploiting them to accelerate and enhance approaches to approximate opportunity cost as discussed in Section 2.2. Consequently, in the following, we compare opportunity cost in traditional settings, that is, in revenue management problems, and opportunity cost in i-DMVRPs and carve out decisive differences.

## 3.3 Generalization of the Concept of Opportunity Cost for i-DMVRPs

In traditional revenue management applications, the concept of opportunity cost bases on two main assumptions (see Weatherford and Bodily, 1992) that cause the opportunity cost to be equivalent to displacement cost (DPC). Those are defined as "[...] the expected loss in future revenue from using the capacity now rather than reserving it for future use." (Talluri and Van Ryzin, 2004, p. 33) In the following, we show that this definition cannot be transferred to i-DMVRPs by stating each of the underlying assumptions and investigating it in the respective context:

Assumption 1 – Supply is inflexible, that is, resource capacities are fixed: In i-DMVRPs, either driver working times, fleet sizes, or loads represent resources with fixed capacities. As expected, such limited resources may cause a displacement of demand (see Example 1 in Appendix B.1). Thus, in most i-DMVRPs, the first assumption is valid.

Assumption 2 – Variable cost associated with the usage of capacity are either negligible or at least directly attributable to individual orders: This does not hold in most i-DMVRPs, which can be shown by considering fuel cost as an example: Because the fuel consumption of a fulfillment tour depends on the specific combination of customer locations in the tour, there is no way to calculate and attribute the share and the resulting cost of each individual customer location (e.g., Vinsensius et al., 2020). Further, in i-DMVRPs, such variable overhead cost are not negligible. The exact same combination of state and customer request of the same problem instance can yield different optimal demand control decisions depending on whether fulfillment cost are taken into consideration or are neglected (see Examples 2a and 2b in Appendix B.2).

Consequently, we must adapt the traditional concept of opportunity cost, which equalizes opportunity cost and expected displacement cost (Talluri and Van Ryzin, 2004, p. 33), for i-DMVRPs. More precisely, a concept is needed that explicitly takes into account variable overhead cost related to order fulfillment: In the literature on i-DMVRPs, some authors already explicitly consider the marginal increase of variable overhead cost caused by the acceptance of a customer request and refer to it as *marginal cost-to-serve (MCTS)* (e.g., Strauss et al., 2021, Vinsensius et al., 2020). For myopic decision-making, that is, when neglecting future orders, we can calculate a request's MCTS by optimizing the tour plan for all accepted customer orders including the current request and comparing its variable fulfillment cost with the cost of the optimal tour plan without the current request (see Example 2b in Appendix B.2). However, such myopic MCTS are not sufficient for optimal decision-making (see Example 3 in Appendix B.3). In summary, for optimal decision-making, opportunity cost for i-DMVRPs cannot only be revenue-related in the form of expected DPC, but also have to take cost-related effects into account in the form of expected MCTS. Correspondingly, we amend the definition of opportunity cost as follows:

**Definition 2.** In i-DMVRPs with variable overhead cost that are not directly attributable to customer requests, opportunity cost comprises two components: DPC as the difference of cumulative

expected future revenue caused by accepting a customer request and MCTS as the difference of expected future fulfillment cost caused by accepting a customer request.

# 4 Properties and Analytical Discussion of Opportunity Cost for i-DMVRPs

We have showed in Section 3 that opportunity cost is calculated as the difference of two value functions. Hence, opportunity cost, as well as its components DPC and MCTS, are recursive functions that are intractable for realistic-sized i-DMVRPs. Consequently, solving i-DMVRPs requires accurate approximations of value functions or opportunity cost. To support the development and selection of respective approximation approaches, we discuss four central properties of the generic i-DMVRP, or more precisely, of the corresponding value function (5) and the derived opportunity cost values. Those are as follows.

- 1. Decomposability into DPC and MCTS
- 2. Potential negativity of DPC and MCTS
- 3. Non-negativity of opportunity cost
- 4. Monotonicity of the value function

Please note, for ease of readability, we move the minor mathematical proofs to Appendix C and only state the final proofs of the central properties throughout our discussions.

## 4.1 Decomposability into Displacement Cost and Marginal Cost-to-Serve

To prove the decomposability of opportunity cost into DPC and MCTS, we first define both terms formally, starting with a definition of *expected future revenue* and *expected future fulfillment cost* of a certain interim state  $s'_{t-1}$ .

**Definition 3.** The expected future revenue  $R'_{t-1}(s'_{t-1})$  of a given interim state  $s'_{t-1}$  at decision epoch t - 1 is defined as:

$$R_{t-1}'(s_{t-1}') = \sum_{c \in C} \lambda_c^t \cdot \left( g_t^* \cdot r_c + R_t'(s_t'|s_{t-1}', \phi_{t-1}^*(g_{t-1}), g_t^*) \right) + (1 - \sum_{c \in C} \lambda_c^t) \cdot R_t'(s_t'|s_{t-1}', \phi_{t-1}^*(g_{t-1}), 0),$$
(11)

with boundary condition:

$$R'_{T}(s'_{T}) = 0 (12)$$

and  $g_t^*$  denoting the optimal demand control decision, given that a customer request of type *c* arrives in state  $s_t$ .

**Definition 4.** The expected future fulfillment cost  $F'_{t-1}(s'_{t-1})$  of a given interim state  $s'_{t-1}$  at decision epoch t - 1 is defined as:

$$F'_{t-1}(s'_{t-1}) = r_{\phi^*_{t-1}(g_{t-1})} + \sum_{c \in C} \lambda^t_c \cdot \left(F'_t(s'_t \mid s'_{t-1}, \phi^*_{t-1}(g_{t-1}), g^*_t)\right) + (1 - \sum_{c \in C} \lambda^t_c) \cdot F'_t(s'_t \mid s'_{t-1}, \phi^*_{t-1}(g_{t-1}), 0),$$
(13)

with boundary condition:

$$F'_T(s'_T) = r_{\phi^*_T(g_T)}.$$
(14)

Based on Definitions 3 and 4, we now formally define DPC and MCTS:

**Definition 5.** DPC of accepting a customer request of type *c* at decision epoch *t* and state  $s_{t-1}$  is defined as:

$$\Delta R_t(s_{t-1}, c) = R'_t(s'_t(0)) - R'_t(s'_t(c)).$$
(15)

**Definition 6.** MCTS of accepting a customer request of type *c* at decision epoch *t* and state  $s_{t-1}$  is defined as:

$$\Delta F_t(s_{t-1}, c) = F'_t(s'_t(0)) - F'_t(s'_t(c)).$$
(16)

DPC and MCTS both depend on the state of the system and all consecutive decisions and transitions. Thus, both suffer from the curse of dimensionality (Powell et al., 2012) in that the number of potential tour planning decisions that must be evaluated is intractable for realistic-sized instances.

Now, we show that there is a valid decomposition of the value function (7) for interim states into two components. In other words, the *value* of any interim state, that is,  $V'_t(s'_t)$ , equals the sum of expected future revenue,  $R'_t(s'_t)$ , and expected future fulfillment cost,  $F'_t(s'_t)$ . This leads to the following lemma, based on which we then define the first property:

**Lemma 1**. The value (function) of an interim state  $s'_t$  can be decomposed into two additive components, one capturing expected future revenue and one capturing expected future fulfillment cost:

$$V'_t(s'_t) = R'_t(s'_t) + F'_t(s'_t).$$
<sup>(17)</sup>

Property 1. Opportunity cost can be decomposed into DPC and MCTS:

$$\Delta V_t(s_{t-1}, c) = \Delta R_t(s_{t-1}, c) + \Delta F_t(s_{t-1}, c).$$
(18)

*Proof.* To prove Property 1, we substitute Lemma 1 into Equation (9). Further, we substitute Definitions 5 and 6, which results in:

$$\Delta V_t(s_{t-1}, c) = V'_t(s'_t(0)) - V'_t(s'_t(c))$$
  
=  $\left(R'_t(s'_t(0)) + F'_t(s'_t(0))\right) - \left(R'_t(s'_t(c)) + F'_t(s'_t(c))\right)$   
=  $R'_t(s'_t(0)) - R'_t(s'_t(c)) + F'_t(s'_t(0)) - F'_t(s'_t(c))$   
=  $\Delta R_t(s_{t-1}, c) + \Delta F_t(s_{t-1}, c)$  (19)

Please note that despite DPC and MCTS can be expressed as two separate terms, the decisions they stem from are still interconnected. More precisely, to optimally calculate one of the components, the other one has to be taken into account because both Equation (11) and (13) incorporate optimal decisions that can only be determined based on the original value function including DPC and MCTS.

## 4.2 Potential Negativity of DPC and MCTS

Contrary to demand management problems of traditional revenue management applications in which DPC can only be non-negative (Talluri and Van Ryzin (2004), p. 217), in i-DVMRPs, DPC and MCTS can be negative, which is the next property we discuss.

*Negative DPC* – The intuition behind negative DPC, as they occur in Example 5 in Appendix B.5, is the following: turning the considered customer request into a customer order enables accepting one or more expected future customer requests in its vicinity that otherwise would not be profitable regarding their fulfillment cost and revenue.

*Negative MCTS* – The intuition behind negative MCTS, as they occur in Example 4 in Appendix B.4, is the following: Accepting a corresponding customer request and following the subsequent optimal decisions leads to expected future fulfillment cost that is lower than the cost generated by optimal decisions following the rejection of the same customer request. In other words, accepting a certain customer request inhibits the acceptance of one or more future customer requests, which would otherwise be accepted with optimal decisions and would lead to a longer tour, that is, larger fulfillment cost.

This is a decisive difference between the traditional concept of opportunity cost and the newly derived concept for i-DMVRPs.

Property 2. DPC and MCTS can both be negative.

*Proof.* By Example 4 and Example 5 in Appendix B. ■

### 4.3 Nonnegativity of Opportunity Cost

Despite the finding that both DPC and MCTS can be negative, we can show that for the generic MDP model with the value functions defined by (5) or (8), opportunity cost, that is, the sum of DPC and MCTS, is always nonnegative. To prove this property, we show that the value of the interim state following the acceptance of a customer request  $c_t$  by action  $g_t = 1$  cannot be greater than the value of the interim state following a rejection of the same customer request  $c_t$  by action  $g_t = 0$ . The corresponding proof builds on three lemmata (Lemma 2, Lemma 3, and Lemma 4), which formalize characteristics that are valid for the i-DMVRP model defined in Section 3.1. First, Lemma 2 concerns the stochastic transition probabilities, that is, the arrival rates  $\lambda_c^t$  in a stage t.

**Lemma 2**. Stochastic transition probabilities are independent of the set of already confirmed customer orders:

 $\forall t \in 1, \dots, T, c \in C: \lambda_c^t \text{ independent of } \mathcal{C}_{t-1}.$ (20)

Second, Lemma 3 concerns the relationship of the action spaces at decision epoch t when starting in any two states  $s_{t-1}$  and  $\hat{s}_{t-1}$ , that only differ in that the latter contains exactly one additional customer order, denoted by  $\hat{c}$ , that is,  $\hat{C}_{t-1} = C_{t-1} \cup {\hat{c}}$ . Starting in those two states, the action space resulting from the latter is a subset of the action space resulting from the former. **Lemma 3.** The action space resulting from any state  $\hat{s}_{t-1} = (C_{t-1} \cup {\hat{c}}, \phi_{t-1})$  is a subset of the action space resulting from a corresponding state  $s_{t-1} = (C_{t-1}, \phi_{t-1})$ :

$$\forall t \in 1, \dots, T, c \in C, \hat{c} \in C: \mathcal{A}(\hat{s}_{t-1}, c) \subseteq \mathcal{A}(s_{t-1}, c).$$

$$(21)$$

Third, Lemma 4 concerns the state space of an i-DMVRP MDP model. More precisely, it claims that, for any state  $\hat{s}_t = (\hat{C}_t, \phi_t)$ , there exists a state  $s_t = (\mathcal{C}_t, \phi_t)$ , which only differs in that it does not include a certain customer order  $\hat{c}$ .

Lemma 4. For every state 
$$\hat{s}_t = (\hat{C}_t, \phi_t)$$
, there exists a state  $s_t = (C_t, \phi_t)$  with  $C_t = \hat{C}_t \setminus \{\hat{c}\}$ :  
 $\forall \hat{s}_t \text{ with } t \in 1, \dots, T: \exists s_t : C_t = \hat{C}_t \setminus \{\hat{c}\}.$ 
(22)

We now consider a certain decision sequence, denoted as  $\pi = (\phi_t(g_t), a_{t+1}, a_{t+2}, \dots, a_T)$ , and apply it to a certain sample path  $\omega = (c_t, c_{t+1}, c_{t+2}, \dots, c_T)$ . Both start in an interim state  $s'_t$ . A sample path is a specific sequence of stochastic realizations throughout the decision epochs. Thus, for this sample path, the respective request arrival probabilities in Equation (5) equal 1 and the probabilities of other realizations  $\omega' \neq \omega$  equal 0.

Given Lemma 2 to 4, two further lemmata regarding the resulting *revenue*, denoted by  $R_t^{\prime\pi\omega}(s_t^{\prime})$ , and regarding the resulting *fulfillment cost*, denoted by  $F_t^{\prime\pi\omega}(s_t^{\prime})$ , can be derived: More precisely, Lemma 5 states that, applying  $\pi$  to  $\omega$ , assuming it starts in the interim state  $s_t^{\prime}(c_t)$ , results in the same cumulative revenue as assuming  $\omega$  starts in the corresponding interim state  $s_t^{\prime}(0)$ . This is because of the circumstance that in the interim state the revenue of  $c_t$  has already been collected. Lemma 6 states that, applying  $\pi$  to  $\omega$ , assuming it starts in the interim state  $s_t^{\prime}(c_t)$ , results in higher or equal fulfillment cost as assuming  $\omega$  starts in the corresponding interim state  $s_t^{\prime}(0)$ .

**Lemma 5.** Applying decision sequence  $\pi$  to sample path  $\omega$ , assuming it starts in interim state  $s'_t(c_t)$ , results in the same cumulative revenue as assuming  $\omega$  starts in the interim state  $s'_t(0)$ :

$$R_t^{\prime \pi \omega}(s_t^{\prime}(c_t)) = R_t^{\prime \pi \omega}(s_t^{\prime}(0)).$$
<sup>(23)</sup>

**Lemma 6**. Applying decision sequence  $\pi$  to sample path  $\omega$ , assuming it starts in the interim state  $s'_t(c_t)$ , results in higher or equal fulfillment cost as assuming  $\omega$  starts in the corresponding interim state  $s'_t(0)$ :

$$F_t^{\prime\pi\omega}(s_t^{\prime}(c_t)) \le F_t^{\prime\pi\omega}(s_t^{\prime}(0)).$$
<sup>(24)</sup>

Combining Lemmata 5 and 6 shows that applying a decision sequence  $\pi$  to sample path  $\omega$  starting in interim state  $s'_t(c_t)$  cannot result in a greater objective value than starting in interim state  $s'_t(0)$ . We formalize this in the following lemma:

**Lemma 7**. If the same decision sequence  $\pi$  is applied to sample path  $\omega$  starting in interim state  $s'_t(c_t)$ , it cannot yield a greater value than it does when starting in interim state  $s'_t(0)$ :

$$V_t^{\prime \pi \omega}(s_t^{\prime}(c_t)) \le V_t^{\prime \pi \omega}(s_t^{\prime}(0)).$$
 (25)

With this in mind, we can formally prove the third opportunity cost property.

Property 3. Opportunity cost is generally non-negative:

 $\forall c \in C, t = 1, ..., T : \Delta V_t(s_{t-1}, c) \ge 0.$ (26)

*Proof.* The proof is by contradiction. For a sample path  $\omega$ , starting in an interim state  $s'_t(c_t)$ , the optimal sequence of decisions, denoted by  $\pi^*(c_t)$ , results in value  $V'_t^{\pi^*(c_t)\omega}(s'_t(c_t))$ . We now assume that this value is higher than any value that we can accrue on the same sample path starting in interim state  $s'_t(0)$ . However, with Lemma 2 to 4, we can feasibly apply  $\pi^*(c_t)$  to the sample path starting in interim state  $s'_t(0)$  and, with Lemma 7, this results in at least the same value. The original assumption is proven wrong. Hence, the following holds:

$$V_t^{\prime\omega}(s_t^{\prime}(c_t)) = V_t^{\prime\pi^*(c_t)\omega}(s_t^{\prime}(c_t)) \le V_t^{\prime\pi^*(c_t)\omega}(s_t^{\prime}(0)) \le V_t^{\prime\pi^*(0)\omega}(s_t^{\prime}(0)) = V_t^{\prime\omega}(s_t^{\prime}(0)), \quad (27)$$

with  $\pi^*(0)$  being the optimal sequence of decisions for sample path  $\omega$ , starting in interim state  $s'_t(0)$ . This proof by contradiction can be replicated for every sample path  $\omega \in \Omega$ . Then, because the Bellman function represents the expected value over all possible sample paths  $\omega \in \Omega$ , following their respective optimal decision sequences  $\pi^*$ , it holds that:

$$V_t'(s_t'(c_t)) = V_t'^{\pi^*(c_t)}(s_t'(c_t)) \le V_t'^{\pi^*(0)}(s_t'(0)) = V_t'(s_t'(0)),$$
(28)

and substituted in Equation (9), this proves that  $\Delta V_t(s_{t-1}, c_t) \ge 0$ .

### 4.4 Monotonicity of the Value Function

We now investigate the monotonicity of the value function in confirmed customer orders, and across decision epochs, as typically done in traditional revenue management (e.g., Gallego and Topaloglu, 2019, p. 10, Adelman, 2007). Property 3 directly implies that the value function is monotonically decreasing in the confirmed customer orders for which fulfillment has not yet started  $c \in C_{t-1}$  at a certain decision epoch t, that is, the following holds:

$$\Delta V_t(s_{t-1}, c) = V'_t(s'_t(0)) - V'_t(s'_t(c)) \ge 0 \Leftrightarrow V'_t(s'_t(0)) \ge V'_t(s'_t(c)).$$
<sup>(29)</sup>

Analogously, we can investigate monotonicity in time across decision epochs. More precisely, we consider the monotonicity of state values of *consecutive* states  $s_t$  and  $s_{t'}$ , with t' > t. For two states  $s_t$  and  $s_{t'}$  to be consecutive, there must exist a sequence of (potentially multiple) stochastic transitions, that is, request arrivals, such that making optimal decisions  $a_t^*(s_t, c_t) = (g_t^*, \phi_t^*(g_t^*))$  causes the decision process to transition from  $s_t$  to  $s_{t'}$  in a finite number of decision epochs. Then, the value function (5) of the generic MDP model presented in Section 3.1 is clearly not monotonically decreasing across consecutive states. This is due to the negative logistics-related rewards arising throughout the decision process. Consider, for example, a decision epoch at which the routing constraints do not allow feasibly accepting any new order until the end of the service horizon. In this case, no future revenue will be collected. However, the tour planning decisions required for the fulfillment of the confirmed orders cause future logistics-related rewards, which are negative. Over the remaining service horizon, these negative rewards realize and the state value increases, that is, becomes less negative, over the remaining decision epochs until it equals zero in the terminal state.

Despite this finding, it is possible to achieve value function monotonicity across consecutive states by modifying the generic MDP model as shown in Appendix D such that it still models the exact same problem, that is, such that the modified model is mathematically equivalent to the original model. Formally, the modified model differs from the original model regarding *cost realization* and *cost modeling*. Cost realization concerns the point of time in which the cost is incurred in the real application. Cost modeling concerns the decision epoch in which the corresponding cost is taken into account within the MDP model as a negative reward. In the original model, cost realization and cost modeling match. In the modified model, we delay cost modeling. To this end, we augment the state of the modified model and adapt the transition and the value function accordingly, whereas all other model components remain unaltered:

State – For the modified model, we add a third state component denoted by  $r_t^{l\ cum}$ . It captures the cumulative logistics-related rewards, that is, the negative of the cumulative fulfillment cost that realized before or at decision epoch t. Thus, we define the state as:  $s_t = (C_t, \phi_t, r_t^{l\ cum})$ . The state space comprises all combinations of possible customer requests and arrival times with potential tour plans and potential cumulative logistics-related rewards.

*Transition* – The transition of the additional state component  $r_t^{l \ cum}$  equals:  $r_t^{l \ cum} = r_{t-1}^{l \ cum} + r_{\phi_t(g_t)}$ . The transition of all other components remains unaltered as introduced in Section 4.

*Value function* – For the modified model, we delay cost modeling to decision epoch *T*. Consequently, during the decision epochs t = 1, ..., T, only rewards  $r_c$  are considered in the value function, which is hence defined as:

$$\widetilde{V}_{t-1}(s_{t-1}) = \sum_{c \in C} \lambda_c^t \cdot \max_{g_t \in \mathcal{G}(s_{t-1},c)} \left( g_t \cdot r_c + \max_{\phi_t(g_t) \in \Phi(s_{t-1},c,g_t)} \widetilde{V}_t \left( s_t \mid s_{t-1}, \phi_t(g_t) \right) \right) \\
+ (1 - \sum_{c \in C} \lambda_c^t) \cdot \max_{\phi_t(0) \in \Phi(s_{t-1},0)} \widetilde{V}_t \left( s_t \mid s_{t-1}, \phi_t(0) \right).$$
(30)

We only consider rewards  $r_{\phi_t(g_t)}$  in the boundary condition such that the salvage value equals the respective state component:

$$\tilde{V}_T(s_T) = r_T^{l \ cum}.$$
(31)

In the following, we show that the value function of the modified model (30), denoted by  $\tilde{V}_t(s_t)$ , is monotonically decreasing across consecutive states, that is,  $\tilde{V}_t(s_t) \ge \tilde{V}'_{t+1}(s'_{t+1}) \ge \tilde{V}_{t'}(s'_t)$ , with  $\tilde{V}'_{t+1}(s'_{t+1})$  denoting the value of interim state  $s'_{t+1}$  at decision epoch t + 1 in the modified model. Thus, we must show that  $\tilde{V}_t(s_t) \ge \tilde{V}'_{t+1}(s'_{t+1})$  and  $\tilde{V}'_{t+1}(s'_{t+1}) \ge \tilde{V}_{t+1}(s_{t+1})$  holds for any pair of consecutive states  $s_t$  and  $s_{t+1}$  as this directly implies that  $\forall t < t' \le T: \tilde{V}'_t(s'_t) \ge \tilde{V}'_{t'}(s_{t'})$ .

**Property 4**. The value function of the modified model is monotonically decreasing in the course of decision epochs for consecutive states  $s_t$  and  $s_{t+1}$ :

$$\forall t = 0, \dots, T - 1 : \tilde{V}_t(s_t) \ge \tilde{V}_{t+1}(s_{t+1}).$$
(32)

*Proof.* The fact that  $\tilde{V}_t(s_t) \ge \tilde{V}'_{t+1}(s'_{t+1})$  follows directly from Equation (10) with the following line of reasoning: Starting in a certain post-decision state  $s_t$  in decision epoch t means that there

is one more customer request  $c_{t+1}$  potentially contributing to the state value compared with starting in the resulting interim state  $s'_{t+1}$ . If all potentially arriving requests are not profitable based on  $s_t$ , the optimal demand control decision is the rejection in any case, and  $\tilde{V}_t(s_t) = \tilde{V}'_{t+1}(s'_{t+1})$ holds, as no revenue can be collected in t + 1. Otherwise, if there is a nonempty subset of potentially arriving requests that is profitable, it is optimal to accept these requests and the associated expected revenue positively contributes to  $\tilde{V}_t(s_t)$ . Then,  $\tilde{V}_t(s_t) > \tilde{V}'_{t+1}(s'_{t+1})$  holds.  $\tilde{V}'_{t+1}(s'_{t+1}) \ge \tilde{V}_{t+1}(s_{t+1})$  directly follows from Equation (7) because for the considered value function  $\forall t \in 0, ..., T - 1 : r_{\phi^*_{t+1}(g_{t+1})} = 0$  holds by definition of the modified MDP formulation. Thus,  $\forall t = 0, ..., T - 1$ , it holds that  $\tilde{V}'_{t+1}(s'_{t+1}) = \tilde{V}_{t+1}(s_{t+1})$ .

Please note, Properties 1 to 3 as well as the monotonicity of the value function in confirmed customer orders also hold for the modified model. The respective proofs are straightforward with  $\tilde{F}'_{t-1}(s'_{t-1}) = F'_{t-1}(s'_{t-1}) + r^{l \ cum}_t$  (analogously to the transformation of  $\tilde{V}'_{t-1}(s'_{t-1})$  as previously described and proven in Appendix D). Because OC, DPC, as well as MCTS are defined as the differences of the respective value functions, (7), (11), and (13) for different interim states on the same stage t, the constant  $r^{l \ cum}_t$  cancels in Equations (9), (15), and (16). Thus, for the modified model, our proofs can be applied as conducted for the original model.

#### 4.5 Generalization for Multiple Fulfillment Options

All properties formulated in Section 4 are also valid for i-DMVRPs with multiple fulfillment options. Compared with accept/reject control, the difference in this case is that the demand control decision consists of selecting an offer set of feasible fulfillment options. Thus, there is an interim state for each fulfillment option the customer can possibly choose, preceded by an additional stochastic transition according to the customer's purchase choice probabilities. Opportunity cost is then defined separately for each fulfillment option as the difference of the interim state value for the customer choosing the particular option compared with choosing the no-purchase option, as shown in Appendix E.

Despite these differences, the two types of rewards, revenue and cost, can still be separated, such that Lemma 1 remains valid and the proof for Property 1 can be conducted as presented for a single fulfillment option. As the multi-option case generalizes the single-option case, Example 4 and Example 5 also prove Property 2 for multiple options. Likewise, modeling multiple options does not affect the validity of Lemma 2 because the additional stochastic transition reflecting the customer's purchase choice is also independent of the set of already confirmed customer orders. Lemma 3 also holds for multiple fulfillment options because the feasibility of each fulfillment option case. Based on the basic lemmata, the remainder of the proof for Property 3 can be conducted in a similar way as presented above for the single-option case. Finally, the line of reasoning in the proof of Property 4 can also be made based on profitable fulfillment options instead of customer requests.

# 5 Computational Study

The theoretical results obtained in Section 3 and Section 4 contribute to a deeper understanding of i-DMVRPs' mathematical structure, which is useful in itself. On top of that, they can also be of direct, practical use because they represent domain knowledge that can be exploited by a variety of solution approaches. In the case of the nonnegativity of opportunity cost (Property 3) and the monotonicity of the value function (Property 4), first promising results in this regard are found by Koch and Klein (2020). They consider a problem with disjoint booking and service horizons and apply a statistical learning approach that imposes structural constraints on policy updates to preserve both properties, which facilitates the learning process. Because we prove that these properties hold in general, such an exploitation is possible for any i-DMVRP, including problems with overlapping horizons based on the modified model (Appendix D). Likewise, our results imply that the opportunity cost approximation approach by Adelman (2007), which is based on linear programming and includes constraints exploiting domain knowledge, can also be validly transferred to i-DMVRP solution approaches.

In contrast to Property 3 and Property 4, to the best of our knowledge, the targeted algorithmic exploitation of the decomposability of opportunity cost into DPC and MCTS (Property 1 and Property 2) has not yet been proposed by existing research. Hence, in this computational study, we systematically explore the potential of three general approaches for exploiting the decomposability. In Section 5.1, we first present the three approaches. In Section 5.2, we then describe the design of the computational study, that is, we define the specific i-DVPRP and the settings we consider. Finally, in Section 5.3, we discuss the computational results and derive insights regarding the potential of each of the three approaches.

## 5.1 General Algorithmic Approaches Exploiting Decomposability

In the following, we explain the algorithmic approaches we analyze in the computational experiments. Because we are interested in the general potential of exploiting decomposability in a certain way rather than in a specific (heuristic) algorithm design, we define the algorithms by formulating the respective variant of the Bellman equation and solve it by means of backwards recursion. Thus, we obtain idealized algorithms by combining each of the approximation approaches with an exact method to compute the values of the approximation. For an in-depth review on existing heuristic solution approaches for specific i-DMVRPs see Fleckenstein et al. (2023).

### 5.1.1 Single-Component Approximations

This approach is based on the idea that, in certain problem settings, one of the opportunity cost components may generally have a considerably greater absolute value than the other. This suggests that completely neglecting the other component in the opportunity cost approximation should be possible with only a small deterioration of performance.

We derive an idealized *DPC-based approximation*  $\Delta \tilde{R}_t(s_{t-1}, c)$  from the following Bellman equation:

$$\tilde{R}_{t-1}(s_{t-1}) = \sum_{c \in \mathcal{C}} \lambda_c^t \cdot \max_{g_t \in \mathcal{G}(s_{t-1},c)} \left( g_t \cdot \left( r_c - \Delta \tilde{R}_t(s_{t-1},c) \right) \right) + \tilde{R}_t'(s_t'(0))$$
(33)

with:

$$r_{\phi_t(s_t')} = 0 \ \forall \ t \in \{1, \dots, T\}.$$
(34)

Here, we set the logistics-related rewards equal to 0 while retaining the revenue as the immediate reward of an acceptance decision. Thus, decisions are optimized comparing the immediate reward, that is, the potential immediate revenue, with the exact future revenue impact of the acceptance decision neglecting its future cost impact.

Similarly, an idealized MCTS-based approximation  $\Delta \tilde{F}_t(s_{t-1}, c)$  results from:

$$\widetilde{F}_{t-1}(s_{t-1}) = \sum_{c \in C} \lambda_c^t \cdot \operatorname*{argmax}_{g_t \in \mathcal{G}(s_{t-1},c)} \left( g_t \cdot \left( r_c - \Delta \widetilde{F}_t(s_{t-1},c) \right) \right) \cdot \left( -\Delta \widetilde{F}_t(s_{t-1},c) \right) \\
+ \widetilde{F}_t'(s_t'(0)).$$
(35)

By using an  $\operatorname{argmax}(\cdot)$  operator in Equation (35), we still consider the revenue for decisionmaking but at the same time prevent it from entering the value function as an actual immediate reward. Thereby, we make sure that  $\Delta \tilde{F}_t(s_{t-1}, c)$  only captures logistics-related rewards, and thus, correctly quantifies the future cost impact of an acceptance decision. This cost impact is then compared with the revenue to derive the currently optimal decision-making in each stage.

### 5.1.2 Hybrid Reward Approximations

Compared with single-component policies, an approximation aiming at both components has structural advantages if both MCTS and DPC have non-negligible values. Many existing approximation approaches of this type are based on value function approximations that return an aggregated estimate such that the components are not distinguishable (e.g., Ulmer, 2020). Alternatively, the decomposability property allows hybrid approximation design. The underlying idea is to compute a separate estimate for DPC and MCTS, aggregate the estimates, and therewith, use both as an input for demand control decision-making. A new sampling-based solution method recently presented by Abdollahi et al. (2023) constitutes a first non-learning-based, step in this direction. They use tentative route planning with sampled orders to compute a DPC estimate and an MCTS estimate separately.

In the class of learning-based approaches, the equivalent are hybrid reward architectures that have been successfully applied to classical reinforcement learning problems (Van Seijen et al., 2017), but, to the best of our knowledge, have not yet been applied to i-DMVRPs. In such approaches, a separate value function approximation is computed for revenue and routing cost, that is, two single-component opportunity cost approximations result, each of which can depend on different features such that learning is expedited. To aggregate the two estimates, several strategies can be used, like a simple additive aggregation (Van Seijen et al., 2017) or a more complex delegation architecture (Russell and Zimdars, 2003).

Because Property 1 suggests that i-DMVRPs lend themselves to hybrid reward architectures, we evaluate the potential of such opportunity cost approximation approaches for i-DMVRPs by applying two idealized implementations: The first, rather naive one, bases on the aggregation of the two single-component approximations by simply summing up their value functions (33) and (35), that is, the corresponding opportunity cost approximation is set to  $\Delta \tilde{R}_t(s_{t-1}, c) + \Delta \tilde{F}_t(s_{t-1}, c)$ . We refer to this approach as naive hybrid reward approximation (naive HR approximation). The other, more sophisticated, approach relies on the additive aggregation of an offline learned component and an online learned component. More precisely, we investigate whether it is promising to approximate the DPC offline and the MCTS online, or vice versa. In the existing literature, offline-online learning as a general concept has already been successfully applied (e.g., Ulmer et al., 2019). However, no approach has been proposed yet that specifically takes advantage of opportunity cost decomposability. For the offline approximation, we draw on the (disaggregate) approximation of the respective single-component approximation ((33) and (35)) for each state, and store them in a look-up table by averaging the respective estimates over a two-dimensional state space representation. The first dimension measures the remaining time in the booking process, and the second dimension represents the remaining logistical capacity. For the online approximation, we use the disaggregate DPC (MCTS) value from the original model ((15) and (16)). Then, we set the corresponding opportunity cost approximation equal to  $\Delta \tilde{F}_t^{aggr}(s_{t-1}, c) +$  $\Delta R_t(s_{t-1}, c)$  (or  $\Delta \tilde{R}_t^{aggr}(s_{t-1}, c) + \Delta F_t(s_{t-1}, c)$ , respectively). We refer to this approach as offline-online hybrid reward approximation (offline-online HR approximation) and distinguish its variants by using the term DPC (MCTS) when the DPC (MCTS) are approximated online.

### 5.1.3 Compact Overview

In summary, this computational study compares decision-making based on the following approximations, which include the ones described above and the true opportunity cost derived from solving the dynamic program (9) as the benchmark:

- DPC-based approximation:  $\Delta \tilde{V}_t(s_{t-1}, c) = \Delta \tilde{R}_t(s_{t-1}, c)$
- MCTS-based approximation:  $\Delta \tilde{V}_t(s_{t-1}, c) = \Delta \tilde{F}_t(s_{t-1}, c)$
- Naive hybrid reward approximation:  $\Delta \tilde{V}_t(s_{t-1}, c) = \Delta \tilde{R}_t(s_{t-1}, c) + \Delta \tilde{F}_t(s_{t-1}, c)$
- Offline-online HR approximation (DPC-based):  $\Delta \tilde{V}_t(s_{t-1}, c) = \Delta \tilde{F}_t^{aggr}(s_{t-1}, c) + \Delta R_t(s_{t-1}, c)$
- Offline-online HR approximation (MCTS-based):  $\Delta \tilde{V}_t(s_{t-1}, c) = \Delta \tilde{R}_t^{aggr}(s_{t-1}, c) + \Delta F_t(s_{t-1}, c)$
- True opportunity cost (Benchmark):  $\Delta \tilde{V}_t(s_{t-1}, c) = \Delta V_t(s_{t-1}, c)$

## 5.2 Study Design and Methodology

To evaluate the approaches introduced above, we apply them for decision-making in various different settings. All settings have in common that they reflect an i-DMVRP with the same problem structure as introduced in Section 3.1 and the following additional assumptions: disjoint booking and service horizons, a single fulfillment vehicle, pure accept/reject decisions, and a booking horizon of T = 10 potential decision epochs. With that, we ensure that the instance size is sufficiently small for being computationally tractable without overly compromising complexity.

The settings differ with regard to four parameters whose realizations emulate characteristics of i-DMVRPs from typical application areas. Two of those parameters define a setting's *customer distribution*, one sets the general *profitability* of a setting, and the last defines the type of *capacity consumption* considered in a setting. In the following, we further describe those parameters and motivate our choice of their realizations.

### 5.2.1 Customer Distribution

The customer distribution of a setting is characterized by two parameter values, which are the customers' *location distribution* and their *revenue distribution*, and is further defined by the mutual interplay of those two parameters.

Location distribution – To vary the level of difficulty of demand consolidation, we draw the customer locations  $l_c$  from two different customer distributions on a line segment of 50 *length units* (*LU*) in the interval [-25,25], with a centrally located depot. The first customer distribution follows a uniform distribution over the entire interval and, hence, mimics an urban area. The second customer distribution is drawn from two truncated normal distributions with means -10 and 20 and the same standard deviation of 2.5 *LU*, from which we draw 50% of the customer locations each. Therewith, we obtain two clusters that could represent two villages in a rural area. **Table 1** Customer distribution parameters

revenue distribution –	location distribution			
	uniform (unif)	clustered (clust)	clustered sorted (clust_sort)	
homogeneous (homog)	$\checkmark$	$\checkmark$		
high-before-low (h-b-l)	$\checkmark$	$\checkmark$	$\checkmark$	
low-before-high (l-b-h)	$\checkmark$	$\checkmark$	$\checkmark$	
random (rand)	$\checkmark$	$\checkmark$	$\checkmark$	

*Revenue distribution* – The customers' revenue distribution is an important characteristic that influences displacement effects, both in a spatial and a temporal sense. Thus, on the one hand, we consider *homogeneous revenues* (no additional displacement effects) with a value of  $r_c = 15$ *monetary units* (*MU*). This corresponds to, for example, next-day parcel delivery with a static, uniform delivery fee. On the other hand, we consider heterogeneous settings with 70% low-revenue customers ( $r_c = 15MU$ ) and 30% high-revenue customers ( $r_c = 25MU$ ). We vary their distribution over time as follows: The sequence of customer arrivals either follows a strict *high-before-low* sorting (low displacement effects), a *random* sorting (medium displacement effects), or a strict *low-before-high* sorting (high displacement effects). Such variations over time can occur, for example, as a result of markup-pricing or markdown-pricing. In addition, for each of those three schemes, we consider a customer setting in which the high-revenue customers only originate from the distant cluster to mimic a *distance-based* pricing scheme, which is common, for example, in mobility-on-demand applications. We refer to this special combination of location distribution and revenue distribution as *clustered sorted*.

Overall, these two parameters and the combinations of their potential realizations yield 11 different customer settings, which are marked with a  $\checkmark$  in Table 1. As mentioned before, for all those 11 customer settings, we vary two more parameters that address the general profitability of a setting as well as the capacity consumption.

# 5.2.2 Profitability

To vary the general level of profitability of a problem setting, we adjust the relation between revenues and cost by solving all previously mentioned settings for three different routing cost factors. More precisely, we consider low-cost settings with cost of 0.2 MU per LU travelled, medium-cost settings with 0.6 MU, and high-cost settings with 1 MU. Thereby, the different profitability settings represent different fields of application. Low-cost settings are dominant in attended home delivery due to minimum order values that cause high revenue relative to cost. On the contrary, high-cost settings occur in mobility-on-demand applications with comparatively low revenues equal to the fare the provider charges.

### 5.2.3 Capacity Consumption

To consider the impact of the marginal capacity consumption per order for all our settings, we additionally assume another parameter. Its value either reflects *route length constraints*, or *physical capacity constraints* such that only one of them is restrictive in a setting. First, for the settings with restrictive physical capacity, we set the maximum capacity to 3 orders and assume unit demand for all orders. Hence, the marginal capacity consumption is both uniform and known apriori. A typical application that is represented by these settings is attended home delivery of bulky goods. Second, for settings with a restrictive route length, we set the maximum capacity to 50 *LU*. With this type of constraint, the marginal capacity consumption is variable among the orders and is unknown until the final routing decision. This is typical for applications that are mainly time-constrained such as same-day delivery.

Because we aim at a full-factorial analysis of our approaches, we solve all the 11 customer settings from Table 1 for all 6 combinations of parameter values defining the profitability and capacity consumption. Hence, we test our approaches in 66 settings. For each setting, we draw 50 instances. Thereby, for each instance, we draw a fixed (deterministic) customer stream of 10 customers from the setting-dependent customer distribution. Then, for each customer, we assume a probability of  $\lambda_0^t = 0.5$  that the respective request does not arrive, which is the only source of stochasticity, once a setting is defined.



(a) Physical capacity-constrained settings





Fig. 2 Objective values resulting from the different opportunity cost approximations – Averaged across 50 instances per setting

### 5.3 Performance Evaluation

To evaluate the performance of the considered approaches, for each of the 66 settings, we calculate the mean objective value, that is, profit after fulfillment, over the 50 different instances drawn. In Fig. 2 and Appendix F, we report the results and discuss them in the following:

### 5.3.1 Single-Component Approximation

The comparison between the two variants of this approximation approach reveals that the DPCbased variant outperforms the MCTS-based variant in high-profitability settings and most route length-constrained settings. Especially in the latter settings, it also achieves very small optimality gaps and is competitive with the more sophisticated approximations. In turn, the MCTS-based variant is among the best-performing approaches in low-profitability, physical capacity-constrained settings. On the downside, the performance of both variants fluctuates strongly across different settings. For example, we observe a very bad performance of the MCTS-based variant for low-profitability settings with distance-dependent revenues and even negative objective values for the DPC-based variant in some low-profitability, physical capacity-constrained settings. However, despite these issues with solution quality fluctuations, single-component approximations can be a viable approach because their practical implementations usually require less computational effort and, if the right variant is applied in the right setting, it can offer competitive solution quality.

### 5.3.2 Naive Hybrid Reward Approximation

The naive hybrid reward approximation is among the best-performing policies in many highprofitability settings and, with some outliers, also in medium-profitability settings. Hence, in these settings, ignoring the interdependency between DPC and MCTS values does not appear particularly harmful to the solution quality. However, it performs weaker and shows severe outliers in low-profitability settings, especially if capacity is route length-constrained. Because of this lack of robustness and the small gains relative to single-component approximations, the additional computational effort may only be justified in a few specific settings.

### 5.3.3 Offline-Online HR Approximation

Comparing the two variants of the offline-online HR approximation, the MCTS-based variant performs superior in physical capacity-constrained settings and in low-profitability settings. With a few exceptions, it is even the best-performing approach for these settings and shows a very robust performance overall. The DPC-based variant is slightly better in some route length-constrained settings with medium or high profitability. However, it can be considered inferior overall because its performance is also more variable. A possible explanation for this result is that the MCTS show a stronger variation over similar states, whereas the differences in DPC are smaller for states with the same decision epoch and remaining capacity. If this is indeed the case, anticipating the MCTS online in disaggreate form should yield more accurate opportunity cost

estimates. Overall, the results are very promising for the MCTS-based variant, even though its practical implementations are expected to require the highest computational effort.

We can conclude that none of the presented approaches exploiting the decomposability of opportunity cost is strictly dominated such that all are worth being investigated further. Our results can serve as a rough guidance as to which approach is most promising in a certain setting. Interestingly, we also observe that the relevance of DPC and MCTS correlates with the performance of the approaches. As an example, the MCTS gain relevance with decreasing profitability because the fulfillment cost becomes larger relative to the revenues, and thus, we would expect a relative performance gain of the MCTS-based approaches. This gain can indeed be observed for both the MCTS-based single component approximation and the MCTS-based offline-online HR approximation.

# 6 Conclusion and Future Research Opportunities

This work constitutes the first formal, generic analysis of opportunity cost in i-DMVRPs. We showed that the original interpretation of opportunity cost from traditional revenue management applications cannot be transferred to i-DMVRPs and, therefore, generalized its definition. Further, we analytically investigated opportunity cost properties with the central property being the decomposability into DPC and MCTS. Finally, we conducted a computational study and applied previously unconsidered approximation approaches as a proof of concept for that the properties can be directly exploited in algorithm design. In the following, we first briefly summarize our theoretical and computational results. Second, we discuss the future research opportunities emerging from our work.

In existing works, insights are mainly derived either from qualitative reasoning and computational studies (e.g., Mackert, 2019 and Vinsensius et al., 2020) or from analytical analyses for specific i-DMVRPs (Asdemir et al., 2009 and Lebedev et al., 2020). In contrast, our approach is both analytical and generic, which is why our results apply to the whole family of the most common types of i-DMVRPs. By showing that it is possible to substantiate existing observations at the modeling level, we not only confirm their general validity, but also highlight the importance of modeling frameworks for understanding and exploiting the problem structure of i-DMVRPs when designing solution concepts. The following summary captures the essence of our theoretical framework:

For i-DMVRPs, opportunity cost measures the impact of selling an order as a consequence of a demand control decision, on both (expected) future revenues, in the form of DPC, and (expected) future cost of fulfillment, in the form of MCTS (e.g., Yang and Strauss, 2017). With our analysis, we can precisely express this impact in a formal way: First, we show that the impact of a demand control decision can be isolated from the impact of the subsequent routing decision. Second, we show that both revenue impacts and cost impacts can be formally expressed and, thus, mathematically decomposed from each other (Property 1). Third, we find that both impacts can have a

positive or a negative sign (Property 2). Fourth, although we can measure the impacts in isolation through MCTS and DPC, the corresponding values are still nonseparable and, in sum, nonnegative, which leads to the general non-negativity of opportunity cost (Property 3). Finally, the value function decreases monotonically in an increasing set of accepted but not yet served customer orders and can also be transformed to decrease monotonically in time (Property 4).

In addition to these theoretical findings, our computational study shows that exploiting the decomposability of opportunity cost allows transferring new algorithmic approaches to i-DMVRPs that have been shown to be beneficial in related fields. This includes different hybrid reward approximation approaches, which are already established in reinforcement learning (Van Seijen et al., 2017). The computational evaluation of idealized implementations of these approaches illustrates their potential in the context of i-DMVRPs. At the same time, we also observe that the relative performance of the approaches can vary considerably even among similar settings. This highlights the importance of continuously expanding the available toolbox of approaches such that a wide range of settings can be suitably tackled.

For future research, we primarily see the opportunity to exploit the decomposability and the other properties in heuristic solution algorithms for specific i-DMVRPs, not only with the aim of improving solution quality but also for reducing runtimes. In particular, we believe that heuristic versions of the presented hybrid reward approximation approaches deserve further investigation. Another future research question arises in connection with Property 4. In the case of problems with overlapping horizons, it only holds for modified models (Appendix D). Interestingly, this model transformation can be viewed as a form of reward shaping (Laud, 2004), which establishes monotonicity at the cost of increasing the delay of rewards. To the best of our knowledge, reward shaping has not been applied to solve i-DMVRPs. Hence, although it is out of scope for the study at hand, our theoretical results suggest that there is potential to investigate its application.

Finally, we see the potential for doing similar theoretical research for other novel demand management problems that feature integrated combinatorial optimization problems to plan service fulfillment, such as scheduling problems (Xu et al., 2015).

# References

- Abdollahi, M., Yang, X., Nasri, M. I., & Fairbank, M. (2023). Demand management in timeslotted last-mile delivery via dynamic routing with forecast orders. *European Journal of Operational Research*, 309(2), 704-718. https://doi.org/10.1016/j.ejor.2023.01.023.
- Adelman, D. (2007). Dynamic bid prices in revenue management. *Operations Research*, 55(4), 647-661. https://doi.org/10.1287/opre.1060.0368.
- Agatz, N. A. H., Campbell, A. M., Fleischmann, M., van Nunen, J., & Savelsbergh, M. W. P. (2013). Revenue management opportunities for internet retailers. *Journal of Revenue and Pricing Management*, 12(2), 128-138. https://doi.org/10.1057/rpm.2012.51.
- Akkerman, F., Mes, M., & Lalla-Ruiz, E. (2022). Dynamic time slot pricing using delivery costs approximations. In *International Conference on Computational Logistics*, 214-230. https://doi.org/10.1007/978-3-031-16579-5 15.
- Al-Kanj, L., Nascimento, J., & Powell, W. B. (2020). Approximate dynamic programming for planning a ride-hailing system using autonomous fleets of electric vehicles. *European Jour*nal of Operational Research, 284(3), 1088-1106. https://doi.org/10.1016/j.ejor.2020.01.033.
- 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.
- Asdemir, K., Jacob, V. S., & Krishnan, R. (2009). Dynamic pricing of multiple home delivery options. *European Journal of Operational Research*, 196(1), 246-257. https://doi.org/10.1016/j.ejor.2008.03.005.
- Atasoy, B., Ikeda, T., Song, X., & Ben-Akiva, M. E. (2015). The concept and impact analysis of a flexible mobility on demand system. *Transportation Research Part C: Emerging Technol*ogies, 56, 373-392. https://doi.org/10.1016/j.trc.2015.04.009.
- Avraham, E., & Raviv, T. (2021). The steady-state mobile personnel booking problem. Transportation Research Part B: Methodological, 154, 266-288. https://doi.org/10.1016/j.trb.2021.10.008.
- Azi, N., Gendreau, M., & Potvin, J.-Y. (2012). A dynamic vehicle routing problem with multiple delivery routes. *Annals of Operations Research*, 199(1), 103-112. https://doi.org/10.1007/s10479-011-0991-3.
- Bertsimas, D., Jaillet, P., & Martin, S. (2019). Online vehicle routing: The edge of optimization in large-scale applications. *Operations Research*, 67(1), 143-162. https://doi.org/10.1287/opre.2018.1763.
- Campbell, A. M., & Savelsbergh, M. W. P. (2005). Decision support for consumer direct grocery initiatives. *Transportation Science*, *39*(3), 313-327. https://doi.org/10.1287/trsc.1040.0105.
- Campbell, A. M., & Savelsbergh, M. W. P. (2006). Incentive schemes for attended home delivery services. *Transportation Science*, 40(3), 327-341. https://doi.org/10.1287/trsc.1050.0136.
- Fleckenstein, D., Klein, R., & Steinhardt, C. (2023). Recent advances in integrating demand management and vehicle routing: A methodological review. *European Journal of Operational Research*, 306(2), 499-518. https://doi.org/10.1016/j.ejor.2022.04.032.
- Gallego, G., & Topaloglu, H. (2019). Revenue management and pricing analytics. New York, NY: Springer. https://doi.org/10.1007/978-1-4939-9606-3.
- Hosni, H., Naoum-Sawaya, J., & Artail, H. (2014). The shared-taxi problem: Formulation and solution methods. *Transportation Research Part B: Methodological*, 70, 303-318. https://doi.org/10.1016/j.trb.2014.09.011.
- Klein, R., Koch, S., Steinhardt, C., & Strauss, A. K. (2020). A review of revenue management: Recent generalizations and advances in industry applications. *European Journal of Operational Research*, 284(2), 397-412. https://doi.org/10.1016/j.ejor.2019.06.034.
- Klein, R., Mackert, J., Neugebauer, M., & Steinhardt, C. (2018). A model-based approximation of opportunity cost for dynamic pricing in attended home delivery. *OR Spectrum*, *40*(4), 969-996. https://doi.org/10.1007/s00291-017-0501-3.

- Klein, V., & Steinhardt, C. (2023). Dynamic demand management and online tour planning for same-day delivery. *European Journal of Operational Research*, 307(2), 860-886. https://doi.org/10.1016/j.ejor.2022.09.011.
- Koch, S. (2017). Least squares approximate policy iteration for learning bid prices in choicebased revenue management. *Computers & Operations Research*, 77, 240-253. https://doi.org/10.1016/j.cor.2016.07.015.
- Koch, S., & Klein, R. (2020). Route-based approximate dynamic programming for dynamic pricing in attended home delivery. *European Journal of Operational Research*, 287(2), 633-652. https://doi.org/10.1016/j.ejor.2020.04.002.
- Kullman, N. D., Cousineau, M., Goodson, J. C., & Mendoza, J. E. (2022). Dynamic ride-hailing with electric vehicles. *Transportation Science*, 56(3), 775-794. https://doi.org/10.1287/trsc.2021.1042.
- Lang, M. A. K., Cleophas, C., & Ehmke, J. F. (2021). Anticipative dynamic slotting for attended home deliveries. *Operations Research Forum*, 2(4), 70. https://doi.org/10.1007/s43069-021-00086-9.
- Laud A. D. (2004). Theory and application of reward shaping in reinforcement learning. Doctoral dissertation, University of Illinois at Urbana-Champaign.
- Lebedev, D., Goulart, P., & Margellos, K. (2021). A dynamic programming framework for optimal delivery time slot pricing. *European Journal of Operational Research*, 292(2), 456-468. https://doi.org/10.1016/j.ejor.2020.11.010.
- Lebedev, D., Margellos, K., & Goulart, P. (2020). Approximate dynamic programming for delivery time slot pricing: A sensitivity analysis. Working paper, University of Oxford. https://doi.org/10.48550/arXiv.2008.00780.
- Lee, T. C., & Hersh, M. (1993). A model for dynamic airline seat inventory control with multiple seat bookings. *Transportation Science*, *27*(3), 252-265. https://doi.org/10.1287/trsc.27.3.252.
- Li, Y., Archetti, C., & Ljubic, I. (2024). Emerging optimization problems for distribution in sameday delivery. Working paper, ESSEC Business School. https://doi.org/10.48550/arXiv.2405.05620.
- Mackert, J. (2019). Choice-based dynamic time slot management in attended home delivery. *Computers* & *Industrial Engineering*, *129*, 333-345. https://doi.org/10.1016/j.cie.2019.01.048.
- Maddah, B., Moussawi-Haidar, L., El-Taha, M., & Rida, H. (2010). Dynamic cruise ship revenue management. *European Journal of Operational Research*, 207(1), 445-455. https://doi.org/10.1016/j.ejor.2010.03.043.
- Powell, W. B. (2011). Approximate dynamic programming: Solving the curses of dimensionality. 2nd ed. Hoboken, NJ: John Wiley & Sons. https://doi.org/10.1002/9781118029176.
- Powell, W. B. (2022). Reinforcement learning and stochastic optimization: A unified framework for sequential decisions. Hoboken, NJ: John Wiley & Sons. https://doi.org/10.1002/9781119815068.
- Powell, W. B., Simao, H. P., & Bouzaiene-Ayari, B. (2012). Approximate dynamic programming in transportation and logistics: a unified framework. *EURO Journal on Transportation and Logistics*, 1(3), 237-284. https://doi.org/10.1007/s13676-012-0015-8.
- Prokhorchuk, A., Dauwels, J., & Jaillet, P. (2019). Stochastic dynamic pricing for same-day delivery routing. Working paper, Nanyang Technological University. https://doi.org/10.48550/arXiv.1912.02946.
- Quante, R., Fleischmann, M., & Meyr, H. (2009). A stochastic dynamic programming approach to revenue management in a make-to-stock production system. Working paper, Vienna University of Economics and Business.
- Russell, S. J., & Zimdars, A. (2003). Q-decomposition for reinforcement learning agents. In *Proceedings of the 20th international conference on machine learning (ICML-03)*, 656-663.

- Snoeck, A., Merchán, D., & Winkenbach, M. (2020). Revenue management in last-mile delivery: state-of-the-art and future research directions. *Transportation Research Procedia*, 46, 109-116. https://doi.org/10.1016/j.trpro.2020.03.170.
- Soeffker, N., Ulmer, M. W., & Mattfeld, D. C. (2022). Stochastic dynamic vehicle routing in the light of prescriptive analytics: A review. *European Journal of Operational Research*, 298(3), 801-820. https://doi.org/10.1016/j.ejor.2021.07.014.
- Strauss, A. K., Gülpinar, N., & Zheng, Y. (2021). Dynamic pricing of flexible time slots for attended home delivery. *European Journal of Operational Research*, 294(3), 1022-1041. https://doi.org/10.1016/j.ejor.2020.03.007.
- Strauss, A. K., Klein, R., & Steinhardt, C. (2018). A review of choice-based revenue management: Theory and methods. *European Journal of Operational Research*, 271(2), 375-387. https://doi.org/10.1016/j.ejor.2018.01.011.
- Subramanian, J., Stidham Jr, S., & Lautenbacher, C. J. (1999). Airline yield management with overbooking, cancellations, and no-shows. *Transportation Science*, *33*(2), 147-167. https://doi.org/10.1287/trsc.33.2.147.
- Talluri, K. T., & Van Ryzin, G. J. (2004). The theory and practice of revenue management. New York, NY: Springer. https://doi.org/10.1007/b139000.
- Ulmer, M. W. (2020). Dynamic pricing and routing for same-day delivery. *Transportation Science*, 54(4), 1016-1033. https://doi.org/10.1287/trsc.2019.0958.
- Ulmer, M. W., Goodson, J. C., Mattfeld, D. C., & Hennig, M. (2019). Offline-online approximate dynamic programming for dynamic vehicle routing with stochastic requests. *Transportation Science*, 53(1), 185-202. https://doi.org/10.1287/trsc.2017.0767.
- Ulmer, M. W., Goodson, J. C., Mattfeld, D. C., & Thomas, B. W. (2020). On modeling stochastic dynamic vehicle routing problems. *EURO Journal on Transportation and Logistics*, 9(2), 100008. https://doi.org/10.1016/j.ejtl.2020.100008.
- Van Seijen, H., Fatemi, M., Romoff, J., Laroche, R., Barnes, T., & Tsang, J. (2017). Hybrid reward architecture for reinforcement learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 5398-5408.
- Vinsensius, A., Wang, Y., Chew, E. P., & Lee, L. H. (2020). Dynamic incentive mechanism for delivery slot management in e-commerce attended home delivery. *Transportation Science*, 54(3), 567-587. https://doi.org/10.1287/trsc.2019.0953.
- Waßmuth, K., Köhler, C., Agatz, N., & Fleischmann, M. (2023). Demand management for attended home delivery – A literature review. *European Journal of Operational Research*, 311(3), 801-815. https://doi.org/10.1016/j.ejor.2023.01.056.
- Weatherford, L. R., & Bodily, S. E. (1992). A taxonomy and research overview of perishableasset revenue management: Yield management, overbooking, and pricing. *Operations Re*search, 40(5), 831-844. https://doi.org/10.1287/opre.40.5.831.
- Xu, L., Wang, Q., & Huang, S. (2015). Dynamic order acceptance and scheduling problem with sequence-dependent setup time. *International Journal of Production Research*, 53(19), 5797-5808. https://doi.org/10.1080/00207543.2015.1005768.
- Yang, X., & Strauss, A. K. (2017). An approximate dynamic programming approach to attended home delivery management. *European Journal of Operational Research*, 263(3), 935-945. https://doi.org/10.1016/j.ejor.2017.06.034.
- Yang, X., Strauss, A. K., Currie, C. S. M., & Eglese, R. (2016). Choice-based demand management and vehicle routing in e-fulfillment. *Transportation Science*, 50(2), 473-488. https://doi.org/10.1287/trsc.2014.0549.

# **Appendix A: Notation**

Table 2 Notation (continued on next page)

$t = 1, \dots, T$	Decision epoch
С	Type of customer request
C <sub>t</sub>	Customer requesting in stage t
С	Set of customer request types
$\lambda_c^t$	Arrival rate of a customer request from customer order type $c$ in stage $t$
l <sub>c</sub>	Location of a customer request from customer request type $c$
$r_c$	Revenue of a customer request from customer request type $c$
τ	Time of an incoming customer request
0	Fulfillment option chosen by a customer of a confirmed customer order
$(l_c, \tau, o)$	Confirmed customer order that is defined by parameters $l_c$ , $\tau$ , and $o$
$\mathcal{C}_t$	Set of customer orders that have been confirmed but for which fulfillment has not yet started until decision epoch $t$
v	Vehicle
ν	Set of vehicles
$\theta_t^v$	Tour of vehicle $v$ currently running in decision epoch $t$
$\phi_t = \{\theta_t^v\}_{v \in \mathcal{V}}$	Overall tour plan in decision epoch t
$\Psi(\phi_t)$	Subset of orders for which the fulfillment process has started according to the new tour plan $\phi_t$
$s_t = (\mathcal{C}_t, \phi_t)$	State in decision epoch t
$s_t' \mid s_{t-1}, c, g_t$	Interim state that is reached when in state $s_{t-1}$ a customer request of type <i>c</i> arrives and demand control decision $g_t$ is taken
$s_t'(c)$	Interim state that is reached when in state $s_{t-1}$ a customer request of type <i>c</i> arrives and is accepted by demand control decision $g_t = 1$
$s_t'(0)$	Interim state that is reached when in state $s_{t-1}$ a customer request of type <i>c</i> arrives and is rejected by demand control decision $g_t = 0$ or in case there is no request at decision epoch <i>t</i>
$g_t$	Demand control decision in decision epoch <i>t</i>
$\mathcal{G}(s_{t-1},c) \subseteq \{0,1\}$	Feasible demand control decisions (1: accept, 0: reject) in state $s_{t-1}$ , when a customer request of type <i>c</i> arrives
$\phi_t(g_t)$	Tour planning decision in decision epoch $t$ , depending on demand control decision $g_t$
$\phi_t^*(g_t)$	Optimal tour planning decision in decision epoch $t$ , depending on demand control decision $g_t$
$r_{\phi_t(g_t)}$	Logistics-related rewards depending on tour planning decision $\phi_t(g_t)$
$r_{\phi_t^*(g_t)}$	Logistics-related rewards depending on optimal tour planning decision $\phi_t^*(g_t)$
$\Phi(s_{t-1}, c, g_t)$	Set of tour plans that are feasible given state $s_{t-1}$ and the demand control decision $g_t$ for the arriving customer request of type $c$
$a_t = (g_t, \phi_t(g_t))$	Action taken at decision epoch t
$\mathcal{A}(s_{t-1}, c)$	Action space in state $s_{t-1}$ for a request arrival of type $c$
x	A policy
π	A certain decision sequence

$\pi^*(c_t)$	The optimal decision sequence when starting in interim state $s'_t(c_t)$
ω	A certain sample path
Ω	Set of all potential sample paths
$V_t(s_t)$	Value of being in state $s_t$
$\tilde{V}_t(s_t)$	Value of being in state $s_t$ of the modified model
$V_t'(s_t')$	Value of being in interim state $s'_t$
$V_t^{\prime\omega}(s_t^\prime)$	Value of being in interim state $s'_t$ when sample path $\omega$ realizes
$V_t^{\prime\pi\omega}(s_t^\prime)$	Value of being in interim state $s'_t$ when sample path $\omega$ realizes and decision sequence $\pi$ is applied
$\Delta V_t(s_{t-1},c)$	Opportunity cost of accepting a request of type $c$ starting in state $s_{t-1}$
$R'_{t-1}(s'_{t-1})$	Expected future revenues of interim state $s'_{t-1}$
$\Delta R_t(s_{t-1},c)$	DPC of accepting a customer request of type $c$ at decision epoch $t$ and state $s_{t-1}$
$F_{t-1}'(s_{t-1}')$	Expected future fulfillment cost of interim state $s'_{t-1}$
$\Delta F_t(s_{t-1}, c)$	MCTS of accepting a customer request of type <i>c</i> at decision epoch <i>t</i> and state $s_{t-1}$

# **Appendix B: Examples**



Fig. 3 Customer locations of the problem instance underlying the discussion of opportunity cost properties in i-DMVRPs

For illustrative reasons and in order to generate a general intuition of opportunity cost in i-DMVRPs and its decisive characteristics, we consider a simple instance of the generic i-DMVRP presented in Section 3.1, which we introduce in the following. However, the respective results can be generalized to more complex problem instances. The problem instance corresponds to a special case of the generic i-DMVRP with disjoint booking and service horizons. Thus, the provider takes a single routing decision at the terminal decision epoch T and demand control decisions in t = 1, ..., T. We encode the accept decision by actions  $a_t = g_t = 1$  and the reject decisions by actions  $a_t = g_t = 0$ . We assume that there are only three potential customer requests, denoted as  $c_1, c_2$ , and  $c_3$ . As depicted in Fig. 3, they are located on a line with one single, centrally located depot. There are three decision epochs, and thus, three stages in which customer requests arrive with time-dependent arrival rates  $\lambda_{c_i}^t = 0.5$ , if i = t and  $\lambda_{c_i}^t = 0$ , else. The potential revenues  $r_{c_i}$  associated with those customer requests are 10, 10 and 20 monetary units (MU) for requests  $c_1$ ,  $c_2$ , and  $c_3$ , respectively. The customer request characteristics are summarized in Table 3. Further, for every customer request, the same physical capacity consumption is assumed. It equals the size of one trunk. To serve customer orders, the provider has a single vehicle available, which can only load two trunks at a time and is not allowed to conduct multiple trips. Fuel cost is assumed to equal 1 MU per length unit (LU).

$c_i$	$l_{c_i}$	$r_{c_i}$	$\lambda_{c_i}^t$
<i>C</i> <sub>1</sub>	(-4)	10	
<i>C</i> <sub>2</sub>	(4.5)	10	$\lambda_{c_i}^t = \begin{cases} 0.5 & \text{if } i = t \\ 0 & else \end{cases}$
<i>C</i> <sub>3</sub>	(-5.5)	20	

Table 3 Customer requests of the problem instance underlying the discussion of opportunity cost properties in i-DMVRPs

### B.1. Example 1

Example 1: We investigate decision epoch t = 2 of the above described problem instance. We assume that a customer request  $c_1$  arrived at the previous decision epoch and turned into a customer order. In t = 2, customer request  $c_2$  realizes. Thus,  $C_1 = \{c_1\}$  and the provider has to decide whether to accept the current customer request  $c_2$  with action  $a_2 = 1$  or reject it with action  $a_2 = 0$ . If the provider accepts  $c_2$ , it turns into a confirmed customer order. Then, it is not possible to also accept the customer request  $c_3$ , which realizes at the subsequent decision epoch with probability  $\lambda_{c_3}^3 = 0.5$ . Consequently, decision  $a_2 = 1$  results in expected DPC that equals  $\lambda_{c_3}^3 \cdot r_{c_3} MU = 0.5 \cdot 20 MU = 10 MU$ . This means that an expected revenue of 10 MU is displaced due to limited vehicle capacities if decision  $a_2 = 1$  is taken.

## B.2. Example 2

Example 2a: Again, we consider the same decision epoch, with equal state and potential actions of the problem instance as described in Example 1. If fuel cost were to be neglected, DPC would equal opportunity cost, i.e.,  $\Delta V_2(s_1, c_2) = 10 MU$ . Since the immediate contribution of action  $a_2 = 1$  also equals  $r_{c_2} = 10 MU$ , both decisions,  $a_2 = 1$  or  $a_2 = 0$ , are equally good decisions for the provider.

Example 2b: However, Fig. 3 shows clearly that the additional fulfillment cost in case the provider accepts customer request  $c_2$  equals 9MU. In turn, rejecting customer request  $c_2$  by action  $a_2 = 0$  and then accepting customer request  $c_3$  instead only leads to additional fulfillment cost of 3MU. Since customer request  $c_3$  realizes with probability  $\lambda_{c_3}^3 = 0.5$ , the expected increase in delivery cost caused by decision  $a_2 = 1$  is calculated as  $9 MU - 0.5 \cdot 3 MU = 7.5 MU$ . Considering this cost additionally to the previously calculated DPC, action  $a_2 = 1$  causes an expected cost of 17.5 MU. Since the immediate contribution of accepting customer request  $c_2$  is below this expected cost, the provider has to decide for  $a_2 = 0$  in order to generate profit.

### B.3. Example 3

Example 3: Again, we consider the same problem instance, with the same potential customer requests regarding locations, revenues, and arrival rates as depicted in Table 3. Fuel cost is again 1 MU/LU. This time, we examine decision epoch t = 1. There are no confirmed customer orders yet, i.e.,  $C_0 = \{\}$ , and a request  $c_1$  realizes. The myopic MCTS equal 8 MU since the distance between  $l_{c_1}$  and the depot is 4 LU.

For decision making, also DPC need to be calculated as in the previous examples. We must calculate the sum of all expected revenues that the provider can accrue under optimal decision-making in the subsequent decision epochs until the terminal decision epoch starting in interim state  $s'_1$ . This sum equals 15 MU as we explain in the following: If the provider rejects customer request  $c_1$ , it is still possible to accept  $c_2$  and  $c_3$  if they realize. Thus, it is possible to accrue their revenues 10 MU and 20 MU with the respective arrival probabilities  $\lambda_{c_2}^2 = \lambda_{c_3}^3 = 0.5$ . From that, we subtract the corresponding sum of expected revenues under optimal decision-making starting in interim state  $s'_1(c_1)$ , which equals 10 MU as described in the following: After accepting customer request  $c_1$ , it would be optimal to reject customer request  $c_2$  if it realizes and to accept customer request  $c_3$  respectively. Overall, this yields DPC = 15 MU - 10 MU = 5 MU. Consequently, if the provider bases decision-making on myopic MCTS and DPC, the resulting optimal decision is to reject customer request  $c_1$  by action  $a_1 = 0$  because the sum of myopic MCTS and DPC exceeds its revenue. Nevertheless, the optimal decision resulting from solving the value function is accepting customer request  $c_1$  by action  $a_1 = 1$ . This is also the intuitive decision when looking at Fig. 3 and considering the vicinity to potential future customer request  $c_3$ , or in more technical terms, when considering that accepting customer request  $c_1$  entails a cost-related opportunity effect.



### B.4. Example 4

Fig. 4 Decision Tree - Example 4

Example 4: We consider the same problem instance as in Example 3, with the same potential customer requests regarding locations, revenues, and arrival rates as depicted in Table 3. Fuel

cost is again assumed to equal 1 MU/LU and decision epoch t = 1 is examined with  $C_0 = \{\}$ . Customer request  $c_1$  realizes. Decision  $a_1 = 1$ , i.e., accepting  $c_1$ , results in value  $V'_1(s'_1(c_1)) = 0.5MU$ . Rejecting it by decision  $a_1 = 0$ , results in value  $V'_1(s'_1(0)) = 5 MU$ . Consequently, the corresponding opportunity cost of decision  $a_1 = 1$  for customer request  $c_1$  when starting in the considered state equals 4.5 MU. DPC are calculated as in Example 3, thus,  $\Delta R_1(s_0, c_1) = 5 MU$ . Exploiting Property 1 yields  $MCTS = \Delta F_1(s_0, c_1) = \Delta V_1(s_0, c_1) - \Delta R_1(s_0, c_1) = 4.5 MU - 5 MU = -0.5 MU < 0 MU$ .

For illustrative purposes, Fig. 4 shows the partial decision tree for this problem instance, originating in state  $s_0$ , assuming a customer request arrives. Random nodes are depicted as circles and represent whether there is a customer request or not. The outgoing upper arc always represents the arrival of a customer request, the outgoing lower arc represents the case that there is no such arrival. Decision nodes are depicted as rectangles and represent demand control decisions. The upper arcs originating in such nodes represent accepting the respective customer request. The corresponding lower arcs represent rejecting the customer request. Optimal decisions in each decision epoch, derived from solving the corresponding value function, are depicted as solid arcs originating in the demand control decision nodes.

#### **Underlying calculations:**

$$V_{2}(s'_{2}(c_{2}) | s'_{1}(c_{1})) = 0.5 \cdot (-17) + 0.5 \cdot (-17) = -17$$

$$V_{2}(s'_{2}(0) | s'_{1}(c_{1})) = 0.5 \cdot (20 - 11) + 0.5 \cdot (-8) = 0.5$$

$$V_{2}(s'_{2}(c_{2}) | s'_{1}(0)) = 0.5 \cdot (20 - 20) + 0.5 \cdot (-9) = -4.5$$

$$V_{2}(s'_{2}(0) | s'_{1}(0)) = 0.5 \cdot (20 - 11) + 0.5 \cdot 0 = 4.5$$

$$V_{1}(s'_{1}(c_{1})) = 0.5 \cdot V_{2}(s'_{2}(0) | s'_{1}(c_{1})) + 0.5 \cdot V_{2}(s'_{2}(0) | s'_{1}(c_{1})) = 0.5$$

$$V_{1}(s'_{1}(0)) = 0.5 \cdot (10 - V_{2}(s'_{2}(c_{2}) | s'_{1}(0))) + 0.5 \cdot V_{2}(s'_{2}(0) | s'_{1}(0)) = 5$$

### **B.5. Example 5**

Example 5: We consider the same problem instance as in the previous examples. Thereby, we assume the same potential customer requests regarding locations and arrival rates, but with potential revenues  $r_{c_1} = 10 MU$ ,  $r_{c_2} = 10 MU$  and  $r_{c_3} = 10.5 MU$ . Furthermore, we now assume that the physical vehicle capacity is unrestricted, and instead, the maximum route length is constrained to 12 LU. Traveling one LU still costs 1 MU. Again, we investigate decision epoch t = 1, and again, we assume that there is no confirmed customer order yet, i.e.,  $C_0 = \{\}$ , and a customer request  $c_1$  realizes. Now,  $R'_1(s'_1(0)) = 5 MU$  as with rejecting  $c_1$  by action  $a_1 = 0$ , the subsequent optimal decisions lead to a future revenue of 10 MU with probability  $\lambda^2_{c_2} = 0.5$ . If request  $c_2$  does not realize, still, a request  $c_3$  will not be accepted. In case the current customer request  $c_1$  converts into a confirmed customer order, in turn, it is optimal to also accept a customer

request  $c_3$  if it realizes. Consequently,  $R'_1(s'_1(c_1)) = 5.25 MU$  and, thus,  $DPC = \Delta R_1(s_0, c_1) = -0.25MU < 0MU$ .



Fig. 5 Decision Tree – Example 5

Analogous to the previous example, Fig. 5 illustrates the corresponding optimal decisions for all decision epochs, derived from solving the value function.

## **Underlying calculations:**

$$V_{2}(s'_{2}(0) | s'_{1}(c_{1})) = 0.5 \cdot (10.5 - 11) + 0.5 \cdot (-8) = -4.25$$
  

$$V_{2}(s'_{2}(c_{2}) | s'_{1}(0)) = 0.5 \cdot (-9) + 0.5 \cdot (-9) = -9$$
  

$$V_{2}(s'_{2}(0) | s'_{1}(0)) = 0.5 \cdot 0 + 0.5 \cdot 0 = 0$$
  

$$V_{1}(s'_{1}(c_{1})) = 0.5 \cdot V_{2}(s'_{2}(0) | s'_{1}(c_{1})) + 0.5 \cdot V_{2}(s'_{2}(0) | s'_{1}(c_{1})) = -4.25$$
  

$$V_{1}(s'_{1}(0)) = 0.5 \cdot (10 - V_{2}(s'_{2}(c_{2}) | s'_{1}(0))) + 0.5 \cdot 0 = 0.5$$

# **Appendix C: Proofs of Lemmata**

## C.1. Proof of Lemma 1

Proof. Based on (7), we must show that the induction hypothesis

$$R'_{t}(s'_{t}) + F'_{t}(s'_{t}) = r_{\phi^{*}_{t}(g_{t})} + V_{t}(s_{t} \mid s_{t-1}, \phi^{*}_{t}(g_{t}))$$
(36)

holds for any interim state  $s'_t$  in any stage t. The proof is by induction over t.

**Initial case**: In the terminal decision epoch, rewards are defined by the boundary conditions (6), (12), and (14). For the left hand side of (36), we have  $R'_T(s'_T) + F'_T(s'_T) = 0 + r_{\phi_T^*(g_T)}$ . For the

right hand side of (36), we get the same result:  $r_{\phi_T^*(g_T)} + V_T(s_T \mid s_{T-1}, \phi_T^*(g_T)) = r_{\phi_T^*(g_T)} + 0$ . Thus, Lemma 1 holds for t = T.

**Induction step**: Now, we show that if the induction hypothesis is valid for t, then it is also valid for t - 1. For this purpose, we apply the following transformations:

- 1. We start with Equation (7) of decision epoch t 1.
- 2. We substitute Equation (8) for  $V_{t-1}(s_{t-1})$ .
- 3. We substitute Equation (7) of decision epoch t and eliminate the maximum operator by inserting the variable for the optimal demand control decision  $g_t^*$ .
- 4. We substitute the induction hypothesis (36).
- 5. We rearrange the terms.
- 6. We substitute Equations (11) and (13).

$$\begin{aligned} V_{t-1}'(s_{t-1}') &\stackrel{!}{=} r_{\phi_{t-1}^{*}(g_{t-1})} + V_{t-1}(s_{t-1}) \\ &\stackrel{?}{=} r_{\phi_{t-1}^{*}(g_{t-1})} + \sum_{c \in C} \lambda_{c}^{t} \cdot \prod_{g_{t} \in \mathcal{G}(S_{t-1,c})} (g_{t} \cdot r_{c} + V_{t}'(s_{t}' \mid s_{t-1}, c, g_{t})) \\ &+ (1 - \sum_{c \in C} \lambda_{c}^{t}) \cdot V_{t}'(s_{t}' \mid s_{t-1}, 0) \\ &\stackrel{?}{=} r_{\phi_{t-1}^{*}(g_{t-1})} + \sum_{c \in C} \lambda_{c}^{t} \cdot (g_{t}^{*} \cdot r_{c} + r_{\phi_{t}^{*}(g_{t}^{*})} + V_{t}(s_{t} \mid s_{t-1}, \phi_{t}^{*}(g_{t}^{*}))) \\ &+ (1 - \sum_{c \in C} \lambda_{c}^{t}) \cdot (r_{\phi_{t}^{*}(0)} + V_{t}(s_{t} \mid s_{t-1}, \phi_{t}^{*}(0))) \\ &\stackrel{!}{=} r_{\phi_{t-1}^{*}(g_{t-1})} + \sum_{c \in C} \lambda_{c}^{t} \cdot (g_{t}^{*} \cdot r_{c} + R_{t}'(s_{t}' \mid s_{t-1}', \phi_{t-1}^{*}(g_{t-1}), g_{t}^{*}) + \\ &F_{t}'(s_{t}' \mid s_{t-1}', \phi_{t-1}^{*}(g_{t-1}), g_{t}^{*})) \\ &+ (1 - \sum_{c \in C} \lambda_{c}^{t}) \cdot (R_{t}'(s_{t}' \mid s_{t-1}', \phi_{t-1}^{*}(g_{t-1}), 0) + F_{t}'(s_{t}' \mid s_{t-1}', \phi_{t-1}^{*}(g_{t-1}), 0)) \\ &\stackrel{!}{=} \sum_{c \in C} \lambda_{c}^{t} \cdot (g_{t}^{*} \cdot r_{c} + R_{t}'(s_{t}' \mid s_{t-1}', \phi_{t-1}^{*}(g_{t-1}), g_{t}^{*})) \\ &+ (1 - \sum_{c \in C} \lambda_{c}^{t}) \cdot R_{t}'(s_{t}' \mid s_{t-1}', \phi_{t-1}^{*}(g_{t-1}), 0) \\ &+ r_{\phi_{t-1}(g_{t-1})} + \sum_{c \in C} \lambda_{c}^{t} \cdot (F_{t}'(s_{t}' \mid s_{t-1}', \phi_{t-1}^{*}(g_{t-1}), g_{t}^{*})) \\ &+ (1 - \sum_{c \in C} \lambda_{c}^{t}) \cdot F_{t}'(s_{t}' \mid s_{t-1}', \phi_{t-1}^{*}(g_{t-1}), 0) \\ &\stackrel{!}{=} R_{t-1}(s_{t-1}') + F_{t-1}'(s_{t-1}'). \quad \blacksquare \qquad (37)
\end{aligned}$$

#### C.2. Proof of Lemma 2

*Proof.* Lemma 2 holds by definition for the generic i-DMVRP model since the customer arrival process is only time-dependent, but not state-dependent, by the definition of the arrival rates  $\lambda_c^t$ .

#### C.3. Proof of Lemma 3

Generally, the demand control component  $\mathcal{G}(s_{t-1}, c)$  of an action space  $\mathcal{A}_t(s_{t-1}, c) = (\mathcal{G}(s_{t-1}, c), \Phi(s_{t-1}, c, g_t))$  strongly depends on the tour planning component  $\Phi(s_{t-1}, c, g_t)$  as follows. It only comprises  $g_t = 1$  if a feasible tour plan  $\phi_t(1)$  exists in  $\Phi(s_{t-1}, c, g_t)$ . Thus, to

prove Lemma 3, it is sufficient to show that the tour planning component of  $\mathcal{A}(\hat{s}_{t-1}, c)$  is a subset of the respective tour planning component in  $\mathcal{A}(s_{t-1}, c)$ .

*Proof.*  $\Phi(s_{t-1}, c, g_t)$  corresponds to the solution space of the constraint satisfaction variant of the underlying VRP. Enforcing the fulfillment of an additional order  $\hat{c}$  requires (at least) one additional constraint compared to fulfilling the set of orders  $C_{t-1}$ , without order  $\hat{c}$ . This constraint is either redundant or further restricts the solution space, as long as the triangle inequality holds, which proves Lemma 3 (see Asdemir et al. (2009) for a similar proof).

#### C.4. Proof of Lemma 4

To prove Lemma 4, it is sufficient to show that, when rejecting a certain customer request  $\hat{c}$ , we can feasibly make the same future decisions as when accepting  $\hat{c}$ , and that both decisions result in the same transitions.

*Proof.* By Lemma 3, the same future decisions can be made by assuming the rejection of a customer request  $\hat{c}$  as well as by assuming its acceptance. Further, by Lemma 2, those decisions result in the same subsequent transitions.

### C.5. Proof of Lemma 5

*Proof.* From Lemmata 3 and 4, it follows that any  $\pi$  that can be feasibly applied to  $\omega$  starting in interim state  $s'_t(c_t)$  can also be feasibly applied to  $\omega$  starting in interim state  $s'_t(0)$ . Then, Lemma 5 directly follows from Lemma 2 since, starting in both interim states, the same set of customer orders are received given the same decisions.

#### C.6. Proof of Lemma 6

*Proof.*  $\pi$  and  $\omega$  start with the same set of confirmed customer orders  $C_{t-1}$ , irrespective of whether their start is assumed from interim state  $s'_t(c_t)$  or  $s'_t(0)$ . Then, assuming  $\pi$  and  $\omega$  start in interim state  $s'_t(c_t)$ , customer order  $c_t$  is added to the set of confirmed customer orders  $C_t$  whereas it is not added assuming  $\pi$  and  $\omega$  start in interim state  $s'_t(0)$ . Afterward, starting in interim states  $s'_t(c_t)$  and  $s'_t(0)$ , respectively, and applying the same decision sequence  $\pi$  to  $\omega$ , which is possible by Lemmata 3 and 4, again, the same customer orders are confirmed, which results from Lemma 2. Consequently, starting  $\pi$  and  $\omega$  in interim state  $s'_t(c_t)$  results in subsequent states  $\hat{s}_{t'} =$  $(\hat{C}_{t'}, \phi_{t'})$  and starting in interim state  $s'_t(0)$  results in subsequent states  $s_{t'} = (C_{t'}, \phi_{t'})$  for t' = $t, \ldots, T$ , with  $\hat{C}_{t'} = C_{t'} \cup \{c_t\}$ . This proves Lemma 6 analogously to the proof of Lemma 3, as the underlying VRP is more restricted and consequently the resulting fulfillment cost cannot be smaller as long as the triangle inequality holds.

#### C.7. Proof of Lemma 7

*Proof.* The proof of Lemma 7 follows directly from Lemma 5 and Lemma 6: Lemma 5 states revenue equality when applying a decision sequence  $\pi$  to sample path  $\omega$  starting in interim state  $s'_t(c_t)$  and starting in interim state  $s'_t(0)$ , i.e.,  $R'^{\pi\omega}(s'_t(c_t)) = R'^{\pi\omega}(s'_t(0))$ . Lemma 6 states that higher or equal fulfillment cost arise for decision sequence  $\pi$  applied to sample path  $\omega$  starting in

interim state  $s'_t(c_t)$  over applying the same decision sequence  $\pi$  to sample path  $\omega$  starting in interim state  $s'_t(0)$ , i.e.,  $F'^{\pi\omega}(s'_t(c_t)) \leq F'^{\pi\omega}(s'_t(0))$ . Hence, substituting Lemma 5 and Lemma 6 yields:

$$V_{t}^{\prime\pi\omega}(s_{t}^{\prime}(c_{t})) = R_{t}^{\prime\pi\omega}(s_{t}^{\prime}(c_{t})) + F_{t}^{\prime\pi\omega}(s_{t}^{\prime}(c_{t})) \leq R_{t}^{\prime\pi\omega}(s_{t}^{\prime}(0)) + F_{t}^{\prime\pi\omega}(s_{t}^{\prime}(0)) = V_{t}^{\prime\pi\omega}(s_{t}^{\prime}(0)).$$
(38)

# Appendix D: Proof of Model Equivalency

Denoting the value function of the original model by  $V_t(s_t)$  and the one of the modified model by  $\tilde{V}_t(s_t)$ , the following relationship holds:  $V_t(s_t) = \tilde{V}_t(s_t) - r_t^{l\ cum}$ . Before proving the model equivalency, we reformulate the corresponding value functions  $V_t(s_t)$  and  $\tilde{V}_t(s_t)$  of the original and the modified model for ease of presentation. Then, we prove model equivalency by induction. First, we replace the maximization operators by the corresponding optimal decisions. We represent the optimal demand control decision, i.e.,  $\max_{g_t \in \mathcal{G}(s_{t-1},c)}$ , by  $g_t^*$ . The optimal tour planning decision, i.e.,  $\max_{\phi_t(g_t) \in \Phi(s_{t-1},c,g_t)}$ , is represented by  $\phi_t^*(g_t)$ . Further, we introduce expressions  $r_t^+$  and  $r_t^-$  to replace expectations over positive rewards (revenues) and negative rewards (costs), respectively.

Value function of the original model:

$$V_{t-1}(s_{t-1}) = \sum_{c \in C} \lambda_c^t \cdot \left( g_t^* \cdot r_c + r_{\phi_t^*(g_t^*)} + V_t \left( s_t \mid s_{t-1}, \phi_t^*(g_t^*) \right) \right) + (1 - \sum_{c \in C} \lambda_c^t) \cdot \left( r_{\phi_t^*(0)} + V_t \left( s_t \mid s_{t-1}, \phi_t^*(0) \right) \right) = \underbrace{\sum_{c \in C} \lambda_c^t \cdot g_t^* \cdot r_c}_{r_t^+} + \underbrace{\sum_{c \in C} \lambda_c^t \cdot r_{\phi_t^*(g_{t-1}^*)} + (1 - \sum_{c \in C} \lambda_c^t) \cdot r_{\phi_t^*(0)}}_{r_t^-} + \underbrace{\sum_{c \in C} \lambda_c^t \cdot V_t \left( s_t \mid s_{t-1}, \phi_t^*(g_t^*) \right) + (1 - \sum_{c \in C} \lambda_c^t) \cdot V_t \left( s_t \mid s_{t-1}, \phi_t^*(0) \right)}_{V_t(s_t)} = r_t^+ + r_t^- + V_t(s_t).$$
(39)

Value function of the modified model:

$$\begin{split} \tilde{V}_{t-1}(s_{t-1}) &= \sum_{c \in C} \lambda_c^t \cdot \left( g_t^* \cdot r_c + \tilde{V}_t \left( s_t \mid s_{t-1}, \phi_t^*(g_{t-1}^*) \right) \right) + (1 - \sum_{c \in C} \lambda_c^t) \cdot \\ \tilde{V}_t \left( s_t \mid s_{t-1}, \phi_t^*(0) \right) \\ &= \underbrace{\sum_{c \in C} \lambda_c^t \cdot (g_t^* \cdot r_c)}_{r_t^+} + \underbrace{\sum_{c \in C} \lambda_c^t \cdot \tilde{V}_t \left( s_t \mid s_{t-1}, \phi_t^*(g_t^*) \right) + (1 - \sum_{c \in C} \lambda_c^t) \cdot \tilde{V}_t \left( s_t \mid s_{t-1}, \phi_t^*(0) \right)}_{\tilde{V}_t(s_t)} \\ &= r_t^+ + \tilde{V}_t(s_t). \end{split}$$
(40)

Model equivalency is given if  $V_t(s_t) = \tilde{V}_t(s_t) - r_t^{l \ cum}$  holds for every t, and  $V_0(s_0) = \tilde{V}_0(s_0)$  holds as well. We prove this by induction, for which we remind the reader of the following two relationships:

$$V_T(s_T) = 0 \tag{41}$$
(42)

 $r_{t-1}^{l\ cum} = r_t^{l\ cum} - r_{\phi_t(g_t)}$ 

The proof starts with the result for the terminal state  $s_T$  to which the system transitioned via a certain sample path denoted by  $\omega$ . Then, independent of which sample path realizes, i.e., for every sample path in the set of all potential sample paths,  $\omega \in \Omega$ , we can conduct the following proof:

Proof. By induction:

#### Initial case:

$$V_T^{\omega}(s_T) = \tilde{V}_T^{\omega}(s_T) - r_T^{\omega \ l \ cum} = 0$$
(43)

Equation (43) holds by definition.

#### Induction hypothesis:

$$V_t^{\omega}(s_t) = \tilde{V}_t^{\omega}(s_t) - r_t^{\omega \ l \ cum}$$

$$\tag{44}$$

**Induction step**: Now, we show that if the induction hypothesis is valid for t, then it is also valid for t - 1. For this purpose, we apply the following transformations:

- 1. From Equation (39), the following relationship can be derived to start with:  $V_{t-1}^{\omega}(s_{t-1}) = V_t^{\omega}(s_t) + r_t^{\omega} + r_t^{\omega}$ .
- 2. We substitute the induction hypothesis (44).
- 3. We substitute Equation (42).
- 4. We substitute Equation (40).
- 5. Since  $r_{\phi_t(g_t)}^{\omega} = r_t^{\omega}$  for  $c_t$  that is observed in sample path  $\omega$ , we can rearrange the terms and eliminate  $r_t^{\omega}$  + and  $r_t^{\omega}$  -.

$$V_{t-1}^{\omega}(s_{t-1}) \stackrel{1}{=} V_t^{\omega}(s_t) + r_t^{\omega} + r_t^{\omega} -$$
(45)

$$\stackrel{2.}{=} \tilde{V}_{t}^{\omega}(s_{t}) - r_{t}^{\omega} \, {}^{l} \, {}^{cum} + r_{t}^{\omega} \, {}^{+} + r_{t}^{\omega} \, {}^{-} \tag{46}$$

$$\stackrel{3.}{=} \tilde{V}_{t}^{\omega}(s_{t}) - \left(r_{t-1}^{\omega \ l \ cum} + r_{\phi_{t}(g_{t})}^{\omega}\right) + r_{t}^{\omega \ +} + r_{t}^{\omega \ -}$$
(47)

$$\stackrel{4.}{=} \left( \tilde{V}_{t-1}^{\omega}(s_{t-1}) - r_t^{\omega} \right) - \left( r_{t-1}^{\omega \ l \ cum} + r_{\phi_t(g_t)}^{\omega} \right) + r_t^{\omega} + r_t^{\omega} - \tag{48}$$

$$\stackrel{5.}{=} \tilde{V}^{\omega}_{t-1}(s_{t-1}) - r^{\omega}_{t-1} \,^{l \ cum}. \tag{49}$$

We have shown that, if the induction hypothesis holds for t, it also holds for t - 1, and thus, also for t = 0, i.e.,  $V_0^{\omega}(s_0) = \tilde{V}_0^{\omega}(s_0) - r_0^{\omega l cum}$ . Further, it holds by definition that  $r_0^{\omega l cum} = r_0^{\omega -} = 0$ . Consequently,  $V_0^{\omega}(s_0) = \tilde{V}_0^{\omega}(s_0)$  holds for any sample paths  $\omega$ , and we conclude that the original and the modified model are equivalent.

#### Appendix E: Model Generalization to Multiple Fulfillment Options

As explained in Section 4.5, the analytical considerations for the single-option model can be readily generalized to models with multiple fulfillment options. In the following, to provide further details on this generalization, we demonstrate it on a formal level. To this end, we first show how the generic MDP model can be adapted to the multi-option case. Second, we present the generalized Bellman equation and derive the generalized definitions of OC, DPC, and MCTS. Action – We define a set  $O(s_{t-1}, c)$  of feasible fulfillment options  $o \in O(s_{t-1}, c)$  that can potentially be offered to a requesting customer of type c starting in state  $s_{t-1}$  and includes the nopurchase option o = 0. Hence, in the multi-option model, a demand control decision is equivalent to selecting a subset of feasible fulfillment options  $g_t \subseteq O(s_{t-1}, c)$ . The corresponding action space  $G(s_{t-1}, c) = 2^{O(s_{t-1}, c)}$  at a decision epoch t starting in state  $s_{t-1}$  with an arriving customer request of type c equals the power set of all feasible fulfillment options. The tour planning decision  $(\phi_t(o))_{o \in O(s_{t-1}, c)}$  comprises a potential tour plan for each feasible fulfillment option with the respective action spaces given by  $\Phi(s_{t-1}, c, o)$  for all  $o \in O(s_{t-1}, c)$ .

*Transition* – Following the selection of an offer set  $g_t$  by the provider, there is an additional stochastic transition because the customer chooses a fulfillment option according to choice probabilities  $P_o(g_t)$ . If the customer chooses a fulfillment option  $o'_t \neq 0$ , an order is confirmed and the tour plan is set to  $\phi_t(o'_t)$ . Overall, the transitions can be summarized as follows:

$$\phi_t = \phi_t(o_t') \tag{50}$$

$$\mathcal{C}_{t} = \begin{cases} \mathcal{C}_{t-1} \setminus \Psi(\phi_{t}), & \text{if there is no customer request arrival or if } o_{t}' = 0\\ (\mathcal{C}_{t-1} \cup \{c_{t}\}) \setminus \Psi(\phi_{t}), & \text{if } o_{t}' \neq 0 \end{cases}$$
(51)

*Rewards* – If a customer of type *c* chooses a fulfillment option  $o \in g_t$ , the provider receives the demand-control-related reward  $r_{co}$ . The logistics-related reward is denoted as  $r_{\phi_t(o)}$ .

Based on the above-described extensions to the MDP model, we now generalize the Bellman equation (5) to the multi-option case:

$$V_{t-1}(s_{t-1}) = \sum_{c \in C} \lambda_c^t \cdot \max_{g_t \in \mathcal{G}(s_{t-1},c)} \left( \sum_{o \in g_t} P_o(g_t) \cdot \left[ r_{co} + \max_{\phi_t(o) \in \Phi(s_{t-1},c,o)} \left( r_{\phi_t(o)} + V_t(s_t \mid s_{t-1},\phi_t(o)) \right) \right] \right) + (1 - \sum_{c \in C} \lambda_c^t) \cdot \max_{\phi_t(0) \in \Phi(s_{t-1},0)} \left( r_{\phi_t(0)} + V_t(s_t \mid s_{t-1},\phi_t(0)) \right) \right)$$
(52)

with boundary condition:

$$V_T(s_T) = 0. (53)$$

Drawing on the interim state  $s'_t | s_{t-1}, c, o$ , we obtain the multi-option variant of Bellman equation (8):

$$V_{t-1}(s_{t-1}) = \sum_{c \in C} \lambda_c^t \cdot \max_{g_t \in \mathcal{G}(s_{t-1},c)} \left( \sum_{o \in g_t} P_o(g_t) \cdot [r_{co} + V_t'(s_t' \mid s_{t-1}, c, o)] \right) + (1 - \sum_{c \in C} \lambda_c^t) \cdot V_t'(s_t' \mid s_{t-1}, 0).$$
(54)

Next, we define opportunity cost, which is now specific to fulfillment options  $o \in O(s_{t-1}, c)$ :

$$\Delta V_t(s_{t-1}, c, o) = V'_t(s'_t(0)) - V'_t(s'_t(o)),$$
(55)

with interim state  $s'_t | s_{t-1}, c, o$  denoted as  $s'_t(o)$  and interim state  $s'_t | s_{t-1}, c, 0$  denoted as  $s'_t(0)$ . Making use of the multi-option definition of opportunity cost, we formulate the multi-option equivalent of Bellman equation (10):

$$V_{t-1}(s_{t-1}) = \sum_{c \in C} \lambda_c^t \cdot \max_{g_t \in \mathcal{G}(s_{t-1},c)} \left( \sum_{o \in g_t} P_o(g_t) \cdot [r_{co} - \Delta V_t(s_{t-1},c,o)] \right) + V_t'(s_t'(0)).$$
(56)

Analogously to the single-option case (Equation (11) and Equation (13)), we can now define the expected future revenue  $R'_{t-1}(s'_{t-1})$  and the expected future fulfillment cost  $F'_{t-1}(s'_{t-1})$  of a given interim state  $s'_{t-1}$  at decision epoch t - 1:

$$R'_{t-1}(s'_{t-1}) = \sum_{c \in C} \lambda^t_c \cdot \left( \sum_{o \in g^*_t} P_o(g^*_t) \cdot [r_{co} + R'_t(s'_t \mid s'_{t-1}, \phi^*_{t-1}(o'_{t-1}), o] \right) + (1 - \sum_{c \in C} \lambda^t_c) \cdot R'_t(s'_t \mid s'_{t-1}, \phi^*_{t-1}(o'_{t-1}), 0),$$
(57)

with boundary condition

$$R'_T(s'_T) = 0, (58)$$

and  $o'_{t-1}$  denoting the fulfillment option chosen by the customer arriving at decision epoch t-1.

$$F'_{t-1}(s'_{t-1}) = r_{\phi^*_{t-1}(o'_{t-1})} + \sum_{c \in C} \lambda^t_c \cdot \sum_{o \in g^*_t} P_o(g^*_t) \cdot \left(F'_t(s'_t \mid s'_{t-1}, \phi^*_{t-1}(o'_{t-1}), o)\right) + (1 - \sum_{c \in C} \lambda^t_c) \cdot F'_t(s'_t \mid s'_{t-1}, \phi^*_{t-1}(o'_{t-1}), 0),$$
(59)

with boundary condition:

$$F'_T(s'_T) = r_{\phi^*_T(o'_T)}.$$
(60)

Finally, we formulate the formal definition of DPC and MCTS for the multi-option case:

$$\Delta R_t(s_{t-1}, c, o) = R'_t(s'_t(0)) - R'_t(s'_t(o)).$$
(61)

$$\Delta F_t(s_{t-1}, c, o) = F'_t(s'_t(0)) - F'_t(s'_t(o)).$$
(62)

# Appendix F: Further Numerical Results: p-Values

 Table 4 p-Values of one-tailed paired t-tests on average objective values resulting from applying the different approximations for physical capacity-constrained settings I/II

		$H_0$ :					
Settings	Profit.	DPC ≤	$MCTS \leq$	$DPC \leq$	HR≤	$MCTS \leq$	HR≤
		MCTS	DPC	HR	DPC	HR	MCTS
(unif, homog)	low	1	4.900e-21	1	4.638e-18	1.015e-06	0.999
(clust, homog)	low	1	2.425e-22	1	6.147e-24	0.013	0.986
(unif, rand)	low	1	5.916e-18	0.999	1.544e-12	0.001	0.998
(clust, rand)	low	1	1.065e-22	0.999	5.734e-13	2.790e-12	0.999
(clust_sort, rand)	low	0.999	0.0002	0.991	0.008	0.008	0.991
(unif, l-b-h)	low	0.999	1.983e-16	0.999	1.979e-12	0.003	0.996
(clust, l-b-h)	low	1	1.259e-23	0.999	5.744e-15	7.575e-11	0.999
(clust_sort, l-b-h)	low	0.999	1.131e-14	1	1.031e-17	0.007	0.992
(unif, h-b-l)	low	0.999	2.156e-15	0.999	5.399e-09	1.466e-05	0.999
(clust, h-b-l)	low	1	1.275e-19	0.999	1.990e-06	6.276e-12	0.999
(clust_sort, h-b-l)	low	7.722e-13	0.999	2.352e-13	0.999	0.164	0.835
(unif, homog)	high	1	1	1	4.103e-20	1	4.103e-20
(clust, homog)	high	1	1	1	8.007e-23	1	8.007e-23
(unif, rand)	high	4.389e-16	0.999	0.999	3.061e-11	1	5.227e-20
(clust, rand)	high	7.178e-17	0.999	0.999	6.156e-10	1	3.768e-21
(clust_sort, rand)	high	5.553e-21	1	0.999	0.0003	1	1.596e-26
(unif, l-b-h)	high	5.246e-32	1	0.999	3.058e-15	1	2.971e-32
(clust, l-b-h)	high	4.779e-43	1	0.999	1.035e-12	1	7.622e-38
(clust_sort, l-b-h)	high	4.655e-34	1	0.993	0.006	1	8.183e-42
(unif, h-b-l)	high	1	1	0.999	1.155e-11	0.999	1.155e-11
(clust, h-b-l)	high	1	1	1	2.154e-17	1	2.154e-17
(clust sort, h-b-l)	high	1	1	0.999	9.676e-08	0.999	9.676e-08

		<i>H</i> <sub>0</sub> :					
Settings	Profit.	DPC ≤	OO_MCTS	$MCTS \leq$	OO_MCTS	$OO_DPC \le$	OO_MCTS
		OO_MCTS	$\leq$ DPC	OO_MCTS	$\leq$ MCTS	OO_MCTS	$\leq$ OO_DPC
(unif, homog)	low	1	4.256e-20	0.004	0.995	1	6.710e-19
(clust, homog)	low	1	1.252e-23	0.982	0.017	0.999	4.664e-09
(unif, rand)	low	1	3.181e-20	0.798	0.201	1	2.579e-17
(clust, rand)	low	1	6.821e-24	0.001	0.998	1	3.191e-20
(clust_sort, rand)	low	0.999	3.430e-13	0.999	6.480e-09	0.999	7.588e-17
(unif, 1-b-h)	low	1	1.109e-19	0.825	0.174	0.999	2.659e-15
(clust, l-b-h)	low	1	1.734e-21	9.905e-05	0.999	1	2.847e-17
(clust_sort, l-b-h)	low	1	2.887e-38	0.999	1.198e-08	1	1.223e-27
(unif, h-b-l)	low	1	2.124e-17	0.721	0.278	1	1.660e-22
(clust, h-b-l)	low	1	6.043e-26	0.984	0.015	1	4.951e-20
(clust_sort, h-b-l)	low	1	5.997e-20	1	2.012e-23	1	1.917e-30
(unif, homog)	high	1	7.345e-21	1	7.345e-21	1	4.879e-21
(clust, homog)	high	1	2.108e-21	1	2.108e-21	1	2.033e-21
(unif, rand)	high	0.999	3.896e-07	1	9.953e-18	0.999	1.356e-05
(clust, rand)	high	0.999	8.823e-06	1	2.490e-17	0.999	1.911e-07
(clust_sort, rand)	high	0.999	0.0002	1	1.305e-22	0.999	7.760e-05
(unif, 1-b-h)	high	0.999	7.929e-17	1	9.466e-31	0.999	5.259e-12
(clust, l-b-h)	high	0.999	5.546e-17	1	1.989e-40	0.999	2.440e-16
(clust_sort, l-b-h)	high	0.825	0.174	1	2.028e-33	0.864	0.135
(unif, h-b-l)	high	0.999	1.385e-11	0.999	1.385e-11	0.999	2.582e-18
(clust, h-b-l)	high	1	2.027e-18	1	2.027e-18	1	2.582e-18
(clust_sort, h-b-l)	high	0.639	0.3602	0.639	0.360	0.219	0.780

**Table 5** p-Values of one-tailed paired t-tests on average objective values resulting from applying the different approximations for physical capacity-constrained settings II/II

**Table 6** p-Values of one-tailed paired t-tests on average objective values resulting from applying the different approximations for route length-constrained settings I/II

		$H_0$ :					
Settings	Profit.	DPC≤	$MCTS \leq$	$DPC \leq$	$\mathrm{HR} \leq$	$MCTS \leq$	$\mathrm{HR} \leq$
		MCTS	DPC	HR	DPC	HR	MCTS
(unif, homog)	low	0.239	0.760	0.106	0.893	0.001	0.998
(clust, homog)	low	2.122e-08	0.999	4.345e-11	0.999	0.012	0.987
(unif, rand)	low	0.006	0.993	0.001	0.998	0.005	0.994
(clust, rand)	low	0.939	0.060	0.028	0.971	0.0001	0.999
(clust_sort, rand)	low	3.886e-18	1	6.508e-24	1	0.029	0.970
(unif, l-b-h)	low	1.324e-05	0.999	1.645e-05	0.999	0.680	0.319
(clust, l-b-h)	low	0.835	0.164	0.230	0.769	0.014	0.985
(clust_sort, l-b-h)	low	5.786e-18	1	1.127e-36	1	0.007	0.992
(unif, h-b-l)	low	0.0003	0.999	2.121e-05	0.999	0.001	0.998
(clust, h-b-l)	low	0.015	0.984	8.219e-11	0.999	1.690e-09	0.999
(clust_sort, h-b-l)	low	5.137e-24	1	5.140e-21	1	0.985	0.014
(unif, homog)	high	0.001	0.998	0.999	0.0001	0.999	0.0003
(clust, homog)	high	0.007	0.992	0.999	5.989e-06	0.999	2.772e-05
(unif, rand)	high	0.0003	0.999	0.995	0.004	0.999	0.0002
(clust, rand)	high	9.636e-07	0.999	0.999	0.0001	0.999	6.101e-08
(clust_sort, rand)	high	5.166e-09	0.999	0.998	0.001	0.999	2.256e-09
(unif, l-b-h)	high	2.742e-06	0.999	0.999	0.0001	0.999	2.050e-06
(clust, l-b-h)	high	4.377e-05	0.999	0.999	0.0001	0.999	4.564e-06
(clust_sort, l-b-h)	high	1.658e-56	1	1	1	1	1.658e-56
(unif, h-b-l)	high	0.005	0.994	0.999	0.0008	0.996	0.003
(clust, h-b-l)	high	0.002	0.997	0.999	2.464e-05	0.999	0.0001
(clust_sort, h-b-l)	high	1.091e-48	1	0.919	0.080	1	2.825e-49

Table 7 p-Values of one-tailed paired t-tests on average	e objective values	s resulting from a	pplying the dif	ferent approx-
imations for route length-constrained settings II/II				

		<i>H</i> <sub>0</sub> :					
Settings	Profit.	DPC ≤	OO_MCTS	$MCTS \leq$	OO_MCTS	$OO_DPC \le$	OO_MCTS
		OO_MCTS	$\leq$ DPC	OO_MCTS	$\leq$ MCTS	OO_MCTS	$\leq$ OO_DPC
(unif, homog)	low	0.114	0.885	0.259	0.740	0.999	8.009e-07
(clust, homog)	low	0.997	0.002	0.999	1.018e-10	1	1.287e-18
(unif, rand)	low	0.001	0.998	0.552	0.447	0.998	0.001
(clust, rand)	low	0.880	0.119	0.272	0.727	0.999	0.0001
(clust_sort, rand)	low	0.024	0.975	0.999	5.385e-14	0.999	2.894e-10
(unif, 1-b-h)	low	0.0001	0.999	0.976	0.023	0.999	3.055e-06
(clust, l-b-h)	low	0.969	0.030	0.571	0.428	0.999	2.334e-05
(clust_sort, l-b-h)	low	0.999	2.645e-05	1	9.137e-19	1	1.491e-19
(unif, h-b-l)	low	0.022	0.977	0.998	0.001	0.976	0.023
(clust, h-b-l)	low	0.933	0.066	0.999	0.0004	0.999	2.029e-08
(clust_sort, h-b-l)	low	0.998	0.001	1	1.394e-25	1	2.352e-22
(unif, homog)	high	0.001	0.998	1	1	0.0005	0.999
(clust, homog)	high	0.005	0.994	0.161	0.838	0.0007	0.999
(unif, rand)	high	0.0003	0.999	1	1	0.0002	0.999
(clust, rand)	high	1.702e-05	0.999	0.356	0.643	1.449e-05	0.999
(clust_sort, rand)	high	7.159e-14	0.999	0.0001	0.999	6.552e-14	0.999
(unif, l-b-h)	high	2.064e-06	0.999	0.100	0.899	1.889e-06	0.999
(clust, l-b-h)	high	2.104e-05	0.999	0.7381	0.261	2.003e-05	0.999
(clust_sort, l-b-h)	high	0.0993	0.900	1	9.713e-56	0.099	0.900
(unif, h-b-l)	high	0.005	0.994	1	1	0.003	0.996
(clust, h-b-l)	high	0.002	0.997	1	1	0.001	0.998
(clust_sort, h-b-l)	high	0.999	0.0007	1	4.116e-39	0.998	0.001

# Article A3: From Approximation Error to Optimality Gap – Explaining the Performance Impact of Opportunity Cost Approximation in Integrated Demand Management and Vehicle Routing

David Fleckenstein<sup>a</sup>, Robert Klein<sup>a</sup>, Vienna Klein<sup>a</sup>, Claudius Steinhardt<sup>b</sup> (<sup>a</sup>University of Augsburg, <sup>b</sup>University of the Bundeswehr Munich) Submitted to Transportation Science (category A), 1<sup>st</sup> submission round

#### Abstract

The widespread adoption of digital distribution channels both enables and forces more and more logistical service providers to manage booking processes actively to maintain competitiveness. As a result, their operational planning is no longer limited to solving vehicle routing problems. Instead, demand management decisions and vehicle routing decisions are optimized integratively with the aim of maximizing revenue and minimizing fulfillment cost. The resulting integrated demand management and vehicle routing problems (i-DMVRPs) can be formulated as Markov decision process models and, theoretically, can be solved via the well-known Bellman equation. Unfortunately, the Bellman equation is intractable for realistic-sized instances. Thus, in the literature, i-DMVRPs are often addressed via decomposition-based solution approaches involving an opportunity cost approximation as a key component. Despite its importance, to the best of our knowledge, there is neither a technique to systematically analyze how the accuracy of the opportunity cost approximation translates into overall solution quality nor are there general guidelines on when to apply which class of approximation approach.

In this work, we address this research gap by proposing an explainability technique that quantifies and visualizes the magnitude of approximation errors, their immediate impact, and their relevance in specific regions of the state space. Exploiting reward decomposition, it further yields a characterization of different types of approximation errors. Applying the technique to a generic i-DMVRP in a full-factorial computational study and comparing the results with observations in existing literature, we show that the technique contributes to better explaining algorithmic performance and provides guidance for the algorithm selection and development process.

Key words: Explainability, Demand Management, Opportunity Cost

# 1 Introduction

The proliferation of e-commerce and the progress of communication technology has led to the emergence and establishment of new business models that allow customers to book on-demand logistical services, mostly the delivery of goods (Waßmuth et al., 2023) or local transportation (Vansteenwegen et al., 2022). Prominent examples of these services are attended home delivery (AHD), same-day delivery (SDD), or mobility-on-demand (MOD). These business models have in common that customers expect a very high service level, e.g., in terms of the deviation from their desired service time (Amorim et al., 2024). Meeting these expectations makes demand consolidation challenging, which entails high fulfillment cost (Ulmer, 2020). To still operate profitably, operational planning for these business models has evolved: Instead of optimizing the associated vehicle routing alone, providers additionally apply demand management to achieve efficient fulfillment operations.

The resulting integrated demand management and vehicle routing problems (i-DMVRPs) are stochastic and dynamic with two types of integrated decisions: For each dynamically arriving customer request, the provider integratively makes a *demand control decision* and a *vehicle routing decision* with the overall objective of maximizing the expected profit, i.e., revenue net of operational fulfillment cost. Such an i-DMVRP can be modeled as a Markov decision process (MDP) and, theoretically, be solved by evaluating the well-known Bellman equation (Puterman, 2014). Practically, however, i-DMVRPs suffer from the curses of dimensionality (Powell, 2011) such that this is not tractable for realistic-sized instances. Consequently, in literature, demand control decisions for i-DMVRPs are often optimized with a decomposition-based solution approach. More precisely, two subproblems are solved sequentially for every incoming customer request (Fleckenstein et al., 2023, Ulmer, 2020, Gallego and Topaloglu, 2019, p. 25, Klein et al., 2018):

- 1.) Approximating opportunity cost (OC) for each potential fulfillment option (e.g., different time windows) to measure the expected profit impact assuming the current customer chooses the respective option, given the state of the system.
- 2.) Solving the actual demand control problem based on the approximated OC, i.e., deciding on the pricing or availability of fulfillment options, or the acceptance/rejection of the request.

This leads to the following conclusion: One of the main avenues for accelerating the development of practical solution approaches for i-DMVRPs is understanding and explaining the relation between the accuracy reached in (1), i.e., the *accuracy of the OC approximation*, and the quality of (2), i.e., the *quality of the resulting demand control decision*. However, despite the maturity of integrated demand management and vehicle routing as a research area, this relation has not been systematically explored, and therefore, largely remains a black box so far. Hence, we now close this research gap by comprehensively analyzing this relationship. To do so, we first introduce a novel explainability technique for i-DMVRPs that combines two building blocks:

- *B1: Chain of influencing factors* The first building block resembles a typical post-hoc explainability technique (Arrieta et al., 2020), i.e., we define metrics to evaluate the behavior of a given policy in certain states and the respective impact on solution quality. Thereby, we aim at answering the central questions that arise along the chain of influencing factors from OC approximation error to objective value loss as depicted in Fig. 1. This allows deriving insights on when (in which states) and why a certain policy performs especially good or bad.
- *B2: Reward decomposition* The second building block incorporates the idea of reward decomposition (Juozapaitis et al., 2019). For that, we exploit the finding that opportunity cost can be decomposed into displacement cost (DPC) and marginal cost-to-serve (MCTS) (Fleck-enstein et al., 2024). More precisely, we propose to apply OC approximations capturing only one of the two components, with the aim of assessing the importance of the respective component for approximation accuracy.

Afterward, in an extensive computational study, we apply the explainability technique to the generic i-DMVRP and the stylized parameter settings from Fleckenstein et al. (2024). Then, we complement the numerical results by an analysis of the existing literature on i-DMVRPs. Therewith, we confirm the validity of our findings and show that our identified OC approximation error types can indeed explain the observed performance of state-of-the-art solution approaches. In summary, our work has four contributions:

- 1.) To the best of our knowledge, we introduce the first explainability technique for the widely established decomposition-based solution approaches for i-DMVRPs.
- 2.) We apply our explainability technique within a comprehensive computational study and identify fundamental OC approximation error types, i.e., OC approximation errors that can occur in a broad variety of real-world i-DMVRPs. Therewith, we are the first to systematically analyze the relation between the accuracy of OC approximation and the objective value.
- 3.) We classify patterns in the occurrence of the fundamental approximation error types we identify, characterize which problem settings are prone to which error type, and propose algorithmic elements to successfully mitigate them. This yields insights that guide the selection and the design of OC approximation algorithms.
- 4.) We compile indications for the occurrence of the identified error types from existing literature and show that our findings improve explainability of the reported results. Thereby, we transform the existing implicit knowledge about specific i-DMVRPs to explicit, high-level knowledge.

The remainder of this paper is structured as follows: In Section 2, we review the related literature both on algorithmic explainability and i-DMVRPs. In Section 3, we introduce and model the generic i-DMVRP under consideration. Then, in Section 4, we present our novel explainability technique for i-DMVRPs in detail, and we present our computational study in Section 5. In Section 6, we derive general insights for algorithm design and summarize our work in Section 7.

Article A3: From Approximation Error to Optimality Gap – Explaining the Performance Impact of Opportunity Cost Approximation in Integrated Demand Management and Vehicle Routing



Fig. 1 Chain of influencing factors from OC approximation error to objective value loss

### 2 Literature Review

Due to the cross-cutting nature of our study, the related literature spans across multiple distinct research areas. In Section 2.1, we review the literature on i-DMVRPs with a special focus on its origins in revenue management, dynamic pricing, and dynamic vehicle routing. In Section 2.2, we then discuss algorithmic explainability techniques, particularly from explainable reinforcement learning (RL), highlighting the techniques that we adapt and apply in the work at hand. Finally, in Section 2.3, we review the descriptive analytics that authors use to explain the observed performance of their i-DMVRP solution approaches.

#### 2.1 Modeling and Solving i-DMVRPs

In logistics, many companies dynamically collect orders for a transportation service that is fulfilled by a given fleet of vehicles. These companies face an i-DMVRP if they can both plan individual offers made in response to customer requests and plan the vehicle routes to feasibly fulfill the resulting orders. Hence, i-DMVRP research synthesizes two originally distinct research areas:

- Research in *revenue management and dynamic pricing* addresses the dynamic optimization of offering decisions under the assumption that fulfillment is already pre-planned. For an extensive overview of this field, we refer the reader to the textbooks by Gallego and Topaloglu (2019) and Talluri and Van Ryzin (2004) as well as the reviews by Klein et al. (2020) and Strauss et al. (2018).
- 2.) Dynamic vehicle routing investigates the optimization of fulfillment assuming given orders that arrive dynamically. For a deeper discussion of this research area, we refer the reader to the textbook by Toth and Vigo (2014) as well as the reviews by Hildebrandt et al. (2023), Soeffker et al. (2022), and Psaraftis et al. (2016).

Starting with the seminal work of Campbell and Savelsbergh (2005) on an AHD system, i-DMVRPs are considered in a variety of applications such as SDD (Azi et al., 2012), MOD (Atasoy et al., 2015), or mobile personnel booking (Avraham and Raviv, 2021). Indicative of the growing importance of this research area, there are several reviews that are either applicationspecific (Li et al., 2024, Waßmuth et al., 2023, and Snoeck et al., 2020), or aim at i-DMVRP literature in general (Fleckenstein et al., 2023).

Since i-DMVRPs are dynamic and stochastic, the natural modeling approach is to formulate a Markov decision process (MDP) model (Puterman, 2014). It is important to note that MDP models not only serve as a formal problem definition. On top of that, model analysis, which can be done analytically or numerically (Bravo and Shaposhnik, 2020), yields domain knowledge that can be exploited by solution approaches. For i-DMVRPs, such model analyses can be found in Fleckenstein et al. (2024), Lebedev et al. (2021), and Asdemir et al. (2009). In particular, we draw on the property that OC can be decomposed into MCTS and DPC (Fleckenstein et al., 2024). The three works have in common that they mainly analyze models analytically. In contrast, our technique focuses on the numerical analysis of solution approaches. It also yields domain knowledge regarding a combination of a solution approach and a model of a specific i-DMVRP.

Although i-DMVRPs can be solved to optimality by exact dynamic programming algorithms, this is impractical for realistic-sized instances. Hence, there exists a wide variety of heuristic solution approaches for specific i-DMVRPs. According to Fleckenstein et al. (2023), they can be classified into two broad solution concepts: First, there are static deterministic approaches that solve auxiliary models with a rolling-horizon to iteratively derive demand management and vehicle routing decisions (e.g., Klapp et al., 2020). Second, there are decomposition- based solution approaches that subdivide the demand management task into two sub-problems as already mentioned: OC approximation and demand management decision-making.

In this work, we only consider the decomposition-based solution approaches, which are adopted by the majority of authors (Fleckenstein et al., 2023). The applied OC approximation approaches can be sub-divided further into sampling-based (e.g., Klein and Steinhardt, 2023) and learningbased (e.g., Ulmer, 2020) and are either targeted at approximating MCTS, DPC, or both (Fleckenstein et al., 2024). To evaluate the performance of these approximation approaches, authors of existing works resort to descriptive analyses. With our work, we aim at explaining the reasons for the observed results by investigating the impact of OC approximation errors.

### 2.2 Explainability of Algorithmic Performance and Behavior

Explainability becomes increasingly relevant in analytics and optimization in general (see, e.g., the recent reviews by De Bock et al. (2024) and Goerigk and Hartisch (2023)). In particular, the field of explainable RL (XRL) has recently gained more attention (Milani et al., 2024). Due to the close relation between RL and the OC approximation approaches observed in i-DMVPR literature, the techniques developed for XRL are also applicable to most approaches tackling i-DMVRPs.

Milani et al. (2024) introduce a two-dimensional taxonomy of explainability techniques tailored to XRL. The first classification dimension proposed by the authors is borrowed from general explainable artificial intelligence:

- 1.) Explainability can be *inherent* to a policy or restored *post-hoc*.
- 2.) We can further distinguish *local* explanations that refer to individual states, and *global* explanations that holistically view the behavior of the policy.
- 3.) Among the post-hoc techniques, a distinction can be made regarding the *degree of portability*, i.e., the range of solution approaches the technique can be readily applied to.

Regarding this classification dimension, our explainability technique is a *post-hoc* explanation since it is applied to a given (decomposition-based) policy. As discussed in Section 2.1, this type of policy is quite common, which makes our technique *portable*. Further, it features *local* (state-level) metrics but also involves *global* considerations since these local metrics are aggregated to explain the global behavior of the policy.

The second classification dimension specifically addresses XRL approaches and distinguishes explainability techniques based on the type of explanations they incorporate as follows:

- Feature importance explanations: explaining individual actions by providing their context, e.g., state features. Typical approaches are, e.g., surrogate policies encoded as decision trees or saliency map explanations.
- 2.) *Learning process and MDP explanations*: exploiting the definition of MDP model elements or training process steps to generate explanations. The aim is to identify critical drivers of the policy's individual decisions.
- 3.) *Policy-level explanations*: identifying recurring sequences of decisions (e.g., by clustering states) to extract patterns of the policy's overall control behavior.

Regarding this second classification dimension, our technique can be viewed as a combination of *policy level explanations* (in B1) and *learning process and MDP explanations* (in B2). In the following, we briefly review the closest related literature for B1 and B2 separately.

B1 of our explainability technique is closely related to a technique called *strategy summarization* by Amir et al. (2019). They suggest identifying states of interest on the basis of importance, coverage, likelihood of encountering, and policy disagreement with the aim of aggregating these states to summarize the behavior of the policy. Applied to i-DMVRPs, measuring the OC approximation error itself can be considered equal to measuring policy disagreement with the optimal policy. To quantify state importance, we measure the impact of an approximation error in a certain state on the quality of the resulting decision. This can also be viewed as a special case of the state importance metric used by Torrey and Taylor (2013). Further, like Amir et al. (2019), we consider the likelihood of encountering a state.

B2 of our explainability technique is a reward decomposition technique. It is first proposed by Russell and Zimdars (2003) with the aim of facilitating the learning process. With the same goal,

it is also applied by Van Seijen et al. (2017) in the form of a hybrid reward architecture. However, as shown by Juozapaitis et al. (2019), reward decomposition can not only be applied for designing hybrid reward architectures but also as an explainability technique. Therefore, they analyze the influence of the different reward components for explaining the behavior of a given policy. In contrast, we analyze approximation errors that result from considering only one reward component for explaining the behavior of a given policy. This idea of analyzing approximation errors in RL is first presented by Mannor et al. (2007) with the aim of computing confidence intervals.

Regarding the application of our explainability technique, a distinguishing feature compared to most existing works in XRL is that we consider a large number of small problem instances and solve them to optimality. Thereby we derive generic domain knowledge, in the form of fundamental OC approximation error types, rather than analyzing heuristic policies for large instances. In this regard, we only found one similar approach by Bravo and Shaposhnik (2020). They use machine learning to analyze optimal policies for small problem instances of, amongst others, traditional revenue management problems.

In summary, our methodology combines a variety of existing RL explainability techniques in a novel way: Besides adapting them to the problem structure of i-DMVRPs, we introduce the new idea of combining strategy summarization and reward decomposition and applying both to derive characterizations of fundamental OC approximation errors.

### 2.3 Performance Metrics in i-DMVRP Literature

In contrast to "pure" revenue management and dynamic pricing, where explainability has already received some attention (e.g., Biggs et al., 2021, Bravo and Shaposhnik, 2020), we find no systematic application of techniques from XRL in the literature on i-DMVRPs. Instead, most authors evaluate the performance of their solution approaches by incorporating descriptive analytics, as we summarize in the following.

*Aggregate metrics* – Apart from the arithmetic mean of profit, which is the objective in most of the considered i-DMVRPs, many authors additionally report the following aggregate metrics describing the performance of policies: Among the most widely reported metrics are average or overall revenue, cost, and number of orders (Campbell and Savelsbergh, 2005). Further, some authors also report revenue per order (Klein et al., 2018), cost per order (Yang et al., 2016), average number of fulfillment options offered to each customer (Mackert, 2019), pooling rate (Anzenhofer et al., 2024), or fleet utilization (Klein and Steinhardt, 2023). In addition to the arithmetic mean, the standard deviation (Yang et al., 2016) or the coefficient of variation (Anzenhofer et al., 2024) are reported in a few studies.

*Decision-making* – For a more detailed analysis of a policies' performance, authors analyze how the resulting decision-making differs over time, i.e., over the course of the booking horizon, or

for different types of requests: For any i-DMVRP, the acceptance rate or conversion rate (Mackert, 2019) or the cumulative revenue over time (Lang et al., 2021) can be reported.

If customers can choose from a set of fulfillment options, the number (Abdollahi et al., 2023) or composition of offered fulfillment options (Klein and Steinhardt, 2023), or the chosen fulfillment options (Anzenhofer et al., 2024) can be analyzed. If dynamic pricing is applied, average prices of offered (Klein et al., 2018) or chosen fulfillment options (Yang et al., 2016) are reported.

*Opportunity cost* – If a parametric OC approximation is used, its parameter values (Lang et al., 2021) or the function values for certain parameter values (Avraham and Raviv, 2021) can be investigated. Only very rarely, authors directly consider approximated OC values for different groups of similar requests (Yang and Strauss, 2017) or over time (Koch and Klein, 2020).

In general, we identify three central problems that limit the explanatory power of the existing descriptive analyses: First, observations of the performance and the behavior of a policy do not provide direct evidence of whether or how exactly an OC approximation error influences the observed performance. Due to a lack of conclusive explanations, the reasoning is often limited to formulating hypotheses. Second, the metrics are only analyzed in an aggregate form, which does not allow distinguishing different types of errors that originate in certain regions of the state space. Third, since typically, a specific solution approach for a specific i-DMVRP is considered, the results are hardly attributable to certain characteristics of the problem structure, the instance structure, or the solution approach. This again limits conclusiveness and transferability.

Overall, there is a clear research gap regarding the development of explainability techniques for i-DMVRPs and the formulation of generalizable explanations for policy performance.

# 3 Problem Definition and Modeling

In this section, we formally characterize i-DMVRPs with a particular focus on the generic MDP model for i-DMVRPs by Fleckenstein et al. (2024).

Typically, an i-DMVRP is structured as follows: During a *booking horizon*, customers log-in to the business platform and place a service request by entering service parameters like pick-up/drop-off locations, desired fulfillment times, or vehicle types. In response, the provider either presents a set of suitable *fulfillment options* with different prices to choose from or accepts/rejects the request. Then, a successfully placed customer request turns into a confirmed *customer order*. All customer orders are eventually served by the provider within the *service horizon*, which can either be *disjoint* or *overlapping* with the booking horizon. The former is typical for AHD, where customer and provider agree on a delivery time window for a certain day in advance. The latter is typical for SDD or MOD, where the customer expects to receive a service on short notice.

In the following, we consider the generic i-DMVRP model as in Fleckenstein et al. (2024) but adapt it specifically for the case of disjoint booking horizons and service horizons. Further, the

underlying demand control subproblem features an accept/reject demand control. However, the generalization to multi-option demand control is straightforward (see Fleckenstein et al., 2024).



Fig. 2 Overview of the MDP model of the i-DMVRP booking and fulfillment process including the interim state (Fleckenstein et al., 2024)

Decision epoch – A decision epoch marks the start of the MDP model's stages. In the considered problem, such stages correspond to (constant) time steps t = 1, ..., T. A customer request of type  $c \in C$  can arrive in stage t with a certain arrival rate  $\lambda_c^t$ . With each customer request of type c, the provider also receives data on the associated location(s)  $l_c$  and revenue  $r_c$ . Individual customer requests are then uniquely identified by combining this information with their request time  $\tau$ . Arrival rates are assumed to be small enough that at most one customer request arrives per stage. State – The system state  $s_t = (C_t, \phi_t)$  comprises two sets. The first set  $C_t$  consists of tuples  $(c, \tau, o)$ , which store customer orders for which fulfillment has not yet started. The second set  $\phi_t$ stores the tour plan. Since we assume disjoint booking and service horizons, in our case,  $\phi_t$  is either preliminary or empty for all t < T. Please note that  $s_t$  defines a post-decision state. The state space of a decision epoch t is denoted as  $S_t$  and comprises all potential realizations of customer orders  $C_t$  and tour plans  $\phi_t$ . Thus,  $\forall t \in \{1, ..., T\}$ :  $s_t \in S_t$ .

Action – An action in response to an arriving customer request of type c integrates an accept/reject decision for demand control  $g_t \in \mathcal{G}(s_{t-1}, c) \subseteq \{0,1\}$ , and a tour planning decision  $\phi_t(g_t) \in \Phi(s_{t-1}, c, g_t)$ . Again,  $\phi_t(g_t)$  is either preliminary or empty for all t < T due to the disjoint horizons. The action space for tour planning, denoted as  $\Phi(s_{t-1}, c, g_t)$ , is defined by the routing constraints of the problem and depends on the preceding state  $s_{t-1}$ , the type c of the arriving request, and the demand control decision  $g_t$ . The action space for demand control, denoted as  $\mathcal{G}(s_{t-1}, c)$ , in turn, depends on  $\Phi(s_{t-1}, c, g_t)$  since  $g_t = 1$  is only feasible if  $\Phi(s_{t-1}, c, g_t) \neq \emptyset$ . Thus,  $\mathcal{A}_t(s_{t-1}, c) = \{(g_t, \phi_t(g_t)) : g_t \in \mathcal{G}(s_{t-1}, c), \phi_t(g_t) \in \Phi(s_{t-1}, c, g_t)\}$ .

*Rewards* – As a consequence of an acceptance decision  $g_t = 1$ , a revenue  $r_c$  is received. A rejection yields no reward. A routing decision  $\phi_t(g_t)$  entails a reward  $r_{\phi_t(g_t)}$ . It equals the newly arising fulfillment cost, which, given the triangle inequality holds, is non-positive. Again, since we assume disjoint booking and service horizons,  $\forall t < T : r_{\phi_t(g_t)} = 0$ .

*Transition* – When transitioning to state  $s_t$ ,  $\phi_t$  is set to  $\phi_t(g_t)$ . The first state component  $C_{t-1}$  also changes. More precisely, if the newly arriving request of type  $c_t$  is accepted, the resulting customer order is added.

*Objective* – The provider aims at maximizing profit after fulfillment. Therefore, it is required to determine a policy  $\pi$  that returns the optimal decision for each state that can potentially be reached. These decisions of a policy  $\pi$  can be denoted as  $a_t^{\pi}(s_{t-1}, c_t) = (g_t^{\pi}(s_{t-1}, c_t), \phi_t^{\pi}(g_t^{\pi}(s_{t-1}, c_t)))$  at decision epoch t. Then, the objective function is:

$$\max_{\pi} \mathbb{E} \left( \sum_{t=1}^{T} \left( r_{c_t} \cdot g_t^{\pi}(s_{t-1}, c_t) + r_{\phi_t^{\pi}(g_t^{\pi}(s_{t-1}, c_t))} \right) \mid s_0 \right).$$
(1)

Bellman equation – The objective function (1) can be expressed in the form of a Bellman equation, which defines a value  $V_t(s_t)$  for each state  $s_t$ . Solving this equation yields the optimal policy  $\pi^*$ .

$$V_{t-1}(s_{t-1}) = \sum_{c \in C} \lambda_c^t \cdot \max_{g_t \in \mathcal{G}(s_{t-1},c)} \left( g_t r_c + \max_{\phi_t(g_t) \in \Phi(s_{t-1},c,g_t)} \left( r_{\phi_t(g_t)} + V_t(s_t \mid s_{t-1},\phi_t(g_t)) \right) \right)$$

$$+(1-\sum_{c\in C}\lambda_{c}^{t})\cdot\max_{\phi_{t}(0)\in\Phi(s_{t-1},0,0)}\left(r_{\phi_{t}(0)}+V_{t}\left(s_{t}\mid s_{t-1},\phi_{t}(0)\right)\right),$$
(2)

(3)

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

In Equation (2), both types of decisions are represented in an integrated form. Thus, an interim state  $s'_t | s_{t-1}, c, g_t$  can be defined to isolate the impact of the demand control decision from the impact of the vehicle routing decision as also depicted in Fig. 2. Further, substituting the OC of accepting a request of type c, i.e.,  $\Delta V_t(s_{t-1}, c)$ , we obtain the following reformulation. Note that we denote interim states  $s'_t | s_{t-1}, c, 1$  by  $s'_t(c)$  and interim states  $s'_t | s_{t-1}, c, 0$ , or  $s'_t | s_{t-1}, 0, 0$ , by  $s'_t(0)$ .

$$V_{t-1}(s_{t-1}) = \sum_{c \in C} \lambda_c^t \cdot \max_{g_t \in \mathcal{G}(s_{t-1},c)} \left( g_t \cdot \left( r_c - \Delta V_t(s_{t-1},c) \right) \right) + V_t'(s_t'(0)), \tag{4}$$

with 
$$V'_t(s'_t | s_{t-1}, c, g_t) = \max_{\phi_t(g_t) \in \Phi(s_{t-1}, c, g_t)} \left( r_{\phi_t(g_t)} + V_t(s_t | s_{t-1}, \phi_t(g_t)) \right)$$
  
=  $r_{\phi_t^*(g_t)} + V_t(s_t | s_{t-1}, \phi_t^*(g_t)),$  (5)

and 
$$\Delta V_t(s_{t-1}, c) = V'_t(s'_t(0)) - V'_t(s'_t(c)) \ge 0.$$
 (6)

Further,  $\phi_t^*(g_t)$  denotes the optimal routing decision for a given demand management decision  $g_t$ .

#### 4 Explainability Technique

In this section, we present our novel explainability technique for i-DMVRPs, which comprises two separate building blocks. Both are later applied for the comprehensive analysis of the relation between OC approximation error and the quality of the resulting demand management decisions.

### 4.1 Building Block 1

The basic idea of B1 is to define metrics for each step in the chain of influencing factors (Fig. 1) behind the losses in objective value observed when following a certain policy. By this, we aim at identifying the regions of the state space that are especially relevant regarding the respective overall objective value loss. Therefore, we analyze the occurrence, the sign, and the magnitude of OC approximation errors in the respective states. Then, by suitable visualizations of the metrics, we compare them over various settings of problem parameter values (in the following referred to as *settings*) resembling different real-world i-DMVRPs. Based on that, we classify fundamental types of approximation errors, i.e., OC approximation errors that a broad variety of real-world i-DMVRPs are prone to. Therewith, we can eventually explain the performance of the considered policy. Both the metrics and the respective visualizations are generally valid, i.e., can be applied to any policies derived from different OC approximation approaches. In the following, we first describe the chain of influencing factors between OC approximation error and objective value loss. Afterwards, we describe the metrics we use to quantify each step in this chain of influencing factors and, finally, we propose visualizations of these metrics.

#### Chain of influencing factors

At the beginning of the chain of influencing factors, there is an approximation error in a certain state  $s_{t-1}$ , which could either be an underestimation or an overestimation of the true OC. Depending on the actual magnitude of such an approximation error, the magnitude of the true OC, i.e.,  $\Delta V_t(s_{t-1}, c)$ , and the immediate reward  $r_c$ , this error can but not necessarily must result in a suboptimal decision.

Generally, a suboptimal decision in a certain state  $s_{t-1}$  can either yield less immediate reward than the optimal decision, transition the system to a lower-valued state than the optimal decision, or both. However, the respective negative effect on the objective value itself can vary from barely notable to considerable.

Whether the chain of influencing factors continues further, depends on the likelihood that  $s_{t-1}$  is encountered and the respective suboptimal decision is made when following the policy under consideration.

We now define disaggregated metrics to quantify "how bad" an OC approximation error is, "how wrong" the resulting decision is, and also, "how likely" this decision is. Additionally, we define a fourth metric that captures the aggregated overall impact of OC underestimations or OC overestimations on the objective value.

#### Metrics

Fig. 3 shows which of the metrics presented in the following corresponds to which step in the previously described chain of influencing factors between OC approximation error and objective

value loss. Please note, we assume that we examine an OC approximation relative to the true OC, i.e., a suboptimal policy relative to the optimal policy.



Fig. 3 Metrics among the chain of influencing factors from OC approximation error to objective value loss

(1) Error magnitude – In general, an OC approximation can be inaccurate in both directions. By separately evaluating overestimation and underestimation, we can analyze whether the consequences are different. Hence, we calculate the magnitude of overestimation errors  $e^o(s_{t-1}, c)$  and underestimation errors  $e^u(s_{t-1}, c)$ , separately, for all individual states  $s_{t-1}$  and request arrivals c as follows:

$$e^{o}(s_{t-1},c) = \Delta \tilde{V}_{t}(s_{t-1},c) - \Delta V_{t}(s_{t-1},c),$$
(7)

$$e^{u}(s_{t-1},c) = \Delta V_{t}(s_{t-1},c) - \Delta \tilde{V}_{t}(s_{t-1},c).$$
(8)

This metric allows us to identify regions of the state space, where an approximation systematically overestimates ( $e^o(s_{t-1}, c) > 0$ ) or underestimates ( $e^u(s_{t-1}, c) > 0$ ) the true OC. Hence, it provides information about where approximation errors originate and how strongly the chain of influencing factors is triggered.

(2) Single decision regret – As is well-known in revenue management (e.g., Talluri and Van Ryzin, 2004), an OC approximation error in itself is not problematic because the resulting decision may still be fairly accurate or even optimal. As previously described, the chain of influencing factors only continues if there is a suboptimal decision. Hence, to quantify the "suboptimality" of a single decision when a request of type c arrives in state  $s_{t-1}$ , we introduce the metric single decision regret, denoted as  $\delta(s_{t-1}, c)$ . It computes the overall reward difference between a single decision based on a, potentially wrong, OC approximation and the optimal decision. In less technical terms, to isolate the regret of one single decision, all future decisions from decision epoch t + 1 onward are assumed to be made based on the optimal policy in both cases. Then, again depending on the observed underlying error magnitude  $e^o(s_{t-1}, c)$  and  $e^u(s_{t-1}, c)$ , we can distinguish between overestimation regret and underestimation regret, even though both are calculated based on the same expression:

$$\delta(s_{t-1},c) = g_t^*(s_{t-1},c) \cdot (r_c - \Delta V_t(s_{t-1},c)) - \tilde{g}_t(s_{t-1},c) \cdot (r_c - \Delta V_t(s_{t-1},c)),$$
(9)

with  $g_t^*(s_{t-1}, c)$  denoting the optimal demand control decision and  $\tilde{g}_t(s_{t-1}, c)$  denoting the demand control decision when following the policy under consideration. Then, we define the overestimation regret as: Article A3: From Approximation Error to Optimality Gap – Explaining the Performance Impact of Opportunity Cost Approximation in Integrated Demand Management and Vehicle Routing

$$\delta^{o}(s_{t-1}, c) = \begin{cases} \delta(s_{t-1}, c), e^{o}(s_{t-1}, c) > 0\\ 0, & \text{otherwise,} \end{cases}$$
(10)

and the underestimation regret as:

$$\delta^{u}(s_{t-1},c) = \begin{cases} \delta(s_{t-1},c), e^{u}(s_{t-1},c) > 0\\ 0, & \text{otherwise.} \end{cases}$$
(11)

With this metric, we can assess whether an overestimation error or underestimation error leads to a suboptimal decision, and by which amount it causes the objective value to deteriorate assuming optimal decisions over the remaining booking process. Please note, for ease of readability, in the following, we refer to the single decision regret as *regret*.

(3) Decision rate – As a third step in the chain of influencing factors, the relevance of a suboptimal decision must be considered. It depends on how likely it is to visit the state in which the decision is made. We measure the likelihood in the form of the decision rate  $P(s_{t-1}, c)$ , which denotes the probability that a policy  $\tilde{\pi}$  visits state  $s_{t-1}$  and decides on the acceptance/rejection of a customer request of type c at decision epoch t. To calculate it, we simulate decision-making based on the considered OC approximation for a sufficiently high number of drawn sample paths. This metric provides information about the relevant areas of the state space, and thus, to what extent the regret in a certain state impacts the objective value.

(4) Weighted error ratio – Additionally to the metrics involved in our chain of influencing factors, we propose an aggregate metric to approximate the share of the loss in objective value that is caused by overestimation and underestimation, respectively. We refer to this metric as *weighted* error ratio and define it for a certain setting as follows:

$$E = \frac{\sum_{s \in \mathcal{S}} \sum_{c \in \mathcal{C}} \delta^o(s, c) \cdot P(s, c)}{\sum_{s \in \mathcal{S}} \sum_{c \in \mathcal{C}} (\delta^o(s, c) + \delta^u(s, c)) \cdot P(s, c)}$$
(12)

Thus, a weighted error ratio of 1 means that the full observed objective value loss of a certain policy is caused by overestimation errors, and a weighted error ratio of 0 means that only underestimation causes objective value loss.

#### Visualizations

Disaggregated metrics (1)-(3) – Due to the curses of dimensionality, analyzing the values of the metrics (1)-(3) for individual states as introduced above is impractical. Hence, we propose a state space aggregation for their visualization and the subsequent analysis. For this, we draw on two-dimensional lookup tables, which are well-known from reinforcement learning (e.g., Powell, 2022). We aggregate states according to the dimensions *decision epoch t* (*y*-axis) and *capacity consumption* (*x*-axis), where we measure the latter in percentage of the available capacity. Since we consider multiple instances per setting, we further aggregate the results of all instances per setting in one look-up table by averaging the respective numbers. In sum, these aggregations allow us to examine the average magnitude of overestimation and underestimation, the associated regret, and the decision rate in the different regions of the state space. Additionally, all instance-

specific effects are averaged out. Further, this two-dimensional aggregation allows us to represent the resulting look-up tables in heatmaps, as exemplarily depicted in Fig. 4.







<sup>(</sup>b) Policy B: Average objective value reached = 12.447



Now, it is possible to qualitatively analyze and compare the results for different policies as exemplarily demonstrated in the following: In the example presented in Fig. 4, we analyze two suboptimal policies A and B that, over the same set of instances, yield entirely different average objective values. However, by analyzing the heatmaps of our metrics (calculated based on the optimal policy), we can now explain these drastic results: Policy A suffers from severe underestimation errors that cause high regret in particularly relevant areas of the state space. Overestimation, in turn, also occurs but does not cause any regret. Contrary, policy B only exhibits mild underestimation with substantially lower regret in rather irrelevant states. Further, for policy B, overestimation also causes regret that is even slightly more relevant. Overall, the aggressive acceptance of (early) customers due to underestimation errors by Policy A causes a severe objective value loss. In comparison, the (slightly too) conservative behavior of Policy B leads to a much better performance. Hence, if we were to develop a policy for an i-DMVRP with this setting structure, we conclude that Policy A is missing crucial information. Hence, in the development process, we can now, e.g., integrate algorithmic elements to tackle the systematic underestimation by Policy A, or draw the conclusion to focus more development effort on Policy B due to its structural advantages.

Aggregated metric (4) – To visualize and interpret the fourth metric, i.e., the weighted error ratio, we propose a scatter plot, where the result of applying a policy to a certain setting is plotted as a point according to the weighted error ratio on the *x*-axis and the relative optimality gap on the *y*-axis (see Fig. 5). This visualization enables us to analyze in which settings a policy is prone to either underestimation or overestimation in combination with the resulting performance impact. The example given in Fig. 5 shows a policy that is mainly affected by underestimation errors in Setting 1 and Setting 2, with a more severe impact in Setting 1. Setting 3 is equally affected by overestimation.



Fig. 5 Weighted error ratio - exemplary for four settings

#### 4.2 Building Block 2

To apply the previously presented technique for gaining general insights on i-DMVRPs within our computational study, we now rely on a second explainability technique, namely on reward decomposition (Juozapaitis et al., 2019). More precisely, although B1 can be applied to *any* OC approximation in comparison to the optimal policy and already yields valuable insights, we apply it to *specific* OC approximations that base on the idea of reward decomposition. This comes with two advantages: First, the resulting policies resemble typical approximation approaches from literature that (predominantly) capture one of the reward components. Second, and more importantly, they reveal which reward component is more relevant in a certain setting of a given i-DMVRP and, therewith, hold further explainability potential.

Thus, to define the policies we analyze in Section 5, we draw on the results of Fleckenstein et al. (2024). They show that the OC of an i-DMVRP,  $\Delta V_t(s_{t-1}, c)$ , can be decomposed into two components: DPC, formally denoted as  $\Delta R_t(s_{t-1}, c)$ , and MCTS, formally denoted as  $\Delta F_t(s_{t-1}, c)$ . While the former captures the loss of future revenue, i.e., a decision's impact on the positive

rewards, the latter measures the respective increase of fulfillment cost, i.e., the impact on the negative rewards. For the formal definition, we first define the expected future revenue of a given interim state  $s'_{t-1}$  at decision epoch t - 1 as  $R'_{t-1}(s'_{t-1})$  and the expected future fulfillment cost of a given interim state  $s'_{t-1}$  at decision epoch t - 1 as  $F'_{t-1}(s'_{t-1})$ . Then, we can define DPC and MCTS as follows:

$$\Delta R_t(s_{t-1}, c) = R'_t(s'_t(0)) - R'_t(s'_t(c))$$
(13)

$$\Delta F_t(s_{t-1}, c) = F'_t(s'_t(0)) - F'_t(s'_t(c)), \tag{14}$$

and  $\Delta V_t(s_{t-1}, c) = \Delta R_t(s_{t-1}, c) + \Delta F_t(s_{t-1}, c)$  holds.

Based on this finding, we derive and analyze OC approximations that only capture DPC or MCTS, i.e., completely neglect the other component. By considering these most extreme cases, we make sure that any systematic errors resulting from inadequately approximating one component occur as clearly as possible. Further, to reduce random errors related to the varying performance of a heuristic solution algorithm, we formulate the Bellman equations corresponding to our policies and solve them in an exact way by backwards recursion. The DPC-based approximation  $\Delta \tilde{R}_t(s_{t-1}, c)$  results from:

$$\tilde{R}_{t-1}(s_{t-1}) = \sum_{c \in C} \lambda_c^t \cdot \max_{g_t \in \mathcal{G}(s_{t-1},c)} \left( g_t \cdot \left( r_c - \Delta \tilde{R}_t(s_{t-1},c) \right) \right) + \tilde{R}_t'(s_t'(0))$$
(15)

Since the DPC-based approximation neglects fulfillment cost, i.e.,  $r_{\phi_t(s'_t)} = 0 \forall t = 1, ..., T$ , the revenue is the only type of reward, and the future cost impact of an acceptance decision is ignored. Solving Equation (15), we obtain the DPC policy  $\pi^R$ , the DPC-based decision  $g_t^R(s_{t-1}, c) = \underset{g_t \in \mathcal{G}(s_{t-1}, c)}{\operatorname{argmax}} \left( g_t \cdot \left( r_c - \Delta \tilde{R}_t(s_{t-1}, c) \right) \right)$  being in state  $s_{t-1}$  and observing the arrival of a request of type c, and the DPC-based objective value  $J^R = \tilde{R}_0(s_0)$ .

The MCTS-based approximation  $\Delta \tilde{F}_t(s_{t-1}, c)$  results from:

$$\tilde{F}_{t-1}(s_{t-1}) = \sum_{c \in C} \lambda_c^t \cdot \operatorname*{argmax}_{g_t \in \mathcal{G}(s_{t-1},c)} \left( g_t \cdot \left( r_c - \Delta \tilde{F}_t(s_{t-1},c) \right) \right) \cdot \left( -\Delta \tilde{F}_t(s_{t-1},c) \right) + \tilde{F}_t'(s_t'(0)).$$
(16)

Since the MCTS-based approximation neglects displacement cost, formulating the Bellman equation of this approximation requires an  $argmax(\cdot)$  operator. This prevents the revenue from being included in the state value while still comparing it to the future cost impact for deciding on each single request's acceptance. Then, the result of the  $argmax(\cdot)$  operator, which encodes the binary demand control decision, is multiplied with the future cost impact, i.e., the MCTS. Solving Equation (16), we obtain the MCTS policy  $\pi^F$ , the MCTS-based decision  $g_t^F(s_{t-1}, c) =$  $\underset{g_t \in \mathcal{G}(s_{t-1}, c)}{argmax} \left(g_t \cdot \left(r_c - \Delta \tilde{F}_t(s_{t-1}, c)\right)\right)$  being in state  $s_{t-1}$  and observing the arrival of a request of type c, and the MCTS-based objective value  $J^F = \tilde{F}_0(s_0)$ .

# 5 Computational Study

We now apply the explainability technique consisting of B1 and B2 to a generic i-DMVRP as introduced in Section 3. The aim of this numerical analysis is to identify and characterize fundamental types of approximation errors that can occur in any real-world i-DMVRP. In Section 5.1, we first present the experimental design, the tested parameter settings, and the benchmarks that we consider next to the already introduced DPC-based and MCTS-based approximations. Then, in Section 5.2, we characterize the observed types of approximation errors. Finally, we investigate the impact of these types of errors on the objective value in Section 5.3 to further refine the characterization.

### 5.1 Experimental Design

Overall, we draw on the full-factorial study design proposed in Fleckenstein et al. (2024), which consists of 66 different settings of an i-DMVRP with disjoint booking horizon and service horizon and pure accept/reject decisions. The integrated VRP is a distance-constrained, capacitated VRP in all settings. Further, all settings have in common that we assume a single fulfillment vehicle and a booking horizon of T = 10 potential decision epochs with 10 potentially arriving customer requests such that at most one customer request arrives per decision epoch.

The settings differ in the following parameters: *location distribution, revenue distribution*, general *profitability* of a setting, and binding *capacity constraints*. The former two parameters are customer-related, the latter two are provider-related. In detail, the following parameter values are possible:

Location distribution – Regarding the customers' location distribution, we consider two realizations: A stream of customer request locations  $l_c$  with c = 1,...,10 are either drawn from (1) a uniform distribution over a line segment with length 50 in the interval [-25,25] representing a single urban area, or (2) in random order from two truncated normal distributions with means -10 and 20 and the same standard deviation of 2.5 to generate two equal-sized clusters representing, e.g., two villages in a rural area. The first location distribution is referred to as *unif* for uniform distribution. The second location distribution is referred to as *clust* since there is two clusters of customers.

*Revenue distribution* – Regarding the customers' revenue distribution, we consider four realizations: Besides (1) homogeneous customer streams with all revenues equal to 15 monetary units referred to as *homog*, (2) heterogeneous customer streams are generated by randomly assigning a revenue of 25 monetary units to 30% of the customers in a stream of customer requests. The other 70% of the customers are assigned revenues equal to 15 monetary units. If heterogeneous customers are considered in a setting, the high-revenue customers can either (2.1) strictly arrive in the beginning (referred to as *h-b-l* for high-before-low), (2.2) randomly (referred to as *rand*), or (2.3) strictly in the end (referred to as *l-b-h*} for low-before-high) of the booking horizon. Further, for each of the realizations (2.1)-(2.3), an additional clustered setting is considered in which all high-revenue customers are located in the distant cluster. We refer to these settings as *clust\_sort*. Also, if required (as in *h-b-l* and *l-b-h* settings), we sort the drawn customer streams. These customer-related parameter realizations yield 11 meaningful combinations (referred to as *customer settings*), as depicted as leaf nodes in Fig. 7 in Appendix A. Each of these settings is then considered six times, according to six combinations of the two provider-related parameters' realizations that we explain in the following.

*Profitability* – To vary the general profitability of the settings, we modify the routing cost factor. More precisely, each customer setting is considered three times with different routing cost factors of 0.2, 0.6, or 1 monetary units per distance unit. We refer to these settings by addressing the profitability as *high*, *med*, or *low*.

*Capacity constraints* – Regarding the logistical capacity constraints, we assume two different realizations. We either limit the route length to 50 length units and refer to these settings as *distance-constrained* (*dist*), or we limit the physical capacity of the fulfillment vehicle to 3 units. In the latter case, we assume unit demand for all customers and refer to the respective settings as *load-constrained* (*load*).

For each of the resulting 66 settings, we consider 50 different instances in our computational study. Each instance is defined by an individual customer stream of 10 customers, sampled according to the above-mentioned customer-related parameter values in advance. Of these customer streams, each customer then places a request with probability  $\lambda_c^t = 0.5$  if c = t and  $\lambda_c^t = 0$ , otherwise. This is the only source of stochasticity once an instance is fully specified. We apply the policies presented in Section 4.2 to each instance, calculate the metrics introduced in Section 4.1, and average the respective results over all 50 instances per setting to derive the setting-specific results we report.

Note, the choice of such a basic problem is deliberate, and typically done in literature, to obtain findings that are valid for a broad variety of real-world problems (e.g., Ulmer and Thomas, 2020). *Benchmark policies* – As a reference for evaluating the performance of the previously introduced DPC- and MCTS-based approximations, we mainly consider two benchmark approaches, which we also apply to all instances as previously described. First, we solve Equation (4) to compute the true OC  $\Delta V_t(s_t)$ . Thereby, we also obtain the optimal policy  $\pi^*$  and the optimal objective value  $J^* = V_0(s_0)$  as well as the optimal decision  $g_t^*(s_{t-1}, c) = \underset{g_t \in \mathcal{G}(s_{t-1}, c)}{\operatorname{argmax}} \left( g_t \cdot (r_c - \Delta V_t(s_{t-1}, c)) \right)$ .

Second, we consider a myopic OC approximation due to its high relevance as a benchmark policy in the i-DMVRP literature (e.g., Arian et al., 2022, Klein and Steinhardt, 2023, or Yang et al., 2016). Instead of a state value difference, the respective OC approximation is defined as the

insertion cost into the current myopic route plan, i.e., a route plan based on only confirmed customer orders.

The objective values of all policies are depicted in Fig. 8 in Appendix B.

### 5.2 Identification of Fundamental Approximation Errors

In the following, we identify and characterize the fundamental types of approximation errors that we observe for the DPC policy and the MCTS policy. This analysis is based on the heatmaps introduced in Section 4.1. Since it is not possible to include the full set of heatmaps for all settings in the paper at hand, we provide two types of supplementary material for the interested reader. First, we present carefully selected heatmaps in Appendix C. These are intended to be used as "textbook" examples that show a certain error and its characteristics particularly clearly. Second, to allow for full reproducibility, the complete set of heatmaps can be viewed or downloaded at zenodo.org.

When discussing an error type's characteristics in the following, especially regarding the influence of setting parameters, we only mention those error types that are reasonably pronounced and occur over many different settings. Thereby, we ensure generalizability beyond the basic settings we consider. Then, we explain the prevalence of these error types in the different settings as well as their interplay.

### **DPC** policy

The DPC policy lacks information on how well a request can be consolidated with other orders to a profitable fulfillment tour. It can only observe this information indirectly if it is reflected in the request's capacity consumption, and thus, affects the displacement of future revenue. Apart from that, the policy assumes perfect consolidation of orders, i.e., MCTS of zero for each request. In the following, we introduce the three distinctive types of approximation errors resulting from this DPC policy. More precisely, we discuss two types of underestimation errors and one type of overestimation error, that occur when applying the DPC policy.

### <u>Underestimation error type 1</u>: neglecting better consolidation of future requests

This error occurs when the acceptance of the current request causes the displacement of future requests with a similar revenue that can be consolidated better with other future requests or orders already received. This type of underestimation error becomes smaller over time as there are less opportunities for collecting orders that allow better consolidation (see Fig. 9a). In the case that future requests have a higher revenue than the current request, as especially visible in our *l-b-h* settings, the policy protects a certain amount of capacity. This reduces the error's regret since the DPC are sufficiently high to correctly cause a rejection (compare Fig. 9a and Fig. 9b). The regret also becomes smaller with more correlation between consolidation and capacity consumption, as observable in our *dist* constrained settings (compare Fig. 9b and Fig. 9c). In general, the error occurs in a broad region of the state space with a relatively high decision rate (see Fig. 9).

#### Article A3: From Approximation Error to Optimality Gap – Explaining the Performance Impact of Opportunity Cost Approximation in Integrated Demand Management and Vehicle Routing

#### <u>Underestimation error type 2</u>: neglecting inherent request unprofitability

This error mainly occurs when the current request is neither profitable based on consolidation with the orders already received nor is it expected to become profitable taking future consolidation opportunities into account. Hence, the effect becomes stronger over time with shrinking future consolidation opportunities and tends to occur in states with low capacity consumption where not much consolidation is already established (see Fig. 9a). The associated regret is only dampened when high-revenue requests arrive in the affected states as especially observable in our *l-b-h* settings (see Fig. 9b). Compared to underestimation error type 1, considering the decision rate reveals that this error type has much less impact because the region of late states with low capacity consumption is hardly visited in any of our settings (see Fig. 9a-9c). However, if revenues reflect the distance from the depot as in our *clust\_sort* settings, the error can occur earlier, i.e., in states with a higher decision rate (see Fig. 9d).

In extreme cases, the underestimation due to these two types of errors can become large enough that an acceptance decision is made in every state.

#### Overestimation error: protecting capacity for high-revenue demand with poor consolidation

The missing information on consolidation opportunities can also lead to an OC overestimation. This error occurs when the policy anticipates demand that has indeed higher revenue but cannot be consolidated as well as the current request. In such a state, it is better to accept the current customer request despite its lower revenue than to reserve capacity for the higher-revenue demand that is less profitable due to poor consolidation. Consequently, a heterogeneous revenue distribution that is not strictly *h-b-l* is prone to this error (compare Fig. 10a and Fig. 10b). We then can locate the error in all states early in the booking horizon but it only leads to regret in states in which displacement effects are sufficiently strong, i.e., states with high capacity consumption in case of load-constrained settings or states with low capacity consumption in case of dist-constrained settings (compare Fig. 10a and Fig. 10c). It remains constant over time until the highrevenue demand is expected to realize. Its magnitude is proportional to the mismatch between revenue and consolidation opportunities offered by the demand, or in other words, between revenue displacement observed by the policy and the actual profit displacement. A particularly strong occurrence can be observed in low, clust sort settings in which the high-revenue requests also cause the highest routing cost (see Fig. 10d). In the region of the state space, in which this error type occurs, the decision rate is comparatively high (see Fig. 10).

#### **MCTS policy**

The MCTS policy has information about all parameters of expected future requests including their revenue. However, it only uses the information about the revenue when making a decision on an individual request, as revenues generally do not enter its value function. Thus, the policy can only anticipate which future requests it will accept in case of any realization of future demand and the

resulting fulfillment cost but not the associated cumulative revenue. For this policy, we find three distinctive estimation errors.

#### <u>Underestimation error type 1</u>: neglecting future high-revenue demand

Since the MCTS policy cannot observe heterogeneity in revenue, it fails to reserve capacity if there is more profitable, high-revenue demand arriving at a later phase of the booking process as especially observable in our l-b-h settings. The resulting regret due to earning less revenue per order is roughly proportional to the magnitude of the underestimation error. The error becomes stronger the more capacity is consumed and the earlier a state is (see Fig. 11a). The error magnitude is insensitive toward the profitability (compare Fig. 11a and Fig. 11b).

#### <u>Underestimation error type 2</u>: missing information on the volume of future demand

Even if the revenues are homogeneous, underestimation occurs when applying the MCTS policy since its OC approximations do not fully reflect the volume of demand to come. Only in the end of the booking horizon, the low volume of demand to come impacts the approximation via the decreasing number of consolidation opportunities. In earlier states, this information can only be derived from the expected revenue to come, and thus, the displacement cost. The consequence of this underestimation is that the policy does not protect capacity for requests offering better consolidation, i.e., are less costly to serve. This can either lead to less efficient routing, which reduces profitability on the cost side, or it may also reduce the number of accepted orders due to the higher capacity consumption, which affects the revenue side. The error and the associated regret occur predominantly in early states (see Fig. 11c). As for underestimation error type 1, the error magnitude does not change with profitability (compare Fig. 11c and Fig. 11d).

Both types of MCTS underestimation errors are usually superimposed on each other. Similar to the type 1 underestimation error of the DPC policy, they occur in a fairly large region of the state space with a high decision rate. In the worst case, the combination of both errors leads to solutions with less accepted orders, less revenue per order, and an inefficient routing.

#### Overestimation error: wrong cost attribution

The missing information on future revenues can also lead to overestimation errors by the MCTS policy. This is caused by the following mechanism: If none or only very few orders are confirmed already, accepting a request momentarily entails a high increase in routing cost compared to rejecting it. If the instance is sufficiently profitable, anticipating optimal future decision-making would reveal that the final cost difference between accepting the current request and rejecting it would not be as high because of future consolidation. The MCTS policy, however, anticipates that in the reject-case, no or only few additional requests would be accepted, again because of their momentarily high increase in routing cost compared to a rejection. In summary, the policy wrongly attributes high routing cost to early arriving requests.

The error and the resulting regret become more severe with lower profitability (compare Fig. 12a and Fig. 12b or Fig. 12c and Fig. 12d), and its recursive character makes this error particularly harmful. At worst, it can lead to the policy stalling in states with no or few orders, which becomes apparent from the plots showing the decision rate (see Fig. 12b and Fig. 12d). The reason is that suboptimal rejections due to overestimation lead to successor states in which the overestimation error is likely to occur again. Thus, the error occurs mainly in states with no or a very short tentative route and becomes more severe with decreasing profitability. By contrast, error and regret are reduced by heterogeneous revenues. In addition, the region of high decision rates shifts away from the states with low capacity consumption, where the error occurs, which indicates that the vicious cycle is broken early in the booking process. Responsible for this are high-revenue requests that are accepted despite the overestimation (compare Fig. 13a and Fig. 13b). These orders then serve as seed customers to establish consolidation independent from their distribution over time. Similarly, a uniform distribution of locations (*unif*) can reduce the error due to orders located close to the depot serving as seed customers (compare Fig. 13a and Fig. 13c).

### 5.3 Resulting Performance Impact

Using the weighted error ratio and its visualization introduced in Section 4.1, we now analyze the contribution of underestimation errors and overestimation errors to the losses in objective value relative to the optimal solution. Again, we consider the DPC policy and the MCTS policy in the 66 different settings. The results are depicted in Fig. 6. As discussed in the following, there are some clearly observable patterns between the optimality gap, the weighted error ratio, and certain setting parameter values. From these patterns, we can derive insights on how location distribution, revenue distribution, and profitability cause underestimation errors or overestimation errors when applying a given policy and how this affects the optimality gap. However, depending on the policy applied, there may hardly be any patterns, as we observe for our MCTS policy. This shows that there can be complex interactions between underestimation and overestimation errors, e.g., both can offset.

#### **DPC** policy

When applying the DPC policy, the average optimality gap generally decreases with increasing profitability, but the weighted error ratio stays almost constant with only a slight shift toward underestimation. The type of the binding capacity constraints also has an influence on performance since the average optimality gap is smaller for *dist*-constrained settings, and there is more observations of settings in which the optimality gap results mainly from overestimation. Both scenario parameters combined, the observations reach from an average optimality gap below 1% in *high*-profitability, *dist*-constrained settings, to even negative objective values in some *low*-profitability, *load*-constrained settings.

Generally, in *clust* settings, the average optimality gap tends to be larger than in *unif* settings. Since we only observe a slight shift toward overestimation, this is mainly caused by more severe regret from underestimation errors. If the requests' revenues are proportional to their distance from the depot, i.e., in *clust\_sort* settings, the share of overestimation increases but without a clear impact on the average optimality gap. For *med*-profitability settings, the gap becomes smaller but it increases for *low*-profitability settings.



(a) DPC policy





Fig. 6 Weighted error ratio and average optimality gap

Regarding the revenue distribution, we find the highest average optimality gap for our *homog* settings and the smallest for *h-b-l* settings. In both types of settings, there is no overestimation in case of *load*-constrained settings and only a small impact of overestimation in *dist*-constrained settings. In case of heterogeneous revenues, i.e., in our *h-b-l*, *l-b-h*, or *rand* settings, the average optimality gap and the share of overestimation decrease the earlier the high revenue arrives, as especially observable in our *h-b-l* settings.

#### **MCTS** policy

For the MCTS policy, the patterns are less clear. In *high*-profitability settings, underestimation is by far the dominant error. Due to the dominating underestimation error, the average optimality gap is higher the more high-revenue demand arrives late, i.e., in *l-b-h* settings, as well as in *load*constrained settings, where less consolidation is possible. Despite the error ratio shifting toward overestimation with the setting becoming less profitable, this has no coherent impact on the average optimality gap because of the complex interactions of underestimation errors and overestimation errors. E.g., we observe average optimality gaps increasing consistently with decreasing profitability in *h-b-l* settings that are *dist*-constrained with *unif* or *clust* locations, and *load*-constrained, *clust\_sort* settings. In *load*-constrained settings with *unif* or *clust* location distribution, however, the average optimality gap decreases between *med*-profitability and *low*-profitability settings due to offsetting errors for *rand* and *l-b-h* settings. Another example for these inconsistencies are settings with *homog* revenues. Here, the average optimality gap is much higher for *unif* settings compared to *clust* settings, when considering *med* profitability. In *low*-profitability settings, we observe an inverse relation.

Regarding the location distribution, the general findings are that *unif* settings tend to have a higher average optimality gap compared to *clust* settings and that *low*-profitability, *clust\_sort* settings are prone to overestimation.

Considering the distribution of revenues, we find that *homog* or *l-b-h* settings tend to have a higher average optimality gap.

# 6 Insights for Algorithm Selection and Algorithm Design

In this section, we formulate five actionable insights based on our findings that support developers and users of OC approximation approaches in selecting and designing algorithmic components for specific i-DMVRPs. By closely incorporating the key computational results of the existing literature, we now show that our findings explain much of the performance differences observed in literature. Therewith, we also compile evidence for that our results derived from the numerical analysis of a generic i-DMVRP are fundamental to the entire family of i-DMVRPs.

**Opportunity cost approximation errors can be grouped into a few distinct types**: Neglecting one of the OC components causes systematic errors that can be grouped into a small number of fundamental types consistently observable across a variety of settings. Furthermore, we find many patterns regarding their occurrence in certain regions of the state space and in different settings. While some of those patterns were already suspected in existing literature based on the decision-making behavior of the respective solution approaches, we are the first to provide direct, numerical evidence to characterize these patterns and are also able to identify entirely novel ones. E.g., the existence of underestimation errors when applying an MCTS policy has been discussed but not conclusively proven by Mackert (2019) and Yang and Strauss (2017). Overestimation, however, has only been briefly mentioned by one author (Mackert, 2019). Overall, our main finding

is that neglecting either component can lead to both overestimation errors and underestimation errors. As we show in the work at hand, both are systematic in the sense that they occur even if an exact algorithm, i.e., applying backwards recursion to the (modified) Bellman equations (4), (15), and (16), is used to compute the respective OC estimate.

Underestimation is the dominant error when neglecting DPC or MCTS: The fundamental issue with neglecting DPC or MCTS is that the resulting approximation does not reach the correct absolute level of the true OC since the contribution of the other component is missing. Hence, underestimation is by far more common than overestimation (see Fig. 6): For the DPC policy, underestimation (weighted error ratio < 0.5) is dominant in 89.4% of all settings. For the MCTS policy, this occurs in 69.7% of the settings. This dominance of underestimation leads to greedy decision-making overall. The fact that many authors report such greedy behavior in studies with realistic-sized instances shows the general validity of this result. This accounts for both availability control (e.g., Arian et al., 2022, Lang et al., 2021, Mackert, 2019, and Campbell and Savelsbergh, 2005), and dynamic pricing (e.g., Abdollahi et al., 2023, Klein and Steinhardt, 2023, Klein et al., 2018, Yang and Strauss, 2017, and Yang et al., 2016). In the former case, the policy accepts too many customers or offers too many, or even all of the feasible fulfillment options to each requesting customer resulting in a first-come-first-served decision-making behavior. In the latter case, the average price level is too low and can drop markedly during the booking horizon. In MOD applications, this observation is also known as the "wild goose chase" (Castillo et al., 2024). In accordance with our findings, which show the potential severity of the underestimation errors, some works find that the resulting greedy behavior can lead to a worse performance than static pricing (Abdollahi et al., 2023, Yang and Strauss, 2017).

Due to the absence of research on DPC policies, there are no direct indications for respective underestimation errors in the literature. However, there is one indirect observation, which can now be better explained in the light of our findings: Koch and Klein (2020) observe that average OC increases quite strongly in the end of the booking horizon. This can be attributed to only unpopular fulfillment options remaining available, for which only few orders are collected already. Thus, collecting an additional order for such an option incurs high MCTS. If MCTS is not adequately captured by the OC approximation, this results in severe underestimation errors.

Another important observation from Fig. 6 is that neglecting either MCTS or DPC tends to cause a higher performance loss in *load*-constrained settings than in *dist*-constrained settings. This points toward a link of the two components by the consumption of logistical capacity associated with an order. In settings in which a high capacity consumption causes both strong displacement effects and a high marginal increase of fulfillment cost, DPC and MCTS will both be high, and vice versa. This means that an approximation based on one of the components can still correctly determine the relative ordering of the OC across different requests or different fulfillment options for one request. However, if the absolute value of the OC approximations is too low due to the high relevance of the missing component, this information gain does not come into effect. In Abdollahi et al. (2023), we find further evidence for this mechanism. They report that their MCTS policy performs well in terms of the number of collected orders and the cost per order, which indicates that the policy correctly captures the "right" orders to achieve high consolidation. However, their policy still fails to improve profit due to the general price level being too low.

Regarding the design of solution algorithms, the explicit consideration of both components is the safest way of avoiding structural underestimation. However, this may be associated with a much higher computational effort. Instead, the described findings related to underestimation suggest two low-threshold algorithmic strategies for mitigating its performance loss:

- 1. The OC approximation can be raised to the correct level by adding a rough estimate of the other component. An example for this can be found in Yang et al. (2016). They introduce a cost penalty as an additional rough DPC estimate if an order cannot be feasibly inserted into the sampled route plan (indicating strong displacement effects). The main challenge with such an approach is the correct adjustment of the rough DPC estimate. Klein et al. (2018) benchmark their own approach, which features a more sophisticated DPC approximation, against the approach by Yang et al. (2016), and the results show that its performance does not decline monotonically with a tighter capacity restriction. Based on our results, the reason for this is that the performance depends on how well the penalty is adjusted to result in the correct level of DPC rather than on the general magnitude of displacement effects.
- 2. The demand control approach can be made more robust against the consequences of underestimation rather than tackling the underestimation error itself. Hence, the goal is to reduce the regret resulting from a given underestimation error. Ulmer (2020) present an example for such an approach. They introduce a basis price that is charged even if the OC approximation would suggest an even lower price. This curtails greedy control behavior resulting from underestimation.

**Neglecting DPC or MCTS can cause severe overestimation errors**: In contrast to underestimation, which is presumed to cause performance differences in several existing works, there is only one publication discussing possible overestimation errors. Mackert (2019) observes that using a less accurate approximation of routing cost leads to the policy making less offers on average, which points toward an increase of average OC, i.e., toward a potential overestimation.

With our computational study, we provide definitive proof that overestimation errors as a consequence of neglecting DPC or MCTS exist. In most settings, overestimation contributes less than underestimation to the optimality gap. However, if overestimation is the dominant error, the optimality gap is particularly large in most cases. This is especially true for the MCTS policy, which can completely stall due to the recursive nature of the type of overestimation error it exhibits. At the same time, this recursiveness is what makes the error manageable by providing some form of anticipatory information on future consolidation opportunities. This explains the success of skeletal or sampled route planning that is applied as a well-performing algorithmic component in several existing publications (e.g., Anzenhofer et al., 2024, Koch and Klein, 2020, and Yang et al., 2016). However, none of these publications contains a conclusive explanation of why it is beneficial.

The early phase of the booking process is critical: Computational experiments in the existing literature suggest that demand control in the early phase of the booking process has a critical influence on solution quality (Anzenhofer et al., 2024, Campbell and Savelsbergh, 2006). Our computational results support this finding as we observe that the regret of a suboptimal demand control decision decreases over time for most types of overestimation and underestimation errors. Since the effect is observable independent of the revenue distribution, we can derive that the underlying main reason is not the loss in immediate reward but the transition to a state that has a much lower value than the successor state resulting from the optimal decision.

Anticipation is not an end in itself: The existing literature studying specific i-DMVRPs consistently emphasizes the importance of anticipation for the performance of solution approaches (Fleckenstein et al., 2023). However, our computational study shows that myopic policies can be hard to beat in certain settings (see Fig. 8 in Appendix B, e.g., *med*-profitability or *low*-profitability, *load*-constrained settings). This result becomes even more relevant given that, in the existing literature, anticipatory policies only achieve maximal improvements of a around 10 - 15%over myopic benchmarks for some specific i-DMVRPs (e.g., Azi et al., 2012, Heitmann et al., 2023, Koch and Klein, 2020), and that these numbers may even be subject to a negative results bias (Fanelli, 2012). From our computational study, we can derive some of the underlying reasons:

- 1. In some settings, an i-DMVRP can be "easy" to solve in general, e.g., *high*-profitability, *h-b-l* settings (see Fig. 8). Highly profitable demand arriving early in the booking process hardly impacts estimation errors but substantially reduces the resulting regret. Therefore, a more accurate OC approximation by an anticipatory policy may not necessarily translate into a reduced regret.
- 2. Myopic policies benefit from the fact that estimation errors can offset. Therefore, higher overestimation errors compared to an anticipatory policy can, counter-intuitively, be beneficial because the underestimation errors are then compensated better (e.g., see Fig. 14).
- 3. In settings with *low* profitability or low consolidation, overestimation, which myopic policies are prone to, is generally less harmful than underestimation because it leads to a more conservative acceptance behavior.

In summary, we conclude that investing a substantial amount of the scarce computation time at each decision epoch into an anticipatory OC approximation may not be the most efficient approach for certain i-DMVRPs. Instead, a greater performance improvement might be possible by

investing it into solving the other complex subproblems of i-DMVRPs, e.g., the demand control subproblem, or the dynamic vehicle routing subproblem.

# 7 Conclusion and Future Research Opportunities

With this work, we present a novel post-hoc explainability technique for i-DMVRPs as well as the results of its first systematic application. The proposed explainability technique allows quantifying and visualizing the extent of OC approximation errors, the regret associated with suboptimal decisions, and the rate at which such decisions are made. Applying this technique to a generic i-DMVRP and considering many settings with a broad variety of parameter values and their combinations, we derive a comprehensive identification and analysis of fundamental types of approximation errors. Finally, we show that this knowledge can be used to explain performance differences observed in the literature and to guide algorithm selection and algorithm design. In the following, we conclude by discussing how developers and users can benefit from our findings:

- Our computational results provide a guideline for how important an accurate consideration of DPC and MCTS is in certain problem settings, and how likely the different fundamental types of errors are to occur. This provides indications for the selection of a basic solution concept.
- 2. The technique presented in Section 4 can be directly applied to small instances of the specific i-DMVRP under consideration to compare variants of the pre-selected solution concept with the optimal policy. Thereby, developers and users can gain detailed insights into which types of estimation errors are most relevant for their problem and make targeted adaptions to their algorithms to tackle them.
- 3. Providing training for users based on our findings can enable them to better diagnose anomalies in the behavior of policies and react accordingly.

Since algorithmic explainability is still an emerging concept that is of high relevance both from a theoretical and a practical perspective, we believe that there is great potential for future research connected to our work:

- 1. It would be interesting to apply post-hoc explainability techniques similar to ours to related problem classes, such as network revenue management problems arising, e.g., in the airline industry (Klein et al., 2020). In particular, also given the promising results of related research (Bravo and Shaposhnik, 2020), we believe that optimal policies for small instances are a valuable source of domain knowledge for problems that cannot be solved to optimality when considering realistic-sized instances.
- 2. Complementary to post-hoc techniques based on OC, we see a similar potential for the development of inherently explainable solution approaches for i-DMVRPs, such as inherently explainable policy function approximations.
- 3. Finally, our numerical results suggest that myopic policies deserve more attention. While they are currently mainly considered for pure benchmarking purposes in comparison to

anticipatory approaches, their potential as carefully designed, full-fledged solution approaches should also be investigated more thoroughly.
### References

- Abdollahi, M., Yang, X., Nasri, M. I., & Fairbank, M. (2023). Demand management in timeslotted last-mile delivery via dynamic routing with forecast orders. *European Journal of Operational Research*, 309(2), 704-718. https://doi.org/10.1016/j.ejor.2023.01.023.
- Amir, O., Doshi-Velez, F., & Sarne, D. (2019). Summarizing agent strategies. Autonomous Agents and Multi-Agent Systems, 33, 628-644. https://doi.org/10.1007/s10458-019-09418-w.
- Amorim, P., DeHoratius, N., Eng-Larsson, F., & Martins, S. (2024). Customer preferences for delivery service attributes in attended home delivery. *Management Science*, 70(11), 7559-7578. https://doi.org/10.1287/mnsc.2020.01274.
- Anzenhofer, F., Fleckenstein, D., Klein, R., & Steinhardt, C. (2024). Analyzing the impact of demand management in rural shared mobility-on-demand systems. Working paper, University of Augsburg. http://dx.doi.org/10.2139/ssrn.4682056.
- 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.
- Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82-115. https://doi.org/10.1016/j.inffus.2019.12.012.
- Asdemir, K., Jacob, V. S., & Krishnan, R. (2009). Dynamic pricing of multiple home delivery options. *European Journal of Operational Research*, 196(1), 246-257. https://doi.org/10.1016/j.ejor.2008.03.005.
- Atasoy, B., Ikeda, T., Song, X., & Ben-Akiva, M. E. (2015). The concept and impact analysis of a flexible mobility on demand system. *Transportation Research Part C: Emerging Technol*ogies, 56, 373-392. https://doi.org/10.1016/j.trc.2015.04.009.
- Avraham, E., & Raviv, T. (2021). The steady-state mobile personnel booking problem. Transportation Research Part B: Methodological, 154, 266-288. https://doi.org/10.1016/j.trb.2021.10.008.
- Azi, N., Gendreau, M., & Potvin, J.-Y. (2012). A dynamic vehicle routing problem with multiple delivery routes. *Annals of Operations Research*, 199(1), 103-112. https://doi.org/10.1007/s10479-011-0991-3.
- Biggs, M., Sun, W., & Ettl, M. (2021). Model distillation for revenue optimization: Interpretable personalized pricing. In *International Conference on Machine Learning*, 946-956.
- Bravo, F., & Shaposhnik, Y. (2020). Mining optimal policies: A pattern recognition approach to model analysis. *INFORMS Journal on Optimization*, 2(3), 145-166. https://doi.org/10.1287/ijoo.2019.0026.
- Campbell, A. M., & Savelsbergh, M. W. P. (2005). Decision support for consumer direct grocery initiatives. *Transportation Science*, *39*(3), 313-327. https://doi.org/10.1287/trsc.1040.0105.
- Campbell, A. M., & Savelsbergh, M. W. P. (2006). Incentive schemes for attended home delivery services. *Transportation Science*, *40*(3), 327-341. https://doi.org/10.1287/trsc.1050.0136.
- Castillo, J. C., Knoepfle, D. T., & Weyl, E. G. (2024). Matching in ride hailing: Wild goose chases and how to solve them. Working paper, University of Pennsylvania.
- De Bock, K. W., Coussement, K., De Caigny, A., Słowiński, R., Baesens, B., Boute, R. N., et al. (2024). Explainable AI for operational research: A defining framework, methods, applications, and a research agenda. *European Journal of Operational Research*, 317(2), 249-272.
- Fanelli, D. (2012). Negative results are disappearing from most disciplines and countries. *Scientometrics*, 90(3), 891-904. https://doi.org/10.1007/s11192-011-0494-7.
- Fleckenstein, D., Klein, R., Klein, V., & Steinhardt, C. (2024). On the concept of opportunity cost in integrated demand management and vehicle routing. *Transportation Science, online first*. https://doi.org/10.1287/trsc.2024.0644.

- Fleckenstein, D., Klein, R., & Steinhardt, C. (2023). Recent advances in integrating demand management and vehicle routing: A methodological review. *European Journal of Operational Research*, 306(2), 499-518. https://doi.org/10.1016/j.ejor.2022.04.032.
- Gallego, G., & Topaloglu, H. (2019). Revenue management and pricing analytics. New York, NY: Springer. https://doi.org/10.1007/978-1-4939-9606-3.
- Goerigk, M., & Hartisch, M. (2023). A framework for inherently interpretable optimization models. *European Journal of Operational Research*, 310(3), 1312-1324. https://doi.org/10.1016/j.ejor.2023.04.013.
- Heitmann, R. J. O., Soeffker, N., Ulmer, M. W., & Mattfeld, D. C. (2023). Combining value function approximation and multiple scenario approach for the effective management of ridehailing services. *EURO Journal on Transportation and Logistics*, 12, 100104. https://doi.org/10.1016/j.ejtl.2023.100104.
- Hildebrandt, F. D., Thomas, B. W., & Ulmer, M. W. (2023). Opportunities for reinforcement learning in stochastic dynamic vehicle routing. *Computers & Operations Research*, 150, 106071. https://doi.org/10.1016/j.cor.2022.106071.
- Juozapaitis, Z., Koul, A., Fern, A., Erwig, M., & Doshi-Velez, F. (2019). Explainable reinforcement learning via reward decomposition. In *IJCAI/ECAI Workshop on Explainable Artificial Intelligence*.
- Klapp, M. A., Erera, A. L., & Toriello, A. (2020). Request acceptance in same-day delivery. *Transportation Research Part E: Logistics and Transportation Review*, 143, 102083. https://doi.org/10.1016/j.tre.2020.102083.
- Klein, R., Koch, S., Steinhardt, C., & Strauss, A. K. (2020). A review of revenue management: Recent generalizations and advances in industry applications. *European Journal of Operational Research*, 284(2), 397-412. https://doi.org/10.1016/j.ejor.2019.06.034.
- Klein, R., Mackert, J., Neugebauer, M., & Steinhardt, C. (2018). A model-based approximation of opportunity cost for dynamic pricing in attended home delivery. *OR Spectrum*, 40(4), 969-996. https://doi.org/10.1007/s00291-017-0501-3.
- Klein, V., & Steinhardt, C. (2023). Dynamic demand management and online tour planning for same-day delivery. *European Journal of Operational Research*, 307(2), 860-886. https://doi.org/10.1016/j.ejor.2022.09.011.
- Koch, S., & Klein, R. (2020). Route-based approximate dynamic programming for dynamic pricing in attended home delivery. *European Journal of Operational Research*, 287(2), 633-652. https://doi.org/10.1016/j.ejor.2020.04.002.
- Lang, M. A. K., Cleophas, C., & Ehmke, J. F. (2021). Anticipative dynamic slotting for attended home deliveries. *Operations Research Forum*, 2(4), 70. https://doi.org/10.1007/s43069-021-00086-9.
- Lebedev, D., Goulart, P., & Margellos, K. (2021). A dynamic programming framework for optimal delivery time slot pricing. *European Journal of Operational Research*, 292(2), 456-468. https://doi.org/10.1016/j.ejor.2020.11.010.
- Li, Y., Archetti, C., & Ljubic, I. (2024). Emerging optimization problems for distribution in sameday delivery. Working paper, ESSEC Business School. https://doi.org/10.48550/arXiv.2405.05620.
- Mackert, J. (2019). Choice-based dynamic time slot management in attended home delivery. *Computers* & *Industrial Engineering*, *129*, 333-345. https://doi.org/10.1016/j.cie.2019.01.048.
- Mannor, S., Simester, D., Sun, P., & Tsitsiklis, J. N. (2007). Bias and variance approximation in value function estimates. *Management Science*, 53(2), 308-322. https://doi.org/10.1287/mnsc.1060.0614.
- Milani, S., Topin, N., Veloso, M., & Fang, F. (2024). Explainable reinforcement learning: A survey and comparative review. ACM Computing Surveys, 56(7), 1-36. https://doi.org/10.1145/3616864.

- Powell, W. B. (2011). Approximate dynamic programming: Solving the curses of dimensionality. 2nd ed. Hoboken, NJ: John Wiley & Sons. https://doi.org/10.1002/9781118029176.
- Powell, W. B. (2022). Reinforcement Learning and Stochastic Optimization: A unified framework for sequential decisions. Hoboken, NJ: John Wiley & Sons. https://doi.org/10.1002/9781119815068.
- Psaraftis, H. N., Wen, M., & Kontovas, C. A. (2016). Dynamic vehicle routing problems: Three decades and counting. *Networks*, 67(1), 3-31. https://doi.org/10.1002/net.21628.
- Puterman, M. L. (2014). Markov decision processes: Discrete stochastic dynamic programming. New York, NY: John Wiley & Sons. https://doi.org/10.1002/9780470316887.
- Russell, S. J., & Zimdars, A. (2003). Q-decomposition for reinforcement learning agents. In *Proceedings of the 20th international conference on machine learning (ICML-03)*, 656-663.
- Snoeck, A., Merchán, D., & Winkenbach, M. (2020). Revenue management in last-mile delivery: state-of-the-art and future research directions. *Transportation Research Procedia*, 46, 109-116. https://doi.org/10.1016/j.trpro.2020.03.170.
- Soeffker, N., Ulmer, M. W., & Mattfeld, D. C. (2022). Stochastic dynamic vehicle routing in the light of prescriptive analytics: A review. *European Journal of Operational Research*, 298(3), 801-820. https://doi.org/10.1016/j.ejor.2021.07.014.
- Strauss, A. K., Klein, R., & Steinhardt, C. (2018). A review of choice-based revenue management: Theory and methods. *European Journal of Operational Research*, 271(2), 375-387. https://doi.org/10.1016/j.ejor.2018.01.011.
- Talluri, K. T., & Van Ryzin, G. J. (2004). The theory and practice of revenue management. New York, NY: Springer. https://doi.org/10.1007/b139000.
- Torrey, L., & Taylor, M. (2013). Teaching on a budget: Agents advising agents in reinforcement learning. In Proceedings of the 2013 International Conference on Autonomous Agents and Multi-Agent Systems, 1053-1060.
- Toth, P., & Vigo, D. (Eds.). (2014). Vehicle routing: Problems, methods, and applications. Philadelphia, PA: Society for Industrial and Applied Mathematics. https://doi.org/10.1137/1.9781611973594.
- Ulmer, M. W. (2020). Dynamic pricing and routing for same-day delivery. *Transportation Science*, 54(4), 1016-1033. https://doi.org/10.1287/trsc.2019.0958.
- Ulmer, M. W., & Thomas, B. W. (2020). Meso-parametric value function approximation for dynamic customer acceptances in delivery routing. *European Journal of Operational Research*, 285(1), 183-195. https://doi.org/10.1016/j.ejor.2019.04.029.
- Van Seijen, H., Fatemi, M., Romoff, J., Laroche, R., Barnes, T., & Tsang, J. (2017). Hybrid reward architecture for reinforcement learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 5398-5408.
- Vansteenwegen, P., Melis, L., Aktaş, D., Montenegro, B. D. G., Vieira, F. S., & Sörensen, K. (2022). A survey on demand-responsive public bus systems. *Transportation Research Part C: Emerging Technologies*, 137, 103573. https://doi.org/10.1016/j.trc.2022.103573.
- Waßmuth, K., Köhler, C., Agatz, N., & Fleischmann, M. (2023). Demand management for attended home delivery – A literature review. *European Journal of Operational Research*, 311(3), 801-815. https://doi.org/10.1016/j.ejor.2023.01.056.
- Yang, X., & Strauss, A. K. (2017). An approximate dynamic programming approach to attended home delivery management. *European Journal of Operational Research*, 263(3), 935-945. https://doi.org/10.1016/j.ejor.2017.06.034.
- Yang, X., Strauss, A. K., Currie, C. S. M., & Eglese, R. (2016). Choice-based demand management and vehicle routing in e-fulfillment. *Transportation Science*, 50(2), 473-488. https://doi.org/10.1287/trsc.2014.0549.

## **Appendix A: Settings**



Fig. 7 11 customer settings which are solved six times each, for all combinations of routing cost factors 0.2, 0.6, 1 with either binding load constraints or binding distance constraints



### **Appendix B: Numerical Results**

(b) Dist-constrained settings

Fig. 8 Objective values resulting from the different opportunity cost approximations – Averaged across 50 instances per setting

Rearet underest

0 0.3 0.6 1 Capacity consumption

7.0

6.0

5.0

4.0

3.0

2.0

1.0

0.0

7.0

6.0

5.0

4.0

3.0

2.0

1.0

0.0

## **Appendix C: Heatmaps**

OC underest



17.5

15.0

12.5

10.0

7.5

5.0

- 2.5 0.0

20.0

15.0

10.0

0.0

ò



0 0.3 0.6 1 Capacity consumption

(b) med | load | unif | l-b-h

Rearet underest







(c) med | dist | clust | l-b-h

0.2 0.4 0.6 0.8 Capacity consumption

0.9







Fig. 9 DPC policy – underestimation errors

0 0.3 0.6 1 Capacity consumption

t<sup>4</sup>

OC underest

0.2 0.4 0.6 0.8 0.9 Capacity consumption

OC underest.

t 4 5

6

9 -

ò



Article A3: From Approximation Error to Optimality Gap – Explaining the Performance Impact of Opportunity Cost Approximation in Integrated Demand Management and Vehicle Routing

Fig. 10 DPC policy – overestimation error

(d) low | load | clust\_sort | l-b-h



0.4

0.3

0.2

0.1

0.0

0.4

0.3

0.2

0.1

0.0

0.4

0.3

0.2

0.1

0.0

0.4

0.3

0.2

0.1

0.0

0 0.3 0.6 1 Capacity consumption

Fig. 11 MCTS policy - underestimation errors

0 0.3 0.6 1 Capacity consumption 0.3 0.6 1 icity consumption

(d) med | load | unif | homog

0 Capa





(d) low | dist | clust\_sort | rand

Fig. 12 MCTS policy - overestimation error





(c) low | load | unif | homog

Fig. 13 MCTS policy - overestimation error (cont'd)





Fig. 14 Offsetting of underestimation errors and overestimation errors in setting low | load | unif | 1-b-h

# Article A4: Extended Booking Horizons in Rural Shared Mobility-on-Demand Systems: Insights and Implications for Demand Management

Fabian Anzenhofer<sup>a</sup>, David Fleckenstein<sup>a</sup> (<sup>a</sup>University of Augsburg) Working paper

#### Abstract

Shared mobility-on-demand systems receive increasing attention as a sustainable transportation option in rural areas. To guarantee maximum planning reliability given the limited alternative modes of public transport, rural providers typically enable customers to request rides several days in advance of their desired day and time. Existing literature shows that extending the booking horizon beyond the service day influences system performance via the degree of dynamism. However, there is no analysis of the special demand structure resulting from allowing advance requests within an extended booking horizon. In this work, we present a methodology for descriptively analyzing real-world data to identify unique demand patterns with the aim of addressing them by suitable demand management approaches. Through a case study using a large dataset provided by our industry partner FLEXIBUS, we uncover four patterns: First, customers may strategically reserve rides early in the booking horizon, which can be addressed by anticipatory demand management. Second, customers predominantly request outward rides in advance, while corresponding return rides are often requested ad-hoc. This observation motivates a flexible product design for round trips. Third, extended booking horizons can exacerbate the issue of cancellations. We demonstrate that a state-of-the-art supervized learning model accurately predicts cancellation probabilities, which can be used for advanced demand management. Fourth, we observe that customers exhibit time flexibility in both directions, i.e., regularly accept to ride earlier or later than originally desired. Hence, the provider can use demand management to steer customers toward alternative pick-up/drop-off times that are more favorable.

Key words: Shared Mobility-on-Demand, Rural Transportation, Extended Booking Horizon, Demand Management

## 1 Introduction

Urbanization is rapidly transforming both urban areas and rural areas. By 2050, 68% of the global population is expected to reside in cities (Sun et al., 2020). Meanwhile, rural areas face declining populations, which reduces mobility demand (Mounce et al., 2020). Traditional scheduled public transport becomes less efficient, leading to a vicious cycle of reduced usage and service quality (Bar-Yosef et al., 2013) and a lack of competitiveness with private motorized transport (Eurostat, 2022). Hence, to contribute to a more sustainable rural transport system, public transport must improve its service quality in terms of flexibility and planning reliability.

Shared mobility-on-demand (SMOD) systems are widely regarded as a promising solution for this challenge. They are demand-oriented, offering the customers more flexibility compared to rigid scheduled services. More precisely, customers can request rides between specific pick-up points and drop-off points at a desired time. At the same time, like scheduled services, SMOD systems aim at consolidating demand. The term "shared" refers to ridepooling, whereby passengers of multiple independent requests can be pooled in a vehicle simultaneously.

Despite these favorable characteristics, SMOD systems have a fundamental weakness: the lack of planning reliability. Unlike scheduled services, which guarantee transport, SMOD systems cannot meet all demand without a substantially oversized fleet and reduced consolidation. As a result, request rejections due to vehicle unavailability are inevitable. Most research (e.g., Arian et al., 2022, Hungerländer et al., 2021) focuses on SMOD systems with short booking horizons, which only allow customers to place *ad-hoc requests* with practically no booking lead time ( $\leq$  5 min) before their desired pick-up time, making it difficult for customers to rely solely on SMOD for their future mobility needs. This is especially problematic in rural areas with limited transport alternatives available on short notice.

To improve planning reliability, providers can extend the booking horizon, allowing customers to request rides hours or even days in advance. This includes *same-day requests* with more than a 5-minute booking lead time and *advance requests* placed even before the service day. Both options eliminate the uncertainty of whether the SMOD service will be available to serve a customer's future travel needs. In case it is not, customers have more time to adjust and find an alternative mode of transport or replan their activities.

Notwithstanding its advantages and practical relevance, academic literature lacks research on how an extended booking horizon impacts the demand structure of an SMOD service and what the implications are for the provider's planning process, in particular for demand management. In general, demand management involves strategic, tactical, and operational decision-making regarding the sales process of a logistical service (Talluri and van Ryzin, 2004, p. 3, Waßmuth et al., 2023). In our work, we consider, on the one hand, strategic demand management regarding (virtual) product definitions made by the SMOD provider, such as regular pooled rides and premium express rides without pooling. On the other hand, we discuss operational demand management: Here we distinguish between availability control and dynamic pricing. Availability control refers to the operational decision on which (alternative) rides to offer in response to a request. Availability control in the context of SMOD can use different mechanisms, such as rejecting a request (not offering a ride at all) or using time shifts (offering alternative times to the originally desired time). Dynamic pricing refers to decisions on the the pricing of offered rides depending on the current system state. Both availability control and dynamic pricing can differ in their use of information, resulting in myopic or anticipatory (using information on future demand) decision-making.

This work uses real-world data from FLEXIBUS, one of the most experienced SMOD providers in Germany, to analyze the demand structure of an SMOD provider offering an extended booking horizon. From the results, we derive implications for demand management. Please note that, for reasons of data confidentiality, we deliberately do not disclose the spatial and temporal extent of the data set, any absolute request numbers, and the actual demand management approach used.

To the best of our knowledge, we are the first to extensively investigate the following research questions:

- Which demand patterns emerge when customers can place requests within an extended booking horizon?
- How can providers use strategic and operational demand management to address potential issues arising from these patterns and improve system performance?

In summary, the scientific contribution of our work lies in the identification of four key demand patterns. For each pattern, we provide evidence from a descriptive analysis and propose demand management approaches leveraging this knowledge:

- **Reservation behavior**: Customers show strategic behavior by reserving rides early in the booking horizon that otherwise are likely to be rejected by SMOD providers. To mitigate these reservations, which can lead to inefficient route plans, providers can apply anticipatory demand management at the operational level.
- Round trips: Customers also show strategic behavior by requesting the outward ride in advance, while the return ride is requested with a shorter booking lead time (same-day or adhoc requests). This may reflect the fact that many customers have an appointment with an uncertain end time, preventing them from requesting the return ride in advance. As a result, customers face uncertainty about whether the SMOD service will have enough capacity for their return ride. SMOD providers can address this at the strategic planning level by designing a product that allows customers to request the return ride in advance without specifying the exact pick-up time but instead guarantees service within a fixed time interval. This may increase planning reliability for customers.

- **Cancellations**: We find that offering an extended booking horizon may exacerbate cancellation issues. At the observed scale, cancellations are likely to substantially harm the system performance. Allowing customers to place advance requests may lead to long dwell times of false demand (soon to be cancelled) in the system as true demand, which can be detrimental to operational planning. A mitigation strategy is to predict cancellation probabilities at the time of request arrival. We show that, using state-of-the-art supervized machine learning algorithms, our predictions are accurate enough to support operational demand management.
- **Time flexibility**: We analyze customers' observed time flexibility by comparing the actual booked times with their originally desired times. Providers can exploit this flexibility by applying time shifts, which are an operational demand management mechanism. Thereby, customers are steered toward choosing an alternative time that is more favorable for the provider (Anzenhofer et al., 2024a).

The remainder of this work is structured as follows: Section 2 reviews the literature on descriptive analyses of SMOD systems with a focus on extended booking horizons and their implications for the provider's planning process. Section 3 provides a descriptive analysis on spatial and temporal characteristics with a focus on advance requests. In Section 4, we discuss the operational impact of the four identified demand patterns resulting from extended booking horizons and propose demand management approaches. Finally, Section 5 concludes the study and discusses future research directions.

### 2 Literature Review

In this section, we review the existing academic literature. Section 2.1 focuses on empirical studies examining (S)MOD systems and their analysis of demand structure. Section 2.2 explores research on SMOD systems that allow customers to place same-day and advance requests, emphasizing the impact of an extended booking horizon on operational planning (vehicle routing and demand management). Finally, Section 2.3 identifies the research gap in the analysis of demand patterns in SMOD systems arising from an extended booking horizon.

### 2.1 Descriptive Analysis of SMOD Customer Demand

This section reviews works that conduct descriptive analyses of customer demand patterns in (S)MOD systems. We subdivide the discussion into two parts, starting with works that consider urban areas, before turning toward research focusing on rural areas.

**Urban areas**: Several studies examine demand for SMOD in urban areas. Zwick and Axhausen (2022) analyze MOIA's SMOD systems in Hamburg and Hannover (Germany) and focus on the spatio-temporal demand structure. Using heatmaps and histograms, they visualize the spatial distribution of rides as well as demand fluctuations over time. Other works, such as Zwick et al. (2023), Gödde et al. (2023), and Weckström et al. (2018), provide similar spatiotemporal analyses. However, these studies differ from ours as they focus on SMOD systems that do not allow

advance requests. Other studies also examine spatial and temporal demand patterns in urban areas, but focus on ride-hailing systems, which only handle ad-hoc requests (e.g., Li et al., 2023, Shulika et al., 2024).

Only one work specifically focuses on the demand structure of an SMOD system that allows advance requests. Chandakas (2020) analyze the urban SMOD provider Tisséo Mobibus in Toulouse, developing a medium-term forecast of demand up to seven days before the service day. They find that the level of "early" advance requests is a good indicator of the final total demand. Tisséo Mobibus allows customers to place advance requests 28 days before the desired service day and up to two hours before the desired time. Their descriptive analysis includes demand levels by month, week, and day, as well as booking curves showing when requests are placed within the booking horizon. Although the authors also descriptively examine an SMOD provider with an extended booking horizon, our work differs significantly from theirs: First, unlike them, we emphasize the interrelation between spatial and temporal patterns, particularly regarding the time of request. Second, while their work focuses on demand prediction, we elaborate on how the provider can draw on the broad "toolbox" of demand management to account for the discovered demand patterns. Third, we examine rural SMOD systems with a fundamentally different demand structure than in urban areas. Fourth, Tisséo Mobibus does not offer ad-hoc requests or same-day requests within two hours before the desired time.

**Rural areas**: Research on rural SMOD systems aligns more closely with our work. For example, Sörensen et al. (2021) examine the SMOD system EcoBus in the Oberharz region (Germany). They use flow maps, i.e., plots that illustrate the spatial distribution of desired origin and destination locations via arrows, to provide insights into average demand levels. They also explore the interrelation between spatial demand patterns and desired times by presenting flow maps at different time intervals throughout the service day. While our descriptive analysis also includes analyses of flow maps and desired times (see Section 3.3), our main contribution lies in analyzing the interrelation between the time of request and spatiotemporal characteristics to identify patterns induced by the extended booking horizon. There are also other works, like Imhof and Blättler (2023), that examine a rural SMOD system (mybuxi in the canton of Berne, Switzerland), but differ greatly from our work as they focus on SMOD systems that do not allow advance requests.

## 2.2 Implications of an Extended Booking Horizon for Operational and Strategic Planning

A body of literature examines the operational and strategic planning of SMOD systems that feature an extended booking horizon. The vast majority of works analyzes the implications of an extended booking horizon for operational planning, without explicitly accounting for specific demand patterns. The review consists of three parts: First, we consider implications for vehicle routing decisions, which involve assigning customer requests to vehicles and planning routes. Second, we consider operational demand management, which refers to decisions regarding ride availability or pricing. Third, we discuss strategic demand management.

**Vehicle routing**: Numerous studies address the implications of an extended booking horizon for vehicle routing. Elting and Ehmke (2021) investigate the rural SMOD system EcoBus in the Oberharz region, using constraint programming to search for feasible solutions of the dynamic dial-a-ride problem. They argue that requests with longer booking lead times disrupt the natural ordering of requests, potentially even reducing the number of served requests. However, this effect can be mitigated by more advanced routing heuristics or rolling horizon optimization. Other studies on SMOD systems with extended booking horizons, both in urban (e.g., Engelhardt et al., 2022, Ma and Koutsopoulos, 2022, Theodoridis et al., 2023) and rural areas (e.g., Lu et al., 2023), show that even requests with relatively short booking lead times (e.g., same-day requests) can improve the performance in terms of the number of served requests and vehicle kilometers traveled.

There are also a few studies that investigate the impact of cancellations in SMOD systems on vehicle routing: Horn (2002) incorporates cancellations into operational planning models without proposing mitigation strategies. Wu et al. (2024) investigate prediction-failure-risk-aware online dial-a-ride scheduling that allows for request selection and cancellations. Unlike prior research, their correction mechanisms for prediction errors and cancellations are embedded into the optimization framework to proactively mitigate negative effects. However, they do not use demand management to proactively reject requests that are very likely to be cancelled. We propose that predicting cancellation probabilities and selectively rejecting such requests through demand management could be a major lever for SMOD providers. Cancellations are more often considered in the context of other types of MOD services, such as ridehailing services (e.g., He et al., 2018, Wang et al., 2019, Wang et al., 2024, Xu et al., 2022, Sun et al., 2023) or nonemergency medical transportation (e.g., Yu et al., 2021).

**Operational demand management**: Research on the demand management implications of extended booking horizons is generally scarce. In the following, we discuss the literature that applies operational demand management to SMOD systems in urban and rural areas, highlighting the few papers that consider an extended booking horizon.

*Dynamic pricing*: There is literature on dynamic pricing for SMOD systems, especially in urban areas (e.g., Qiu et al., 2018, Sharif Azadeh et al., 2022), but only few works consider dynamic pricing in rural areas (e.g., Arian et al., 2022, Anzenhofer et al., 2024b). More specifically, Anzenhofer et al. (2024b) is the only work that applies dynamic pricing to a rural SMOD system with an extended booking horizon. In addition, there is literature on ride-hailing systems with dynamic pricing (Al-Kanj et al., 2020, Chen et al., 2019, Haliem et al., 2021, Ni et al., 2021, Liu et al., 2021). However, none of them consider an extended booking horizon.

*Availability control*: The literature on availability control in SMOD systems is limited: While there are few works on SMOD systems in urban areas (Atasoy et al., 2015, Huang et al., 2020,

Haferkamp et al., 2024), only Anzenhofer et al. (2024a) consider a rural SMOD system with an extended booking horizon, and find that customers' time flexibility (the deviation of the booked time from the originally desired time) can be exploited to steer customers toward more favorable pick-up times. In the study at hand, the descriptive analysis of the observed time flexibility is much more detailed.

In addition, there is also a body of literature that examines demand management via simple accept-reject mechanisms in urban areas (Haferkamp and Ehmke, 2022, Heitmann et al., 2023, Hosni et al., 2014, Jung et al., 2016, Lotfi and Abdelghany, 2022). However, these works exclusively consider ad-hoc requests.

**Strategic demand management**: Research related to strategic demand management for SMOD systems is scarce. There is a small body of literature that focuses primarily on operational demand management but considers SMOD systems with differentiated products. For example, Arian et al. (2022) and Sharif Azadeh et al. (2022) study SMOD systems that offer several products (rides) with varying service quality, particularly with regard to waiting times. In addition, Atasoy et al. (2015) investigate a SMOD system that allows customers to choose between shared and non-shared rides.

### 2.3 Research Gap

We are the first to conduct a detailed descriptive analysis of an SMOD system with an extended booking horizon, focusing specifically on the time of request. On the one hand, an extended booking horizon allows SMOD providers to offer a higher level of planning reliability to potential customers. On the other hand, it results in unique demand patterns that may present challenges and opportunities for both SMOD providers and customers. Our study provides guidance to SMOD providers on how to leverage these patterns through tailored demand management approaches. We address a research gap in two key dimensions:

**Descriptive analysis focusing on advance requests**: While spatial and temporal demand distributions in SMOD systems have been extensively studied, we are the first to provide detailed insights into how an extended booking horizon impacts the demand structure and induces demand patterns in rural areas.

### Tailored demand management approaches addressing demand patterns:

- Reservations: So far, there is no literature on rural SMOD areas providing evidence of strategic reservation behavior by customers. Its existence is only hypothesized by Sörensen et al. (2021). Nevertheless, there is already literature presenting demand management approaches suitable for addressing this demand pattern (e.g., Anzenhofer et al., 2024a). However, none of these publications explicitly refers to it.
- Round trips: There is neither specific descriptive analysis of round trips in SMOD systems nor research on demand management approaches for addressing this demand pattern.

- Cancellations: We are the first to analyze cancelled demand with a focus on discussing and visualizing the importance of the ghost demand dwell time, i.e., the time interval between time of request and the actual cancellation, in SMOD systems. While there is some literature on the implications of cancellations on vehicle routing, there is no literature on tailored demand management approaches that account for this demand pattern.
- Time flexibility: There is limited literature that examines SMOD data to observe time flexibility in customer requests. Only Anzenhofer et al. (2024a) provide some evidence of observed customer time flexibility in rural SMOD services, but not at the same level of detail as in this paper. Further, they propose a demand management approach that exploits the time flexibility by means of time shifts.

These gaps point to a clear need for further research into how extended booking horizons influence demand patterns and how providers can better manage demand in response.

### 3 Descriptive Analysis of the Customer Demand

This section provides a descriptive analysis of customer demand in one of the service areas of our industry partner. Section 3.1 provides a general summary of the booking process. In Section 3.2, we briefly introduce the service area under consideration. Then, Section 3.3 and Section 3.4 comprise a detailed analysis of customer demand with particular emphasis on the demand patterns associated with the extended booking horizon.

### 3.1 Booking Process

In this section, we outline the prototypical booking process employed by rural SMOD providers, along with the corresponding operational planning decisions. During the booking horizon, requests are placed dynamically. Each request follows a four-step procedure: request arrival, feasibility check (or demand management decision), order confirmation, and vehicle routing decision. Fig. 1 illustrates the process for an individual request, followed by a step-by-step explanation:

- **Request arrival**: Initially, the customer specifies their desired ride via a smartphone app (or telephone). This includes parameters regarding the desired time, the origin-destination pair (OD pair), and the number of passengers. The provider tracks the time of request placement (timestamp).
- Feasibility check (or demand management decision): Next, the provider assesses the feasibility of fulfilling the request based on the current route plan for all vehicles. Alternative ride times with small deviations from the desired time, either earlier or later, are also checked for feasibility. This approach, referred to as the *feasibility check*, generally follows a firstcome-first-served policy, commonly used by SMOD providers. A more advanced method additionally involves a demand management decision, assessing whether the request is profitable, e.g., considering revenue net of routing costs (see Anzenhofer et al., 2024a). The

outcome of this step is an offer set, a list of feasible (and profitable) rides with different desired times. If no feasible (and profitable) ride exists, the request is rejected by the provider.

- Order confirmation: In this step, the provider presents the offer set from the previous step. The customer chooses their preferred ride and confirms the order. If unsatisfied with the available rides, the customer can still choose to abandon the booking process.
- Vehicle routing decision: Finally, the provider updates the route plan to incorporate the new order. This step is also required for advance requests to ensure feasibility for subsequent requests. However, if the customer abandons the process, no vehicle routing decision is made.



Fig. 1 Booking process for an exemplary request (C: customer, P: provider)

As our primary goal is to analyze customer demand, we concentrate on the first step, request arrival, and provide an explanation of the key request parameters below:

**Time of request**: This parameter refers to the timestamp of a customer's request submission during the booking horizon. In FLEXIBUS' system, customers can place advance requests up to 14 days before the start of the service horizon. Generally, customers can submit requests at any time before their desired time (also during the service horizon).

**Desired time**: The desired time refers to the time in the service horizon when the customer wishes to be transported. We classify requests based on the time of request in relation to the desired time: Advance requests are submitted the day before the start of the service horizon or earlier. Same-day requests are submitted during the service horizon but with a booking lead time, i.e., the time span between time of request and desired time, of more than 5 min. Ad-hoc requests are placed with less than 5 min of lead time. The service horizon, which can vary by service area and day type (e.g., weekday or weekend), defines the available time range for the desired pick-up or drop-off.

**OD pair**: Customers specify an OD pair, consisting of two locations within the service area. The provider uses predefined virtual stops across the service area and assigns each specified location to the nearest virtual stop.

**Number of passengers**: Providers allow customers to request rides for multiple passengers in a single request.

#### 3.2 Service Area and Service Horizon

The service area under consideration features a central town, where the depot is located, surrounded by smaller villages, which we collectively refer to as the peripheral area.



Fig. 2 Level of demand

In our analysis, we focus on working days (Monday to Thursday), as these days share the same service horizon (5 a.m. to 9 p.m.) and have similar demand levels as illustrated in Fig. 2. Please note that we exclude Fridays, Saturdays, Sundays, and public holidays from our analysis, as they have a different service horizon and demand levels than working days.

### 3.3 Descriptive Analysis of Customer Demand

In this section, we analyze the customer demand based on the key request parameters: time of request, desired time, and desired OD pair. We first conduct a separate descriptive analysis of each parameter, followed by an assessment of their interdependencies with varying booking lead times to identify specific demand patterns in the next subsection.

Fig. 3 provides insights into time of request, desired time, and spatial distribution, forming a basis for understanding the structure of customer demand.

**Time of request**: Fig. 3a illustrates the booking curve, showing the distribution of the time of request (aggregated on a day-level). A notable peak occurs at -14 days, with 8.75% of requests made shortly after the start of the booking horizon. The curve then rises moderately, reaching another peak at -1 day, by which time 53.09% of the total requests have been submitted as advance requests (represented by green and light-yellow bars).

Article A4: Extended Booking Horizons in Rural Shared Mobility-on-Demand Systems: Insights and Implications for Demand Management



Fig. 3 Descriptive analysis of time of request, desired times, and origin-destination pairs

These advance requests can be further divided into early advance requests (with a booking lead time of -8 to -14 days, representing 20.25%), and late advance requests (with a booking lead time of -1 to -7 days, representing 32.84%). Overall, advance requests account for the largest share compared to same-day requests (22.86%, dark-yellow bar) and ad-hoc requests (24.05%, red bar). This highlights the importance of advance requests, which are more frequent than same-day requests and ad-hoc requests combined.

**Desired time**: Fig. 3b shows the distribution of hourly desired times for the service horizon, which ranges from 5:00 to 21:00. The demand levels reveal distinct patterns. The data shows an early peak between 9:00 and 12:00, capturing 33.73% of the total demand (light-yellow bars), followed by an afternoon peak from 13:00 to 16:00, which accounts for 30.53% (dark-yellow bars). Two off-peak periods are also identifiable: the morning hours from 5:00 to 8:00, comprising 19.79% of the total demand, and the evening hours from 17:00 to 21:00, making up 15.95%. These patterns highlight a concentration of demand during the midday hours, with significantly lower demand during the early morning and late evening periods.

**Spatial distribution**: Lastly, we examine the spatial demand distribution (Fig. 3c), which is closely tied to the service area. For greater clarity, each virtual stop and the associated demand is mapped to a cluster (given by the provider). The demand is visualized using a heatmap, in which each circle corresponds to one of the clusters. Its size indicates the volume of overall demand in that cluster. The color bar to the right shows whether a cluster is more frequently requested as a pick-up (light-yellow), drop-off (dark-purple), or is relatively balanced between the two (green). The central area of the service area, which we refer to as the center, stands out. It is both the most frequent origin (55.75%) and destination (58.42%), indicated by the balanced yellow color. In contrast, some peripheral clusters lean slightly towards more pick-ups or drop-offs respectively, though lying in moderate ranges. Additionally, we illustrate demand flows with arrows representing ride direction, where arrow thickness corresponds to ride frequency. A hub-and-spoke pattern is evident, with three clusters showing significant bi-directional traffic, highlighted by double-headed arrows.

To provide further insights into the service area under consideration, we report the average requested distances in Appendix A. In this analysis, we categorize all requests based on the lengths of the desired OD pairs and present them in a histogram. The most frequent requests are for distances of less than 5 km (46.35%) or between 5 and 10 km (41.37%). For a more detailed breakdown of the distance distribution, we refer readers to Appendix A.1.

We further enhance interpretability by categorizing requests into four distinct spatial subsets. This aggregation allows for a clearer understanding of spatial demand patterns:

• Rides within the center (29.14%): These are rides with both the origin and destination being within the center.

- Rides from the center (26.92%): These rides start in the center and end at a stop in the surrounding peripheral areas.
- Rides to the center (29.63%): These involve rides originating from peripheral stops with the destination in the center.
- Rides within peripheral area (14.31%): These rides start and end at stops located in the peripheral area.

This classification highlights that only around 15% of requests are "spoke-to-spoke" rides (rides within peripheral area), while the vast majority involve the center as either the origin, destination, or both. This result underscores the hub-and-spoke spatial pattern prevalent in the service area.

Further analysis of the average distances for each spatial type is provided in Appendix A.2. As expected, rides within the center are shortest, with 100% of these rides being less than 5 km. Rides to and from the center show similar distributions, with the majority of rides (64.74% for rides from the center and 67.41% for rides to the center) falling between 5 and 10 km, and smaller percentages (8.93% and 9.34%, respectively) covering distances between 10 and 15 km. In contrast, rides within peripheral areas show a different pattern, with an average distance that is 65.81% longer than the average of the total demand: 36.16% of these rides fall between 10 and 15 km.

#### 3.4 Customer Demand with a Spotlight on the Extended Booking Horizon

In this section, we highlight the unique patterns that result from an extended booking horizon. First, we explore specific trends in the time of request related to desired times and spatial distribution. Thereby, we identify strategic booking behavior in the form of reservations. Second, we examine round trip booking behavior. We start with a brief introduction, followed by an in-depth look at differences in time of request for the outward rides and return rides of round trips. Third, we analyze the characteristic cancellation behavior enabled by the extended booking horizon. Lastly, we examine the time flexibility that we can observe for individual requests in the booking process.

#### 3.4.1 Reservation Behavior

First, we investigate whether an extended booking horizon exhibits specific patterns in desired times. To do so, we utilize the four request types identified in the general analysis: early advance requests, late advance requests, same-day requests, and ad-hoc requests.

Early advance Late advance Same-day requests Ad-hoc requests requests requests 45% Off-peak morning 29% 3% 22% 22% Peak morning 16% 36% 26% Peak afternoon 19% 28% 29% 24% Off-peak evening 14% 22% 29% 35%

 Table 1: Relative frequency of request types during the service horizon

**Time of request – Desired times**: As shown in Fig. 4, there is a general trend that advance requests (both early and late) dominate during the off-peak morning hours, while same-day and adhoc requests increase substantially during the peak periods and ultimately prevail in the off-peak evening (see Table 1). Focusing on late advance and same-day requests, we find that the share of late advance requests is particularly high in the first service hours, while same-day requests are comparatively low. This is largely due to the classification because any request submitted shortly before the start of service is categorized as a late advance request.



Fig. 4 Desired times for different request types

However, the demand patterns of early advance requests are particularly notable. These requests are most frequent during off-peak morning hours (e.g., 6: 00: 43%, 7: 00: 39%), around the afternoon peak at 16: 00 (28%), and in the final service hour (28%). Customers may prefer to reserve rides well in advance, particularly those who, e.g., commute for work.

As the day progresses, same-day and ad-hoc requests rise, peaking in the evening. This indicates a trend toward increased spontaneity among customers later in the day. This spontaneity may be associated with customers' need to plan flexible round trips. For instance, a customer might use an SMOD service for appointments or work. However, due to unpredictable return times the return ride is often requested with a short booking lead time. We explore this pattern in greater detail in Section 3.4.2, where we specifically analyze round trips.

**Time of request – Spatial distribution**: We now examine how the time of request relates to desired OD pairs. Using the four OD pair types and direct OD distances, we aim to identify patterns. As shown in Fig. 5a, the highest proportion of advance requests is observed for peripheral rides, with a combined share of 59.77% (Early: 28.91%, Late: 30.88%). Notably, early advance requests for peripheral rides (28.91%) are higher compared to other OD pair types: within the center (16.53%), from the center (20.85%), and to the center (19.29%). This suggests a moderate preference for early planning for peripheral rides. While rides within and from the center exhibit

relatively balanced booking patterns, an interesting trend emerges for rides to the center. Here, advance requests dominate with a share of 55.58% (Early advance: 19.29%, Late advance: 36.29%), indicating a preference for advance planning, likely driven by commuting, appointments, or shopping trips to the center. This finding implies that advance booking can be crucial for these types of rides. We expect that long rides outside of the center can be pooled much worse with other rides that predominantly occur in and around the center.



Fig. 5 OD-pair types and distances for different request types

Fig. 5b highlights a clear trend that longer rides are more frequently requested in advance. E.g., early advance requests make up 15.00% of the shortest rides (0 - 5 km) but increase to 22.26% for 15 - 20 km rides. Late advance requests also show an upward trend, accounting for 34.16% of the shortest rides and increasing to 41.07% for 15 - 20 km rides. The longest distance range (20 - 25 km), which contains only a small number of requests, is exclusively composed of late advance requests (100%).

This analysis reveals a strong tendency for customers to submit advance requests for longer rides, peripheral rides, and, to some extent, rides to the center. These trends suggest that customers may

prioritize booking longer or peripheral rides in advance, possibly due to greater uncertainty in availability for these ride types.

In Appendix A.3, we combine these patterns and analyze the demand structure of different request types for peripheral rides and their respective desired times. When comparing peripheral rides to requests associated with the other three categories (within, to, and from the center), we find additional evidence that early reservation peaks for specific desired times (e.g., 6:00, 7:00, 16:00, 20:00) are even more pronounced for peripheral rides (49%, 52%, 40%, 37%) compared to rides involving the center (42%, 36%, 25%, 24%). Beyond these desired time peaks, early advance requests generally occur at higher levels for peripheral rides (off-peak morning: 32%, peak morning: 18%) than for rides involving the center (off-peak morning: 29%, peak morning: 16%, peak afternoon: 17%, off-peak evening: 13%).

#### 3.4.2 Round Trips

Now, we focus on customer behavior concerning round trips. A round trip is defined as a pair of requests (outward ride and return ride) that satisfies the following conditions:

- Outward ride and return ride are requested by the same customer.
- The pick-up (drop-off) of the outward ride occurs in the same cluster as the drop-off (pick-up) of the return ride.
- The desired time of the outward ride is earlier than that of the return ride. Both rides take place at the same day.

We provide a general overview of the relevance of round trips and their distinguishing features. Additionally, we explore specific patterns in booking lead times that differ between outward rides and return rides. Overall, 46.76% of all requested rides are part of round trips, making them a crucial consideration for SMOD providers. In this case study, nearly every second customer requests a return ride to their original location.

Fig. 6a visualizes the shift in desired times between outward and return rides. Naturally, the desired time of outward rides lies earlier in the day, while return rides occur closer to the end of the service horizon. On average, the time span between outward and return rides is 227.85 min (around 3.8 hours), with most time spans ranging from 90 min to 334.5 min. A more detailed analysis can be found in Appendix A.4.

In terms of spatial patterns, Fig. 6b shows that most round trips involve the center: 43.79% are within the center, and 34.20% between the center and the peripheral area (back or to the center). This suggests that many customers use the service for rides to the center and back, for example for appointments.

A key observation is the difference in booking lead times between outward and return rides. Outward rides are generally booked with more lead time, averaging 105.35 hours, compared to 77.91 hours for return rides. This suggests that outward rides are often linked to fixed activities, such as

doctor's appointments or commuting, where the start time is precisely known. In contrast, the timing for return rides is often uncertain, as it depends on the length of the stay, resulting in more spontaneous requests.

This difference is also evident in the request types for outward and return rides: 68.02% of outward rides are requested in advance. Considering return rides, only 44.60% of return rides are requested in advance. 29.97% of return rides are requested same-day and 25.43% ad-hoc. This discrepancy can lead to an undesirable situation from the customer's perspective: Assuming a provider applies a first-come-first-served policy without demand management, outward rides are more likely to be accepted as early requests are prioritized. Hence, the return ride corresponding to an accepted outward ride may be rejected, leaving customers "stranded" at their destination. Given that nearly half of all requested rides are part of round trips, this issue is relevant.

#### 3.4.3 Cancellations

In this section, we explore the relevance of cancellations within SMOD systems with extended booking horizons. We define a cancellation as a request that initially became a confirmed order (see Fig. 1) but was later cancelled unilaterally by the customer.

In the previous analyses, we excluded cancelled requests because we assume that they do not reflect real demand. Hence, we have considered only 79.93% of all received requests so far. The remaining 20.07% were cancelled before the desired time or at the desired time (no-show), making it a relevant share. In this section, we analyze the key characteristics of this cancelled demand, including the time of request, the desired time, and the spatial distribution. This descriptive analysis helps identify patterns that may enable SMOD providers to predict the likelihood of a cancellation at the time of request with high accuracy.

Comparing Fig. 3a and Fig. 7a, we observe substantial differences between cancelled and noncancelled demand. Notably, the booking curve for cancelled requests skews toward longer booking lead times. E.g., 38.89% of cancellations are early advance requests (compared to only 20.25% early advance requests for non-cancelled demand), and 40.27% are late advance requests (vs. 32.84%). In contrast, orders resulting from ad-hoc requests and same-day requests are rarely cancelled (3.05% and 17.79%, respectively, vs. 24.05% and 22.86% for non-cancelled demand). These patterns underscore the particularly high relevance of cancellations in SMOD systems with extended booking horizons.

In terms of desired times (Fig. 3b and Fig. 7b), the distribution of cancelled and non-cancelled requests is generally similar. However, certain hours, such as 12:00 and 16:00, are more prone to cancellations. Spatially, cancelled rides are more common in the peripheral area, with a notable shift toward drop-offs in these regions. While the center remains the most common origin (54.02% vs. 55.75%) and destination (50.39% vs. 58.42%), cancellations are more prevalent in rides to the peripheral area, as indicated by the purple dots around the hub in Fig. 7c. These

distinct patterns revealed by comparing cancelled and non-cancelled demand (Fig. 3c and Fig. 7c) suggest potential for machine learning algorithms to predict cancellations at the time of request. Because of an extended booking horizon, a cancelled request may remain in the system, mistakenly treated as true demand, for a certain dwell time, which we call ghost demand dwell time. The provider treats the request as true demand from the time of request until the time of cancellation, although it is actually ghost demand (see Fig. 9). In a didactic example (Fig. 8), the provider receives a request with a lead time of -10 days. The time of cancellation is -2 days before the desired time (cancellation lead time of 2 days). The period between time of request and time of cancellation spans 8 days during which the request, being ghost demand, could have a negative impact on operational planning.



Fig. 6 Temporal and spatial characteristics of round trips



Article A4: Extended Booking Horizons in Rural Shared Mobility-on-Demand Systems: Insights and Implications for Demand Management

c) Service area

Fig. 7 Descriptive analysis of time of request, desired times and origin-destination pairs (cancelled demand)

Article A4: Extended Booking Horizons in Rural Shared Mobility-on-Demand Systems: Insights and Implications for Demand Management



Fig. 8 Illustrative example of an individual cancelled request (ghost demand dwell time, cancellation lead time)

Fig. 9 shows that a considerable proportion of ghost demand remains in the system for up to two weeks, which is only possible because of the extended booking horizon. While 42.17% of all cancelled requests remain in the system for less than one day (0d), 57.83% remain in the system for more than one day, with 29.38% remaining in the system for more than 7 days.



Fig. 9 Distribution of ghost demand dwell time for different request types

Compared to the average ghost demand dwell time (105.24 hours), the average cancellation lead time is lower (36.29 hours). Looking at Fig. 10, we find that the majority of cancellations occur on the same day as the desired day with 72.15% and very few requests are cancelled with a longer cancellation lead time (only 6.86% with at least seven days lead time).

#### 3.4.4 Time Flexibility

In this section, we examine the observed time flexibility of customers when their desired time cannot be met. Time flexibility could be a crucial factor in improving SMOD systems, especially in rural areas with dispersed demand offering few opportunities for ridepooling. Hence, understanding customers' willingness to accept deviations from their desired time can provide valuable insights. Please note that, theoretically, customers could also be asked to be more flexible regarding the waiting time or the maximum added ride time. In practice, however, these parameters are

set uniformly for any request to guarantee a certain service level. Hence, offering alternative pickup/drop-off times is the only practical way of exploiting customers' time flexibility.



Fig. 10 Average cancellation lead time for different request types

By analyzing customer behavior when the ride at the desired time is not offered, we can gain insights into their willingness to accept alternative times. Our analysis provides a lower bound on this flexibility, as it only includes cases where customers place a request that transforms into an order for an alternative time. Due to data limitations, we cannot observe the full offer set presented to customers. This means that we cannot see which options customers have declined. Further, it is possible that customers may exhibit even greater flexibility if they are offered additional options with larger deviations from their desired times. Finally, a substantial proportion of requests (32.28%) are offered their exact desired time and therefore are never "asked" to reveal any time flexibility, let alone their true flexibility.

Our findings are visualized in Fig. 11 using a density function. The results show that most requests exhibit some degree of flexibility, with an overall mean of 29.98 min. Notably, the extended booking horizon allows SMOD providers to offer alternative times both before and after the original desired time, resulting in observations of both positive and negative flexibility. The average positive flexibility (willingness to depart later) is 67.02 min, while the average negative flexibility (willingness to depart earlier) is -18.81 min. Requests are almost uniformly distributed across three categories: 31.96% show flexibility below the median, 35.76% show flexibility above the median, and 32.28% match the median flexibility (0 min).

For further analysis, we segment the requests into three distinct subsets: low-flexibility requests with deviations of fewer than 5 min from the desired time (49.40% of all requests), moderate-flexibility requests with absolute flexibility between 5 min and 30 min (34.72% of all requests), and high-flexibility requests (15.88% of all requests) with over 30 min of flexibility.

Article A4: Extended Booking Horizons in Rural Shared Mobility-on-Demand Systems: Insights and Implications for Demand Management



Fig. 11 Observed time flexibility - Density of observed accepted time deviations from desired time

As Fig. 12 shows, requests submitted early in the booking horizon often show low time flexibility. In contrast, same-day and ad-hoc requests exhibit higher flexibility. Especially at the end of the booking horizon, the share of requests with high observed time flexibility increases markedly, while moderately flexible requests decline.

However, due to the abovementioned data limitations and possible additional biases due to the demand management applied by FLEXIBUS, these findings do not allow the conclusion that customers submitting advance requests are systematically less flexible. However, despite its limitations, our analysis reveals a significant finding: some customers even display negative flexibility, indicating that many customers are willing to adjust their desired times either earlier or later.



Fig. 12 Different flexibility levels over the booking horizon

### 4 Implications for Operational Demand Management

SMOD providers can improve planning reliability for customers by extending the booking horizon. However, descriptive analysis of the demand structure (Section 3.4) reveals four key patterns associated with the extended booking horizon that are prevalent in the considered service area. These patterns impact system performance, which suggests a potential for demand management approaches to both mitigate undesirable effects and take advantage of desirable effects to enhance system performance. In this section, we provide a compact summary of each pattern's impact on operational planning along with suggestions for demand management.

#### **Reservation behavior**

*Impact on operational planning*: We observe that customers disproportionately submit advance requests for certain rides based on the following parameters:

- 1. Desired time (e.g., 6:00-8:00)
- 2. OD pair (peripheral rides)
- 3. Direct OD distance (longer-distance rides).

On the one hand, such strategic reservation behavior is not necessarily detrimental: In general, it could be argued that advance requests are beneficial for the operators' operational planning since they increase the deterministic information known before the service horizon starts. On the other hand, it is also advantageous for the customer since they can reserve rides that may be particularly important to them by planning their trip and requesting their ride earlier than other customers do. On the other hand, strategic reservation behavior can have negative consequences for the provider and other customers. From the provider's point of view, customers taking advantage of early reservations may particularly request "unfavorable rides" (not necessarily in monetary terms). These rides could then displace other customers requesting later in the booking horizon whose rides may be more favorable. Also, from a customer perspective, ordering rides for equally important but spontaneously arising trips may become difficult because the fleet is largely occupied already several days in advance.

For example, it seems likely that peripheral rides have a lower average chance of being accepted, given that they represent only a small portion of total requests (14.31%) and are on average longer than other rides involving the center (Appendix A.2). Therefore, they allow for less consolidation. Whether serving peripheral rides aligns with the provider's objectives depends on factors such as the available supply and the given demand of the service day.

*Demand management suggestion*: If the SMOD provider considers reservation behavior an issue, we suggest the adoption of anticipatory demand management. The SMOD provider may use probabilistic information about future demand when a request arrives to decide on the availability and/or pricing of rides. This involves approximating the opportunity cost of serving a request, which captures the expected displacement of future revenue and the expected increase in routing cost (Fleckenstein et al., 2024).

Anticipatory demand management allows an SMOD provider to steer all demand, including early advance requests, according to their operational objectives. E.g., by applying availability control, a provider can decide which ride options a customer is offered in response to a given request, depending on the opportunity cost of possible ride options (e.g., alternative pick-up times). As a result, the offers made are less dependent on the time of request (Anzenhofer et al., 2024a).

While availability control allows the provider to manage the access to the system very effectively, it is also quite coercive from the customers' perspective. Frequently, customer requests are rejected due to high opportunity cost despite the existence of feasible rides. To give all customers a chance to access the service, dynamic pricing can be used for demand management (Arian et al., 2022). With dynamic pricing, rides with high opportunity cost are still offered but at an increased price. Then, the customer can still decide whether their desired ride is important enough for them to pay the higher price (Eliasson, 2021).

#### **Round trips**

*Impact on operational planning*: A substantial portion of demand (46.76%) involves round trips, underscoring their importance. We observe notable patterns in the request types for outward and return rides: outward rides are more frequently requested in advance (68.02%) than return rides (44.60%). A potential explanation is that customers can precisely specify the desired time for the outward ride well in advance if their planned activity begins at a fixed time. In contrast, return times tend to be more variable and unknown to the customer in advance. If a provider applies a first-come-first-served policy, this could lead to "stranded" customers that are forced to switch to another mode of transport for the return ride on short notice. This might lead to many customers not even considering the SMOD service as a suitable mode of transport. Then, a lot of potential demand is lost.

*Demand management suggestion*: If an SMOD provider wants to facilitate round trips for planned activities with an uncertain end time, our suggestion is twofold:

First, as with reservation behavior, anticipatory demand management can help achieve consistent offer quality across the booking horizon, increasing the likelihood that customers can order their return ride ad-hoc.

However, since this alone does not guarantee the necessary planning reliability to fully avoid "stranding", SMOD providers can additionally apply strategic demand management: E.g., the provider can design a premium ride option that allows the customer to specify a larger, provisional time window within which they are guaranteed a return ride. This would give the customer the opportunity to confirm the exact desired time of their return ride shortly in advance, possibly with a certain minimum booking lead time. E.g., if a customer has a doctor's appointment with a known start time (9: 30), the outward ride can be ordered in advance, while a return ride is guaranteed within a larger timeframe (e.g., 10:00 - 11:00) to account for the uncertain desired return time. With a minimum booking lead time of, e.g., 30 min, the customer could specify their desired

return time within the provisional time window, such as confirming at 09:55 for a 10:25 pickup. While this product design is certainly associated with operational challenges for the provider, it could offer a compromise between sufficient planning reliability for customers and system performance.

Note that the provider can clearly also address reservation and round trip issues by expanding the vehicle fleet. However, this solution comes at considerable cost, including capital expenditures for vehicles and operational expenses such as driver wages. Thus, increasing supply should be viewed as a subordinate solution compared to applying demand management, which requires much lower expenses.

#### Cancellations

*Impact on operational planning*: Unilateral cancellations by customers are inherently problematic for the operational planning of SMOD systems. In our descriptive analysis, their relevance is emphasized by their relative frequency during the observation period (20.07% of all requests submitted result in cancelled orders).

With an extended booking horizon, the issue of cancellations is further exacerbated: Unlike urban providers, which primarily receive ad-hoc requests and mainly deal with ad-hoc cancellations (no-shows), rural providers must consider this "ghost demand" in operational planning, potentially over a large portion of the booking horizon (Fig. 10).

This could distort operational demand management decisions: First, whenever a new request arrives, this ghost demand is considered as true demand when approximating the opportunity cost of the new request. Hence, the consolidation with ghost demand, which is in fact never realized, can lead to an opportunity cost estimation error. E.g., the opportunity cost of a newly arriving request for a ride within the peripheral area could be underestimated due to compatible peripheral ghost demand.

Second, ghost demand blocks fleet resources over its dwell time. Meanwhile, requests that are in fact feasible (were it not for the ghost demand) are rejected. The longer the ghost demand dwell time, the more time there is for such demand displacement to realize. Thus, we conclude that the dwell time of the ghost demand is critical, and the damage likely increases with dwell time.

Independent of the impact on operational demand management, there are also implications for vehicle routing, which we only touch on briefly because it is not the focus of our work: For very late cancellations, routing costs realize because, e.g., the pick-up location is visited by a vehicle. In addition, as it is the case for demand management, route planning is distorted by unrealized consolidation with ghost demand.

*Demand management suggestion*: Unlike reservation behavior or round trips, anticipatory demand management (e.g., Anzenhofer et al., 2024a) cannot mitigate the impact of cancellations unless they are explicitly considered. Hence, a tailored prediction model must be integrated to assess each request's cancellation probability, using information available at request arrival. The
discovery of patterns in cancellations by our descriptive analysis suggests that such a prediction should be possible with high accuracy.

To investigate this hypothesis, we apply XGBoost, a widely used gradient boosting method (Chen and Guestrin, 2016), to predict cancellations. We train the model using a train-test split with 80% of the data set for training and 20% for testing. As hyperparameter values, we set 1000 estimator trees, a maximum depth of 6 for each tree, and a learning rate of 0.1. The set of features includes the requested day of week, the desired hour, the OD type, and the anonymized customer ID. With this model, we achieve a high accuracy of 93.66%. These promising results suggest that a cancellation predictor can, e.g., be incorporated into an anticipatory demand management framework, enabling providers to proactively reject requests based on the cancellation prediction or charge an increased price (Wu et al., 2024).

The high relevance of cancellations also suggests the application of overbooking, i.e., the practice of selling more rides than the available capacity allows as a hedge against customer cancellations. If less demand is cancelled than predicted, the provider must bump customers at pre-determined or negotiated penalties and additional goodwill losses (Klein et al., 2020). Although overbooking is well-established in other applications, particularly the airline industry, there is little literature on (S)MOD systems. The only work to our knowledge is Liu et al. (2021), who develop a general optimization framework for a ridehailing system that allows for the consideration of cancellations via an overbooking policy.

#### Time flexibility

*Impact on operational planning*: Our findings suggest that an extended booking horizon allows SMOD providers to leverage time flexibility both before and after the desired time, which is not possible in systems limited to ad-hoc requests. Hence, SMOD providers have greater opportunities to steer customers toward choosing alternative rides instead of the originally desired ride.

We observe a substantial absolute time flexibility (29.98 min). It is important to note that this observed time flexibility is only a lower bound and the true time flexibility is likely much higher. The average positive flexibility (willingness to accept a later pick-up) is +67.02 min, while the average negative flexibility (willingness to accept an earlier pick-up) is -18.81 min.

*Demand management suggestion*: We suggest that SMOD providers actively leverage the time flexibility of customer requests. Specifically, we suggest implementing demand management using availability control (Anzenhofer et al., 2024a) or dynamic pricing (Arian et al., 2022). E.g., at demand peaks during the service horizon such as the morning peak, it may be beneficial for the provider to actively steer advance requests to the earlier off-peak morning hours, where more supply is available. Anzenhofer et al. (2024a) conduct a computational analysis on the value of time flexibility. They find that demand management approaches utilizing time shifts can improve service reliability (increasing the number of served requests) and environmental sustainability

(improving distance savings relative to private motorized transport), with the benefits increasing the more time flexibility customers have.

# 5 Conclusion and Future Research Directions

This study examined the patterns in customer demand arising from an extended booking horizon in SMOD systems. These patterns offer valuable insights into customer behavior when customers are given the crucial opportunity to plan rides early in advance. Our descriptive analyses revealed clear patterns in reservation behavior, round trip requests, cancellations, and observable time flexibility. Based on these findings, we proposed tailored demand management approaches that can improve the system performance both from the perspective of providers and customers.

Our results give rise to several research questions that require further investigation:

- **Demand management with cancellations**: Despite the rich body of literature on demand management for SMOD systems, approaches explicitly considering cancellations are rare. Future research should focus on developing algorithms that can readily exploit probabilistic information about cancellations.
- Enhanced cancellation prediction models: While we show that simple machine learning models already achieve a high accuracy in predicting cancellations, future research could investigate more sophisticated models that integrate additional features, such as customer preferences, external factors (e.g., weather, events), or customer-specific features. This is especially beneficial if overbooking strategies are used for demand management since erroneous predictions of cancellations would then lead to confirmed orders becoming infeasible. To avoid losses in customer goodwill, expensive short-term measures would be necessary, such as resorting to regular taxis.
- Survey-based research on time flexibility: While we measure the observable time flexibility based on historical orders, future research should aim at determining more accurate approximations of the true time flexibility. Due to historical order data being biased by the provider's decision-making, conducting customer surveys appears to be a promising approach. In particular, such research could reveal the true flexibility of customers placing advance requests.
- Strategic decision support and operational solution algorithms for round trip product: The proposed premium product for round trips deserves further investigation both from the strategic perspective as well as the operational perspective. At the strategic planning level, future research could provide decision support regarding the product design, i.e., the size of the provisional time window for the return ride or the minimum booking lead time for confirming the exact desired pick-up time. At the operational planning level, future research should explore demand management approaches that incorporate the premium product for

round trips. The main challenge is to ensure the guaranteed service given the provisional time window and the minimum booking lead time.

# Acknowledgement

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.

## References

- Al-Kanj, L., Nascimento, J., & Powell, W. B. (2020). Approximate dynamic programming for planning a ride-hailing system using autonomous fleets of electric vehicles. *European Jour*nal of Operational Research, 284(3), 1088-1106. https://doi.org/10.1016/j.ejor.2020.01.033.
- Anzenhofer, F., Fleckenstein, D., Klein, R., & Steinhardt, C. (2024a). Analyzing the impact of demand management in rural shared mobility-on-demand systems. Working paper, University of Augsburg. http://dx.doi.org/10.2139/ssrn.4682056.
- Anzenhofer, F., Fleckenstein, D., Klein, R., & Steinhardt, C. (2024b). Sustainable dynamic pricing for rural shared mobility-on-demand systems. Working paper, University of Augsburg.
- 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.
- Atasoy, B., Ikeda, T., Song, X., & Ben-Akiva, M. E. (2015). The concept and impact analysis of a flexible mobility on demand system. *Transportation Research Part C: Emerging Technol*ogies, 56, 373-392. https://doi.org/10.1016/j.trc.2015.04.009.
- Bar-Yosef, A., Martens, K., & Benenson, I. (2013). A model of the vicious cycle of a bus line. *Transportation Research Part B: Methodological*, 54, 37-50. https://doi.org/10.1016/j.trb.2013.03.010.
- Chandakas, E. (2020). On demand forecasting of demand-responsive paratransit services with prior reservations. *Transportation Research Part C: Emerging Technologies*, *120*, 102817. https://doi.org/10.1016/j.trc.2020.102817.
- Chen, T., & Guestrin, C. (2016). Xgboost: A scalable tree boosting system. In Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining, 785-794. https://doi.org/10.1145/2939672.2939785.
- Chen, H., Jiao, Y., Qin, Z., Tang, X., Li, H., An, B., Hongtu, Z., & Ye, J. (2019). InBEDE: Integrating contextual bandit with TD learning for joint pricing and dispatch of ride-hailing platforms. In 2019 IEEE International Conference on Data Mining (ICDM), 61-70. https://doi.org/10.1109/ICDM.2019.00016.
- Eliasson, J. (2021). Efficient transport pricing Why, what, and when?. *Communications in Transportation Research*, *1*, 100006. https://doi.org/10.1016/j.commtr.2021.100006.
- Elting, S., & Ehmke, J. F. (2021). Potential of shared taxi services in rural areas A case study. *Transportation Research Procedia*, *52*, 661-668. https://doi.org/10.1016/j.trpro.2021.01.079.
- Engelhardt, R., Dandl, F., & Bogenberger, K. (2022). Simulating ride-pooling services with prebooking and on-demand customers. https://doi.org/10.48550/arXiv.2210.06972.
- 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: 13/11/2024.
- Fleckenstein, D., Klein, R., Klein, V., & Steinhardt, C. (2024). On the concept of opportunity cost in integrated demand management and vehicle routing. *Transportation Science, online first*. https://doi.org/10.1287/trsc.2024.0644.
- Gödde, J., Ruhrort, L., Allert, V., & Scheiner, J. (2023). User characteristics and spatial correlates of ride-pooling demand – Evidence from Berlin and Munich. *Journal of Transport Geography*, 109, 103596. https://doi.org/10.1016/j.jtrangeo.2023.103596.
- Haferkamp, J., & Ehmke, J. F. (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.
- Haferkamp, J., Ulmer, M. W., & Ehmke, J. F. (2024). Heatmap-based decision support for repositioning in ride-sharing systems. *Transportation Science*, 58(1), 110-130. https://doi.org/10.1287/trsc.2023.1202.
- Haliem, M., Mani, G., Aggarwal, V., & Bhargava, B. (2021). A distributed model-free ride-sharing approach for joint matching, pricing, and dispatching using deep reinforcement learning.

*IEEE Transactions on Intelligent Transportation Systems*, 22(12), 7931-7942. https://doi.org/10.1109/TITS.2021.3096537.

- He, F., Wang, X., Lin, X., & Tang, X. (2018). Pricing and penalty/compensation strategies of a taxi-hailing platform. *Transportation Research Part C: Emerging Technologies*, 86, 263-279. https://doi.org/10.1016/j.trc.2017.11.003.
- Heitmann, R. J. O., Soeffker, N., Ulmer, M. W., & Mattfeld, D. C. (2023). Combining value function approximation and multiple scenario approach for the effective management of ridehailing services. *EURO Journal on Transportation and Logistics*, 12, 100104. https://doi.org/10.1016/j.ejtl.2023.100104.
- Horn, M. E. (2002). Fleet scheduling and dispatching for demand-responsive passenger services. *Transportation Research Part C: Emerging Technologies*, 10(1), 35-63. 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. *Transportation Research Part B: Methodological*, 70, 303-318. https://doi.org/10.1016/j.trb.2014.09.011.
- Huang, D., Gu, Y., Wang, S., Liu, Z., & Zhang, W. (2020). A two-phase optimization model for the demand-responsive customized bus network design. *Transportation Research Part C: Emerging Technologies*, 111, 1-21. https://doi.org/10.1016/j.trc.2019.12.004.
- Hungerländer, P., Maier, K., Pachatz, V., & Truden, C. (2021). Improving sharing rates of a diala-ride problem implemented for an Austrian mobility provider. *Transportation Research Procedia*, 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. *Research in Transportation Economics*, 99, 101301. https://doi.org/10.1016/j.retrec.2023.101301.
- Jung, J., Jayakrishnan, R., & Park, J. Y. (2016). Dynamic shared-taxi dispatch algorithm with hybrid-simulated annealing. *Computer-Aided Civil and Infrastructure Engineering*, 31(4), 275-291. https://doi.org/10.1111/mice.12157.
- Klein, R., Koch, S., Steinhardt, C., & Strauss, A. K. (2020). A review of revenue management: Recent generalizations and advances in industry applications. *European Journal of Operational Research*, 284(2), 397-412. https://doi.org/10.1016/j.ejor.2019.06.034.
- Li, X., Xu, J., Du, M., Liu, D., & Kwan, M. P. (2023). Understanding the spatiotemporal variation of ride-hailing orders under different travel distances. *Travel Behaviour and Society*, *32*, 100581. https://doi.org/10.1016/j.tbs.2023.100581.
- Liu, Y., Luo, Q., Gopalakrishnan, R., & Samaranayake, S. (2021). A framework for the joint optimization of assignment and pricing in mobility-on-demand systems with shared rides. https://doi.org/10.48550/arXiv.2112.14297.
- Lotfi, S., & Abdelghany, K. (2022). Ride matching and vehicle routing for on-demand mobility services. *Journal of Heuristics*, *28*(3), 235-258. https://doi.org/10.1007/s10732-022-09491-7.
- Lu, C., Maciejewski, M., Wu, H., & Nagel, K. (2023). Optimization of demand-responsive transport: The rolling horizon approach. *Procedia Computer Science*, 220, 145-153. https://doi.org/10.1016/j.procs.2023.03.021.
- Ma, Z., & Koutsopoulos, H. N. (2022). Near-on-demand mobility. The benefits of user flexibility for ride-pooling services. *Transportation Research Part C: Emerging Technologies*, 135, 103530. https://doi.org/10.1016/j.trc.2021.103530.
- Mounce, R., Beecroft, M., & Nelson, J. D. (2020). On the role of frameworks and smart mobility in addressing the rural mobility problem. *Research in Transportation Economics*, *83*, 100956. https://doi.org/10.1016/j.retrec.2020.100956.
- Ni, L., Sun, B., Wang, S., & Tsang, D. H. (2021). Dynamic pricing mechanism design for electric mobility-on-demand systems. *IEEE Transactions on Intelligent Transportation Systems*, 23(8), 11361-11375. https://doi.org/10.1109/TITS.2021.3103199.

- Qiu, H., Li, R., & Zhao, J. (2018). Dynamic pricing in shared mobility on demand service. Working paper, Massachusetts Institute of Technology.
- Sharif Azadeh, S., Atasoy, B., Ben-Akiva, M. E., Bierlaire, M., & Maknoon, M. Y. (2022). Choice-driven dial-a-ride problem for demand responsive mobility service. *Transportation Research Part B: Methodological*, 161, 128-149. https://doi.org/10.1016/j.trb.2022.04.008.
- Shulika, O., Bujak, M., Ghasemi, F., & Kucharski, R. (2024). Spatiotemporal variability of ridepooling potential – Half a year New York City experiment. *Journal of Transport Geography*, 114, 103767. https://doi.org/10.1016/j.jtrangeo.2023.103767.
- Sörensen, L., Bossert, A., Jokinen, J. P., & Schlüter, J. (2021). How much flexibility does rural public transport need? – Implications from a fully flexible DRT system. *Transport Policy*, 100, 5-20. https://doi.org/10.1016/j.tranpol.2020.09.005.
- Sun, L., Chen, J., Li, Q., & Huang, D. (2020). Dramatic uneven urbanization of large cities throughout the world in recent decades. *Nature Communications*, 11(1), 5366. https://doi.org/10.1038/s41467-020-19158-1.
- Sun, H., Lv, Z., Li, J., Xu, Z., & Sheng, Z. (2023). Will the order be canceled? Order cancellation probability prediction based on deep residual model. *Transportation Research Record*, 2677(6), 142-160. https://doi.org/10.1177/03611981221144279.
- Talluri, K. T., & Van Ryzin, G. J. (2004). The theory and practice of revenue management. New York, NY: Springer. https://doi.org/10.1007/b139000.
- Theodoridis, N., Fielbaum, A., Alonso-Mora, J., & Atasoy, B. (2023). Quantifying the benefits of short-term reservations in on-demand ridepooling. In *Australasian Transport Research Forum (ATRF), 44th*.
- Waßmuth, K., Köhler, C., Agatz, N., & Fleischmann, M. (2023). Demand management for attended home delivery – A literature review. *European Journal of Operational Research*, 311(3), 801-815. https://doi.org/10.1016/j.ejor.2023.01.056.
- Wang, X., Liu, W., Yang, H., Wang, D., & Ye, J. (2019). Customer behavioural modelling of order cancellation in coupled ride-sourcing and taxi markets. *Transportation Research Procedia*, 38, 853-873. https://doi.org/10.1016/j.trpro.2019.05.044.
- Wang, G., Zhang, H., & Zhang, J. (2024). On-demand ride-matching in a spatial model with abandonment and cancellation. *Operations Research*, 72(3), 1278-1297. https://doi.org/10.1287/opre.2022.2399.
- Weckström, C., Mladenović, M. N., Ullah, W., Nelson, J. D., Givoni, M., & Bussman, S. (2018). User perspectives on emerging mobility services: Ex post analysis of Kutsuplus pilot. *Research in Transportation Business & Management*, 27, 84-97. https://doi.org/10.1016/j.rtbm.2018.06.003.
- Wu, W., Zou, H., & Liu, R. (2024). Prediction-failure-risk-aware online dial-a-ride scheduling considering spatial demand correlation via approximate dynamic programming and scenario approach. *Transportation Research Part C: Emerging Technologies*, 169, 104801. https://doi.org/10.1016/j.trc.2024.104801.
- Xu, K., Saberi, M., & Liu, W. (2022). Dynamic pricing and penalty strategies in a coupled market with ridesourcing service and taxi considering time-dependent order cancellation behaviour. *Transportation Research Part C: Emerging Technologies*, 138, 103621. https://doi.org/10.1016/j.trc.2022.103621.
- Yu, X., Shen, S., & Wang, H. (2021). Integrated vehicle routing and service scheduling under time and cancellation uncertainties with application in nonemergency medical transportation. *Service Science*, 13(3), 172-191. https://doi.org/10.1287/serv.2021.0277.
- Zwick, F., & Axhausen, K. W. (2022). Ride-pooling demand prediction: A spatiotemporal assessment in Germany. *Journal of Transport Geography*, 100, 103307. https://doi.org/10.1016/j.jtrangeo.2022.103307.

Zwick, F., Fraedrich, E., & Axhausen, K. W. (2023). Ride-pooling in the light of COVID-19: Determining spatiotemporal demand characteristics on the example of MOIA. *IET Intelligent Transport Systems*, 17(6), 1166-1181. https://doi.org/10.1049/itr2.12293.

# Appendix A.1 – Distribution of Direct Origin-Destination Pair Distance (Total Demand)



Fig. 13 Relative frequency of requests associated with a specific origin-destination pair distance

Appendix A.2 – Distribution of Origin-Destination Pair Distance (Spatial Type)

Article A4: Extended Booking Horizons in Rural Shared Mobility-on-Demand Systems: Insights and Implications for Demand Management



d) Rides within peripheral areas

Fig. 14 Relative frequency of requests with a specific origin-destination pair distance per spatial request type





Fig. 15 Relative proportions per request type for different desired times





Fig. 16: Boxplot on time difference between outward and return rides in min

# Article A5: Analyzing the Impact of Demand Management in Rural Shared Mobility-on-Demand Systems

Fabian Anzenhofer<sup>a</sup>, David Fleckenstein<sup>a</sup>, Robert Klein<sup>a</sup>, Claudius Steinhardt<sup>b</sup>
(<sup>a</sup>University of Augsburg, <sup>b</sup>University of the Bundeswehr Munich)
To appear in OR Spectrum (category A)

#### 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.

Key words: 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, 2024, EPA, 2024), 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 longterm 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 FLEXI-BUS 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 proofof-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 Section 2, we first review the literature and distinguish our work from the existing publications. In Section 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 Section 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 multimodal 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 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 adhoc 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.

	Problem as	em and instance structure			Solution concept		Computational study	
Authors	Demand control	Advance requests	Same-day requests	Ad-hoc requests	Rural instances	MDP	Solution concept	Data
Araldo et al. (2019)	FE	Х	Х	$\checkmark$	Х	Х	М	А
Arian et al. (2022)	PR	Х	Х	$\checkmark$	$\checkmark$	$\checkmark$	L	А
Atasoy et al. (2015)	AV	Х	$\checkmark$	$\checkmark$	Х	х	М	А
Attanasio et al. (2004)	FE	(√)	$\checkmark$	Х	Х	х	М	D
Bischoff et al. (2017)	FE	Х	Х	$\checkmark$	Х	х	М	Ο
Elting and Ehmke (2021)	FE	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	х	М	А
Haferkamp and Ehmke (2020)	FE	Х	Х	$\checkmark$	Х	х	М	А
Haferkamp and Ehmke (2022)	AR	Х	Х	$\checkmark$	Х	$\checkmark$	(S)	D
Heitmann et al. (2023)	AR	Х	Х	$\checkmark$	Х	$\checkmark$	L	D
Horn (2002)	FE	Х	$\checkmark$	$\checkmark$	Х	х	М	О
Hosni et al. (2014)	AR	Х	Х	$\checkmark$	Х	х	М	А
Hungerländer et al. (2021)	FE	Х	Х	$\checkmark$	$\checkmark$	х	М	Ο
Jung et al. (2012)	FE	Х	Х	$\checkmark$	Х	х	М	D
Jung et al. (2016)	AR	Х	Х	$\checkmark$	Х	х	М	D
Lotfi and Abdelghany (2022)	AR	Х	Х	$\checkmark$	Х	х	М	А
Lotze et al. (2023)	(FE)	Х	$\checkmark$	$\checkmark$	$\checkmark$	х	М	D
Lu et al. (2023)	FE	(√)	Х	$\checkmark$	$\checkmark$	х	М	Ο
Qiu et al. (2018)	PR	Х	Х	$\checkmark$	Х	$\checkmark$	L	D
Sharif Azadeh et al. (2022)	AV/PR	Х	$\checkmark$	$\checkmark$	Х	х	М	0
Our work	AV	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	S	Ο

Table 1 Literature overview; 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)

## 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 Section 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 Section 3.2 and the solution concept in Section 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).





Fig. 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 Section 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 (Section 3.2.2).

#### 3.2.1 Operational Markov Decision Process Formulation

#### General Notation and Assumptions:

We formalize the r-SMCP as a Markov decision process (Puterman, 2014), using a consider-thenchoose 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, ..., t<sup>s</sup>, ..., T} with t ∈ T denoting each individual stage. t<sup>s</sup> 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 *H*. Formally, a request of type c ∈ C is characterized by the following attributes:
  - Pick-up stop:  $p_c \in \mathcal{H}$
  - Drop-off stop:  $d_c \in \mathcal{H}$
  - Number of passengers:  $m_c$
  - Desired time:  $t_c \in \{t^s, ..., T\}$
  - 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  $f_{c_i}$ .

Additionally, each request has a time of request  $\tau_i \in \mathcal{T}$ , which represents the stage during which the request is placed. Based on the desired time and the time of request, we can classify requests into three categories:

- Advance requests:  $\tau_i < t^s$  (request is placed before the service horizon starts)
- Same-day requests:  $\tau_i \ge t^s$  (request is placed on the same day but with a booking lead time)

• Ad-hoc requests:  $\tau_i = t_{c_i}$  (request is placed for immediate service)

Finally, to model the case of no request arrival, we introduce a dummy request type c = 0.

- Fulfillment options: A fulfillment option o ∈ O<sub>c</sub> 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<sub>c</sub>. The set O<sub>c</sub> includes all fulfillment options that can potentially be offered. When a customer places a request with a desired time t<sub>c</sub>, 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 abadon the booking process)

Thus, the set of potential fulfillment options  $O_c$  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  $g \subseteq O_c$ , 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 ∈ g based on probabilities P<sub>c,o</sub>(g) reflecting their time preferences. If an option o ≠ 0 is chosen, the request i ∈ J becomes an order j ∈ J with i = j. Thus, the set of orders J is a subset of the set of requests J (J ⊆ J). In addition to the attributes of the corresponding request, an order j is further characterized by its associated fulfillment option o<sub>j</sub> ∈ O<sub>c<sub>i</sub></sub>.
- Order fulfillment: To fulfill the orders, the provider deploys vehicles v ∈ V from a given fleet
   V. Each vehicle has the following attributes:
  - Seat capacity:  $Q_v$  (maximum number of passengers per vehicle)
  - Start time:  $t_{v}^{b} \in \{t^{s}, ..., T\}$  (start of the service horizon)
  - End time:  $t_v^r \in \{t^s, ..., T\}$  (end of the service horizon)
  - Start and end location: h = 0 (the vehicle starts and ends its route at the depot)

of vehicle The planned route each is encoded as the set  $\theta_n =$  $\{(j_1, h_{j_1}, a_{j_1}^-, a_{j_1}^+), (j_2, h_{j_2}, a_{j_2}^-, a_{j_2}^+), \dots, (j_n, h_{j_n}, a_{j_n}^-, a_{j_n}^+)\}, \text{ where } h_{j_n} \in \mathcal{H} \text{ is the } n\text{-th stop of } n\text{-th stop of } n\text{-th stop of } n\text{-th stop of } n\text{-th stop } n\text{-th st} n\text{-th stop } n\text{-th stop } n\text{-th stop } n\text{-th stop } n\text$ the route.  $j_n \in \mathcal{J}$  encodes the corresponding order the vehicle picks up or drops off.  $a_{j_n} \in \mathcal{J}$  $\{t^s, ..., T\}$  and  $a_{j_n}^+ \in \{t^s, ..., T\}$  encode the vehicle's arrival time at and the departure time from the stop, respectively.



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

#### 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. Fig. 2 depicts an intuitive visualization in the form of a decision tree. Fig. 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 ∈ T = {1, ..., t<sup>s</sup>, ..., 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<sup>t</sup><sub>c</sub> 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  $v \in \mathcal{V}$ )

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<sub>t</sub> = (g<sub>t</sub>, (φ<sub>t</sub>(o))<sub>o∈g<sub>t</sub></sub>), 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<sub>t</sub> is associated with integrated vehicle routing decisions (φ<sub>t</sub>(o))<sub>o∈g<sub>t</sub></sub> for any option o ∈ g<sub>t</sub> the customer could potentially choose when presented the offer set g<sub>t</sub>. Note that both the control decision and the integrated vehicle routing decision are interdependent.

#### Availability Control Decision:

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  $2^{\mathcal{O}_c}$  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), 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  $C_{t-1}$ .
- $\circ$  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.,  $\mathcal{G}(s_{t-1}, 0) = \{\{0\}\}.$ 

#### Integrated Vehicle Routing Decisions:

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<sub>t</sub><sup>pre</sup> when a new request i<sub>t</sub> ∈ J arrives (also see Powell, 2022). If a request arrives (c<sub>it</sub> ≠ 0), Step 2 follows. If no request arrives (c<sub>it</sub> = 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 gt, the customer either confirms their order jt by choosing a option ojt ∈ gt, or they abandon the booking process. The order confirmation follows customer-specific choice probabilities Pcito(gt). If the order jt is confirmed, it is added to the set of orders Ct-1.
  - Step 4 (Integrated vehicle routing decision): Finally, the process reaches the succeeding 0 post-decision state  $s_t$  by a deterministic update of the route plan. The route plan  $\phi_t(o_{j_t})$ determined as part of the routing decision replaces the route plan  $\phi_{t-1}$  in the system state and is potentially partly executed. As already defined,  $\phi_t(o_{j_t})$  is the route plan, pre-determined in the vehicle routing decision, specifically for the case that the customer chooses fulfillment option  $o_{j_t}$ . If in  $\phi_t(o_{j_t})$  vehicle movements are planned to start until decision epoch t + 1, the respective stops  $\psi_{v,t} \left( \phi_t(o_{j_t}) \right) = \{ (j, h_j, a_j^-, a_j^+) \in \theta_{v,t} : \theta_{v,t} \in \theta_{v,t} \}$  $\phi_t(o_{j_t}), a_j^+ = t + 1$  are removed from the individual routes  $\theta_{v,t} \in \phi_t(o_{j_t})$  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  $\Psi_t(\phi_t(o_{j_t})) = \{ j \in C_{t-1} : (j, h_j, a_j^-, a_j^+) \in \bigcup_{v \in \mathcal{V}} \psi_{v,t}(\phi_t(o_{j_t})), h_j = d_{c_j} \}, \text{ i.e., based}$ on the drop-offs  $(h_j = d_{c_i})$  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:

$$C_t = (C_{t-1} \cup \{j_t\}) \setminus \Psi_t \left( \phi_t(o_{j_t}) \right)$$
(1)

$$\phi_t = \left\{ \theta_{\nu,t} \setminus \psi_{\nu,t} \left( \phi_t(o_{j_t}) \right) : \theta_{\nu,t} \in \phi_t(o_{j_t}) \right\}$$
(2)

- *Rewards:* The provider collects two types of rewards:
  - Availability-control-related rewards ( $r_c \ge 0$ ): 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  $\psi_{v,t}(\phi_t(o_{j_t}))$ as follows:

$$r_{\phi_t(o)} = \begin{cases} -\sum_{v \in \mathcal{V}} \sum_{h: (j,h,a_j^-,a_j^+) \in \psi_{v,t}(\phi_t(o))} \rho_{h,h'}, & \text{if } \exists \psi_{v,t}(\phi_t(o)) \neq \emptyset \\ 0, & \text{otherwise} \end{cases},$$
(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<sub>t</sub>(s<sub>t</sub>), i.e., the expected future reward, for each state s<sub>t</sub> and decision epoch t ∈ T.

$$V_{t-1}(s_{t-1}) = \sum_{c \in \mathcal{C}} \lambda_c^t \max_{g_t \in \mathcal{G}(s_{t-1},c)} \left( \sum_{o \in g_t} P_{c,o}(g_t) \left[ r_c \cdot \mathbf{1}_{o \neq 0} + \max_{\phi_t(o) \in \Phi(s_{t-1},c,o)} \left( r_{\phi_t(o)} + V_t(s_t | s_{t-1}, c, \phi_t(o)) \right) \right] \right) + \lambda_0^t \max_{\phi_t(0) \in \Phi(s_{t-1},0,0)} \left( r_{\phi_t(0)} + V_t(s_t | s_{t-1}, 0, \phi_t(0)) \right),$$
(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  $\lambda_c^t$ . 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 provider must also decide on a routing  $\phi_t(o) \in \Phi(s_{t-1},c,o)$  (·). 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:

$$V_{t-1}(s_{t-1}) = \sum_{c \in \mathcal{C}} \lambda_c^t \max_{g_t \in \mathcal{G}(s_{t-1},c)} \left( \sum_{o \in g_t} P_{c,o}(g_t) [r_c \cdot \mathbf{1}_{o \neq 0} + V_t'(s_t' | s_{t-1}, c, o)] \right) + \lambda_0^t \cdot V_t'(s_t' | s_{t-1}, 0, 0),$$
(5)  
with  $V_t'(s_t' | s_{t-1}, c, o) = \max_{\phi_t(o) \in \Phi(s_{t-1},c,o)} \left( r_{\phi_t(o)} + V_t(s_t | s_{t-1}, c, \phi_t(o)) \right).$ 

With another transformation, we can reformulate the availability control subproblem in (5) based on the opportunity  $\cot \Delta V_t(s_{t-1}, c, o) = V'_t(s'_t|s_{t-1}, c, 0) - V'_t(s'_t|s_{t-1}, c, o)$  of converting a request of type *c* with option *o* into an order:

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

#### Consider-then-choose model for r-SMCP:

In our model, the consideration set  $S_{l,c}$  for each customer segment  $l \in \mathcal{L}$  and request types  $c \in \mathcal{C}$  is determined by two quality cut-offs:  $\Delta_l^+$  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  $S_{l,c} = \{o \in \mathcal{O}_c : o - t_c \leq \Delta_l^+ \land o - t_c \geq -\Delta_l^-\}$ .

Further, we assume a unique ranking function  $\zeta$  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:

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

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

with  $\gamma_l$  denoting the share of segment *l* in the customer population.

#### 3.2.2 Semi-Perfect Information Model

#### Motivation and Outline:

While the operational MDP introduced in Section 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* ∈ *T*. This depends on the arrival rate λ<sup>t</sup><sub>c</sub>.
- *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 requests 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  $\lambda_c^t$  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 Section 4.4).

#### Model formulation:

We now explain the resulting mathematical formulation of the semi-perfect information model and how it differs from the operational MDP formulation (Section 3.2.1). 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  $l_{i_t}$  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.



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

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  $\lambda_c^t$  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:

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

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  $\zeta$  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 G'(s<sub>t-1</sub>, c) = (O<sub>c</sub><sup>f</sup> ∩ S<sub>l,c</sub>) ∪ {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<sub>jt</sub> = g<sub>t</sub> ∈ (O<sub>c</sub><sup>f</sup> ∩ S<sub>l,c</sub>)), or the customer is rejected (o<sub>jt</sub> = g<sub>t</sub> = 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  $s_t^{\text{pre}}$  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:

$$V_{t-1}(s_{t-1}) = \sum_{c \in \mathcal{C}} \lambda_c^t \cdot \max_{g_t \in \mathcal{G}'(s_{t-1},c)} \left( r_c \cdot \mathbf{1}_{g_t \neq 0} - \Delta V_t(s_{t-1},c,g_t) \right) + V_t'(s_t'|s_{t-1},0,0).$$
(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  $\lambda_c^t$  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 equation (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 exogeneous information. Thus, as part of our methodology, we define heuristic availability control policies for the r-SMCP (Section 3.3.1) and compare the performance of state-of-the-art solution algorithms that are representative of each policy (Section 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 semiperfect 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 Section 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.

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.

1  $C_0^{act} \coloneqq \emptyset$ 2  $\phi_0^{act}, \phi_0^{sam} \coloneqq \{\theta_v = \{(0,0,t_v^b,t_v^b), (0,0,t_v^r,t_v^r)\}: v \in \mathcal{V}\}$  $C_0^{sam} \coloneqq draw_sample(day_type, AR^{sam})$ 3  $\phi_0^{sam} \coloneqq parallel_insertion(\phi_0^{sam}, C_0^{sam})$ 4 5 forall  $t \in \mathcal{T}$  do  $\phi_t^{act}, C_t^{act} \coloneqq execute\_route\_plan(\phi_{t-1}^{act}, C_{t-1}^{act}, \tau_{i_t})$ 6  $\phi_t^{sam}, C_t^{sam} \coloneqq synchronize\_route\_plans(\phi_t^{act}, C_{t-1}^{sam})$ 7  $\mathcal{O}^f_{c_{i_*}} \coloneqq \emptyset$ 8 forall  $o \in \mathcal{O}_{c_{i_t}}$  do 9  $\phi_t^{act}(o), \mathcal{O}_{c_{i_*}}^f \coloneqq feasibility\_check\left(\phi_t^{act}, \mathcal{O}_{c_{i_*}}^f, i_t, o\right)$ 10 if  $o \in \mathcal{O}_{c_i}^f$  do 11  $\Delta \widetilde{V}_t(s_{t-1}, c_{i_t}, o), \phi_t^{sam}(o) \coloneqq opportunity\_cost\_estimation(\phi_t^{act}, \phi_t^{act}(o), \phi_t^{sam}, (c_{i_t}, \tau_{i_t}, o))$ <u>12</u>  $o_{j_t} \coloneqq demand\_control\left(i_t, \left\{\Delta \widetilde{V}_t(s_{t-1}, c_{i_t}, o): o \in \mathcal{O}_{c_{i_t}}^f\right\}\right)$ 13 if  $o_{i_t} \neq 0$  then 14  $C_t^{act} \coloneqq C_t^{act} \cup \{j_t\}$ 15  $\phi_t^{act} \coloneqq routing\_control(\phi_t^{act}(o_{i_t}), C_t^{act})$ 16  $\phi_t^{sam}, C_t^{sam} \coloneqq update\_sampled\_route\_plan(\phi_t^{sam}(o_{i_t}), C_t^{act})$ 17 Fig. 4 Basic solution algorithm

Statements with italic line numbers are only needed for the AC. In statements with an underlined line number, the variables  $\phi_t^{sam}$  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  $\phi_t^{act}$ , 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  $\phi_t^{act}$  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  $o \in O_c \cap S_{l,c}$  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  $O_{c_{i_t}}^f$ . 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  $\Delta \tilde{V}_t(s_{t-1}, c, o) \coloneqq 0$ . 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<sub>t-1</sub>. 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., ΔV<sub>t</sub>(s<sub>t-1</sub>, c<sub>it</sub>, o) := cost(φ<sub>t</sub><sup>act</sup>(o)) cost(φ<sub>t</sub><sup>act</sup>). 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  $\phi_t^{sam}$ . At the beginning of the booking horizon, the sampled route plan  $\phi_0^{sam}$  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  $\phi_t^{sam}$  with the actual route plan  $\phi_t^{act}$  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  $o_{j_t} \coloneqq \max_{o \in \mathcal{G}'(s_{t-1},c_{i_t})} \left( r_{c_{i_t}} \cdot \mathbf{1}_{o \neq 0} - \Delta \tilde{V}_t(s_{t-1},c_{i_t},o) \right)$ . 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  $\max_{\phi_t(o_{j_t})\in\Phi(s_{t-1},c,o_{j_t})} \left(r_{\phi_t(o_{j_t})} + V_t\left(s_t|s_{t-1},c,\phi_t(o_{j_t})\right)\right)$ . 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  $\phi_t^{act}$  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  $o \in O_c$  for any possible offer set  $g \in$  $2^{O_c} \setminus \emptyset$ . 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 Section 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 Section 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 Section 4.2. After that, we discuss the results for the base scenario (Section 4.3) as well as the sensitivity analyses of demand-side flexibility (Section 4.4), level of profitability (Section 4.5), and supply-side capacity (Section 4.6). Finally, we demonstrate that our methodology can also be applied to support other strategic decisions (Section 4.7). We performed all computations on an Intel<sup>®</sup> Core<sup>®</sup> 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-hour 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.) 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. Fig. 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 minutes. 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.



**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-hour 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-hour 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.

#### 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  $i \in J$ . 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, dayspecific 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,... or a later pick-up at o = 10:10,10:20,... 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  $c \in C$ , 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  $\in$  (one zone) to 9.9  $\in$  (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  $0.3 \frac{\epsilon}{km}$ , 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 (Section 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 in Section 3.2.2. 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 Section 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 (Section 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 Section 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 Section 4.3.1, we point out the general performance differences and discuss explanations for and implications from them. In Section 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 0.003s and 0.005s for the non-anticipatory policies and only increases to around 0.007s for the anticipatory policies. This indicates that our methodology can provide results even for considerably larger SMOD systems than that of FLEX-IBUS in a reasonable time frame.

Doliou -	Profit [€]		Number of orders		Distance savings [km]		Pooling rate	
Policy -	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

 Table 2 Results overview of the base scenario

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  $6 \in$  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 (Section 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.

#### **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 nontime-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.



**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.

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.



**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.

#### **Desired time:**

For this analysis, we group the requests into 1-hour 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.

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.



**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.

#### **OD-pair length:**

Now, we investigate the control behavior in spatial terms by considering requests with an ODpair 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 km 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).



**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.

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.



**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 (b), and profitability time shift per order in minutes (c). Each series corresponds to one of the policies FC, MC, and AC.

#### 4.4 Sensitivity Analysis: Demand-side Flexibility

As discussed in Section 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.

Fig. 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 FLEXI-BUS 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.



**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.

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) 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. **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

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.



**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.

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 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 (Section 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. Fig. 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 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.



**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 gray indicates the supply of shift hours.

# 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 (Section 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 (Section 5.2). Finally, we outline possible future research directions (Section 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 minutes yields profit gains between 3% and 5% per 10 minutes of additional flexibility. However, after 120 minutes, 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-minute 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-topay, 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 0.003s 0.005s for non-anticipatory policies and 0.007s 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 Section 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.

# Acknowledgment

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.

### References

- Allcott, H. (2013). The welfare effects of misperceived product costs: Data and calibrations from the automobile market. *American Economic Journal: Economic Policy*, 5(3), 30-66. https://doi.org/10.1257/pol.5.3.30.
- Alonso-González, M. J., 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. *Transportation Research Record*, 2672(8), 879-889. https://doi.org/10.1177/0361198118790842.
- Alonso-Mora, J., Samaranayake, S., Wallar, A., Frazzoli, E., & Rus, D. (2017). On-demand highcapacity ride-sharing via dynamic trip-vehicle assignment. *Proceedings of the National Academy of Sciences of the United States of America*, 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. *Management Science*, 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. In 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. https://doi.org/10.48550/arXiv.2302.01709.
- Atasoy, B., Ikeda, T., Song, X., & Ben-Akiva, M. E. (2015). The concept and impact analysis of a flexible mobility on demand system. *Transportation Research Part C: Emerging Technol*ogies, 56, 373-392. https://doi.org/10.1016/j.trc.2015.04.009.
- Attanasio, A., Cordeau, J. F., Ghiani, G., & Laporte, G. (2004). Parallel tabu search heuristics for the dynamic multi-vehicle dial-a-ride problem. *Parallel Computing*, 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. Transportation Research Part B: Methodological, 154, 266-288. https://doi.org/10.1016/j.trb.2021.10.008.
- Banister, D. (2011). Cities, mobility and climate change. *Journal of Transport Geography*, 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. *Transportation Research Part B: Methodological*, 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. *Operations Research*, 67(1), 143-162. https://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, A. M., & Savelsbergh, M. W. P. (2005). Decision support for consumer direct grocery initiatives. *Transportation Science*, 39(3), 313-327. https://doi.org/10.1287/trsc.1040.0105.
- Campbell, A. M., & Savelsbergh, M. W. P. (2006). Incentive schemes for attended home delivery services. *Transportation Science*, 40(3), 327-341. https://doi.org/10.1287/trsc.1050.0136.
- Cordeau, J. F. (2006). A branch-and-cut algorithm for the dial-a-ride problem. *Operations Research*, 54(3), 573-586. https://doi.org/10.1287/opre.1060.0283.

- Coutinho, F. M., van Oort, N., Christoforou, Z., Alonso-González, M. J., 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. *Research in Transportation Economics*, 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. *Research in Transportation Economics*, *83*, 100895. https://doi.org/10.1016/j.retrec.2020.100895.
- EEA (2024). Greenhouse gas emissions from transport in Europe. Retrieved from https://www.eea.europa.eu/en/analysis/indicators/greenhouse-gas-emissions-from-transport. Last accessed: 23/12/2024.
- Elting, S., & Ehmke, J. F. (2021). Potential of shared taxi services in rural areas A case study. *Transportation Research Procedia*, *52*, 661-668. https://doi.org/10.1016/j.trpro.2021.01.079.
- EPA (2024). Fast facts on transportation greenhouse gas emissions. Retrieved from https://www.epa.gov/greenvehicles/fast-facts-transportation-greenhouse-gas-emissions. Last accessed: 23/12/2024.
- European Parliament (2021). Transport infrastructure in low-density and depopulating areas. Retrieved from https://www.europarl.europa.eu/Reg-Data/etudes/STUD/2021/652227/IPOL\_STU(2021)652227\_EN.pdf. Last accessed: 22/12/2024.
- 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: 13/11/2024.
- Fleckenstein, D., Klein, R., Klein, V., & Steinhardt, C. (2024). On the concept of opportunity cost in integrated demand management and vehicle routing. *Transportation Science*, *online first*. https://doi.org/10.1287/trsc.2024.0644.
- Fleckenstein, D., Klein, R., & Steinhardt, C. (2023). Recent advances in integrating demand management and vehicle routing: A methodological review. *European Journal of Operational Research*, 306(2), 499-518. https://doi.org/10.1016/j.ejor.2022.04.032.
- FLEXIBUS (2024). Der FLEXIBUS in Krumbach. Retrieved from https://www.flexibus.net/krumbach. Last accessed: 22/12/2024.
- Gaul, D., Klamroth, K., & Stiglmayr, M. (2022). Event-based MILP models for ridepooling applications. *European Journal of Operational Research*, 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. *Transportation Research Procedia*, 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. *Ecological Economics*, 194, 107335. https://doi.org/10.1016/j.ecolecon.2021.107335.
- Haferkamp, J., & Ehmke, J. F. (2020). An efficient insertion heuristic for on-demand ridesharing services. *Transportation Research Procedia*, 47, 107-114. https://doi.org/10.1016/j.trpro.2020.03.074.
- Haferkamp, J., & Ehmke, J. F. (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, J. R. (2014). Consideration-set heuristics. *Journal of Business Research*, 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 Spectrum*, 46, 1099-1161. https://doi.org/10.1007/s00291-024-00752-4.
- Heitmann, R. J. O., Soeffker, N., Ulmer, M. W., & Mattfeld, D. C. (2023). Combining value function approximation and multiple scenario approach for the effective management of ridehailing services. *EURO Journal on Transportation and Logistics*, 12, 100104. https://doi.org/10.1016/j.ejtl.2023.100104.

- Horn, M. E. (2002). Fleet scheduling and dispatching for demand-responsive passenger services. *Transportation Research Part C: Emerging Technologies*, 10(1), 35-63. 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. *Transportation Research Part B: Methodological*, 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 diala-ride problem implemented for an Austrian mobility provider. *Transportation Research Procedia*, 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. *Research in Transportation Economics*, 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.
- Jaw, J. J., Odoni, A. R., Psaraftis, H. N., & Wilson, N. H. (1986). A heuristic algorithm for the multi-vehicle advance request dial-a-ride problem with time windows. *Transportation Research Part B: Methodological*, 20(3), 243-257. https://doi.org/10.1016/0191-2615(86)90020-2.
- Jung, J., Jayakrishnan, R., & Choi, K. (2012). Shared-taxi operations with electric vehicles. Working paper, University of California.
- Jung, J., Jayakrishnan, R., & Park, J. Y. (2016). Dynamic shared-taxi dispatch algorithm with hybrid-simulated annealing. *Computer-Aided Civil and Infrastructure Engineering*, 31(4), 275-291. https://doi.org/10.1111/mice.12157.
- Klein, R., Koch, S., Steinhardt, C., & Strauss, A. K. (2020). A review of revenue management: Recent generalizations and advances in industry applications. *European Journal of Operational Research*, 284(2), 397-412. https://doi.org/10.1016/j.ejor.2019.06.034.
- Klein, V., & Steinhardt, C. (2023). Dynamic demand management and online tour planning for same-day delivery. *European Journal of Operational Research*, 307(2), 860-886. https://doi.org/10.1016/j.ejor.2022.09.011.
- Koch, S., & Klein, R. (2020). Route-based approximate dynamic programming for dynamic pricing in attended home delivery. *European Journal of Operational Research*, 287(2), 633-652. https://doi.org/10.1016/j.ejor.2020.04.002.
- Köhler, C., Campbell, A. M., & Ehmke, J. F. (2024). Data-driven customer acceptance for attended home delivery. OR Spectrum, 46(2), 295-330. https://doi.org/10.1007/s00291-023-00712-4.
- Lotfi, S., & Abdelghany, K. (2022). Ride matching and vehicle routing for on-demand mobility services. *Journal of 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. *Transportation Research Part A: Policy* and Practice, 178, 103837. https://doi.org/10.1016/j.tra.2023.103837.
- Ma, Z., & Koutsopoulos, H. N. (2022). Near-on-demand mobility. The benefits of user flexibility for ride-pooling services. *Transportation Research Part C: Emerging Technologies*, 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. Computers & Operations Research, 147, 105941. https://doi.org/10.1016/j.cor.2022.105941.

- Mortazavi, A., Ghasri, M., & Ray, T. (2024). Integrated Demand Responsive transport in Low-Demand Areas: A case study of Canberra, Australia. *Transportation Research Part D: Transport and Environment*, 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: 23/12/2024.
- OSRM (n.d.). Open source routing machine. https://project-osrm.org/. Last accessed: 23/12/2024.
- Padam Mobility (2024). Our references. Retrieved from https://www.padam-mobility.com/en/success-stories. Last accessed: 22/10/2024.
- Poltimäe, H., Rehema, M., Raun, J., & Poom, A. (2022). In search of sustainable and inclusive mobility solutions for rural areas. *European Transport Research Review*, 14(1), 13. https://doi.org/10.1186/s12544-022-00536-3.
- Powell, W. B. (2019). A unified framework for stochastic optimization. European Journal of Operational Research, 275(3), 795-821. https://doi.org/10.1016/j.ejor.2018.07.014.
- Powell, W. B. (2022). Reinforcement learning and stochastic optimization: A unified framework for sequential decisions. Hoboken, NJ: John Wiley & Sons. https://doi.org/10.1002/9781119815068.
- Prud'homme, J., Josselin, D., & Aryal, J. (2011). Quantitative analysis of pollutant emissions in the context of demand responsive transport. In *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, M. L. (2014). Markov decision processes: Discrete stochastic dynamic programming. New York, NY: John Wiley & Sons. https://doi.org/10.1002/9780470316887.
- Qiu, H., Li, R., & Zhao, J. (2018). Dynamic pricing in shared mobility on demand service. Working paper, Massachusetts Institute of Technology.
- Schasché, S. E., Sposato, R. G., & Hampl, N. (2022). The dilemma of demand-responsive transport services in rural areas: Conflicting expectations and weak user acceptance. *Transport Policy*, 126, 43-54. https://doi.org/10.1016/j.tranpol.2022.06.015.
- Schilde, M., Doerner, K. F., & Hartl, R. F. (2011). Metaheuristics for the dynamic stochastic diala-ride problem with expected return transports. *Computers & Operations Research*, 38(12), 1719-1730. https://doi.org/10.1016/j.cor.2011.02.006.
- Sharif Azadeh, S., Atasoy, B., Ben-Akiva, M. E., Bierlaire, M., & Maknoon, M. Y. (2022). Choice-driven dial-a-ride problem for demand responsive mobility service. *Transportation Research Part B: Methodological*, 161, 128-149. https://doi.org/10.1016/j.trb.2022.04.008.
- Sieber, L., Ruch, C., Hörl, S., Axhausen, K. W., & Frazzoli, E. (2020). Improved public transportation in rural areas with self-driving cars: A study on the operation of Swiss train lines. *Transportation Research Part A: Policy and Practice*, 134, 35-51. https://doi.org/10.1016/j.tra.2020.01.020.
- Sörensen, L., Bossert, A., Jokinen, J. P., & Schlüter, J. (2021). How much flexibility does rural public transport need? – Implications from a fully flexible DRT system. *Transport Policy*, 100, 5-20. https://doi.org/10.1016/j.tranpol.2020.09.005.
- Strauss, A. K., Klein, R., & Steinhardt, C. (2018). A review of choice-based revenue management: Theory and methods. *European Journal of Operational Research*, 271(2), 375-387. https://doi.org/10.1016/j.ejor.2018.01.011.
- Talluri, K. T., & van Ryzin, G. (2004). Revenue management under a general discrete choice model of consumer behavior. *Management Science*, 50(1), 15-33. https://doi.org/10.1287/mnsc.1030.0147.
- Ulmer, M. W., Goodson, J. C., Mattfeld, D. C., & Thomas, B. W. (2020). On modeling stochastic dynamic vehicle routing problems. *EURO Journal on Transportation and Logistics*, 9(2), 100008. https://doi.org/10.1016/j.ejtl.2020.100008.

- Vazifeh, M. M., Santi, P., Resta, G., Strogatz, S. H., & 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, Germany. *Procedia Computer Science*, 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. *Case Studies on Transport Policy*, 3(2), 129-136. https://doi.org/10.1016/j.cstp.2014.12.006.
- Yang, X., Strauss, A. K., Currie, C. S. M., & Eglese, R. (2016). Choice-based demand management and vehicle routing in e-fulfillment. *Transportation Science*, 50(2), 473-488. https://doi.org/10.1287/trsc.2014.0549.
- Zwick, F., Kuehnel, N., Moeckel, R., & Axhausen, K. W. (2021). Ride-pooling efficiency in large, medium-sized and small towns-simulation assessment in the Munich metropolitan region. *Procedia Computer Science*, 184, 662-667. https://doi.org/10.1016/j.procs.2021.03.083.

# **Appendix A: Notation**

Table 4 Notation Markov decision process model

$t\in \mathcal{T}=\{1,\ldots,t^s,\ldots,T\}$	Decision epoch
$t^s$	Start of the service horizon
$c \in C$	Customer request type
$\lambda_c^t$	Arrival rate of request type $c$ in stage $t$
$p_c, d_c \in \mathcal{H}$	Pick-up (drop-off) stop of request type <i>c</i>
m <sub>c</sub>	Number of passengers of request type <i>c</i>
$t_c \in \{t^s, \dots, T\}$	Desired time of request type <i>c</i>
$f_c \in \{0,1\}$	Indicator for time window type of request type <i>c</i>
r <sub>c</sub>	Revenue of request type <i>c</i>
$i \in \mathcal{I}$	Request
$\tau_i \in \mathcal{T}$	Time of request for request <i>i</i>
$o\in \mathcal{O}_c$	Fulfillment option defined for request type <i>c</i>
$(\tau_{c,o}^{e+}, \tau_{c,o}^{l+}), (\tau_{c,o}^{e-}, \tau_{c,o}^{l-})$	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 $v$
$t^b_{v}$ , $t^r_{v}$	Start (end) of operations for vehicle $v$
$\phi_t$	Route plan at decision epoch t
$\theta_{v,t} \in \phi_t$	Planned route of vehicle $v$ according to route plan $\phi_t$
$a_{j_n}^-$ , $a_{j_n}^+$	Vehicle arrival (departure) time at the <i>n</i> -th stop in a route
s <sub>t</sub>	Post-decision state at decision epoch t
C <sub>t</sub>	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 $o$
$\left(\phi_t(o)\right)_{o\in g_t}\in$	Vehicle routing decision at decision epoch $t$ for request type $c$ availability control de-
$\prod_{o \in g_t} \Phi(s_{t-1}, c, o)$	cision $g_t$
$O_c^f$	Set of feasible fulfillment options for request type <i>c</i>
st	Pre-decision state at decision epoch t
o <sub>jt</sub>	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 $\phi_t(o)$
$ ho_{h,h'}$	Routing cost for traveling from stop $h$ to stop $h'$

$V_t(s_t)$	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 <i>l</i>
ζ	Ranking function over fulfillment options
o <sub>cgl</sub>	Fulfillment option chosen by request type $c$ and segment $l$ from offer set $g$
γı	Share of segment <i>l</i> in the customer population

#### Table 5 Notation of Markov decision process model (continued)

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
$\left( au_{k}^{e}, au_{k}^{l} ight)$	Time window of node <i>k</i>
k <sub>v</sub>	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
$\phi_t^{act}$	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
$\phi_t^{sam}$	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 $o$
<i>AR<sup>sam</sup></i>	Sampling acceptance rate
$\Delta \widetilde{V}_t(s_{t-1},c,o)$	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 $j$
<i>j</i> *	Sampled order with the highest cost saving
$\mathcal{O}^{pro}_{c_{i_t}}$	Subset of profit-maximizing fulfillment options
$\mathcal{O}^{clo}_{c_{i_t}}$	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 Section 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  $\phi_{t-1}$ , 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 dropoff 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  $(k, k') \in \mathcal{A} = \{(k, k'): k = 0, k' \in \mathcal{P} \lor (k, k' \in \mathcal{P} \cup \mathcal{D} \land k \neq k' \land k \neq k' + |\mathcal{J}|) \lor k \in \mathcal{D}, k' = 2|\mathcal{J}| + 1\}$ . The time window of each node k is defined by two time points marking its start  $\tau_k^e$  and its end  $\tau_k^l$ . It is equal to the order's pick-up time window defined by the earliest pick-up time  $\tau_k^{e+}$  and the latest pick-up time  $\tau_k^{l+}$  in case of  $k \in \mathcal{P}$  or the drop-off time window defined by the earliest drop-off time  $\tau_k^{e-}$  and the latest drop-off time  $\tau_k^{l-}$  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  $\phi_{t-1}$  such that only indices of orders remain that have not yet been picked up:

$$\mathcal{P} = C_{t-1} \setminus \left\{ k \in C_{t-1} : a_k^+ \le t, (k, h_k, a_k^-, a_k^+) \in \theta_{\nu, t-1}, \theta_{\nu, t-1} \in \phi_{t-1} \right\}$$
(11)

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

$$\mathcal{D} = \{j + |\mathcal{J}| : j \in \mathcal{C}_{t-1}\}$$
(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  $\phi_{t-1}$ :

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

$$q_{k_{\nu}} \coloneqq q_{k_{\nu}} + \sum_{\{k \in C_{t-1}: a_{k}^{+} \le t, (k, h_{k}, a_{k}^{-}, a_{k}^{+}) \in \theta_{t-1, \nu}\}} m_{c_{k}}$$
(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.

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

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

$$\sum_{k'\in\mathcal{N}} x_{k_{\nu},k'\nu} = 1 \qquad \forall \nu \in \mathcal{V}$$
(16)

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

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

$$B_{k'\nu} \ge (B_{k\nu} + \delta_{kk'}) x_{kk'\nu} \qquad \forall k \in \mathcal{N}, k' \in \mathcal{N}, \nu \in \mathcal{V}$$

$$(19)$$

$$Q_{k'\nu} \ge (Q_{k\nu} + q_{k'})x_{kk'\nu} \qquad \forall k \in \mathcal{N}, k' \in \mathcal{N}, \nu \in \mathcal{V}$$
(20)

$$B_{|\mathcal{J}|+k,\nu} - B_{k,\nu} \ge \delta_{k,|\mathcal{J}|+k} \qquad \forall k \in \mathcal{P}, \nu \in \mathcal{V}$$
(21)

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

$$a_{\nu} \le B_{k_{\nu},\nu} \qquad \qquad \nu \in \mathcal{V} \tag{23}$$

$$\max\{0, q_k\} \le Q_{kv} \le \min\{\kappa_v, \kappa_v + q_k\} \qquad \forall k \in \mathcal{N}, v \in \mathcal{V}$$

$$x_{kk'v} \in \{0, 1\} \qquad \forall k \in \mathcal{N}, k' \in \mathcal{N}, v \in \mathcal{V}$$
(24)
(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 dropoff 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  $(\tau_{c,o}^{e}, \tau_{c,o}^{l+})$  and drop-off  $(\tau_{c,o}^{e-}, \tau_{c,o}^{l-})$  are computed based on the desired time  $t_c$  (Jaw et al., 1986). To guarantee a certain service level for all types of requests  $c \in C$ , we define a maximum added time  $t_c^{mart}$  to the direct ride time  $\delta_{p_c,d_c}$  from the pick-up stop  $p_c$  to the drop-off stop  $d_c$ . It consists of a constant waiting time  $\omega$  and a certain fraction  $\mu$  of the direct ride time. Thus,  $t_c^{mart} = \omega + (1 + \mu)\delta_{j,j+|\mathcal{J}|}$ . Further, each node is assigned a time window  $(\tau_k^e, \tau_k^l)$  resulting from  $\delta_{j,j+|\mathcal{J}|}, \omega, \mu$ , and the trip type  $f_{c_j}$ . 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.

 Table 8 Computation of time windows

	$f_c = 0$ (outbound trip)	$f_c = 1$ (inbound trip)
$k \in \mathcal{P}$ (nick-up node)	$\tau_{c,o}^{e+} = o - \delta_{j,j+ \mathcal{J} } - t_c^{mart}$	$\tau^{e+}_{c,o}=o$
k C 5 (pick-up hode)	$\tau_{c,o}^{l+} = o - \delta_{j,j+ \mathcal{J} }$	$\tau^{l+}_{c,o}=o+\omega$
$k \in \mathcal{D}$ (drop off node)	$\tau^{e-}_{c,o}=o-\omega$	$\tau^{e-}_{c,o}=o+\delta_{j,j+ \mathcal{J} }$
$k \in D$ (drop-off field)	$ au_{c,o}^{l-}=o$	$\tau_{c,o}^{l-} = o + \delta_{j,j+ \mathcal{J} } + t_c^{mart}$

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

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  $\gamma_1$  and  $\gamma_2$ , 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  $(P_{1,0}(\{0,1\}) =$  $P_{1,2}(\{0,1\}) = P_{1,0}(\{0,1\}) = 0)$ . 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

**Input**: Current actual route plan  $\phi_t^{act}$ , set of feasible options  $\mathcal{O}_{c_{i_t}}^f$ , current request  $i_t$ , option o

- 1  $\phi_t^{act}(o) \coloneqq parallel_insertion(\phi_t^{act}, (i_t, o))$
- 2 if  $(i_t, o)$  in  $\phi_t^{act}(o)$  do

3 
$$\mathcal{O}_{c_{i_t}}^f \coloneqq \mathcal{O}_{c_{i_t}}^f \cup \{o\}$$

Fig. 16 Feasibility check function

## Appendix F: Availability Control

**Input**: Current customer request  $i_t$ , set of cost estimates  $\left\{ \Delta \widetilde{V}_t(s_{t-1}, c_{i_t}, o) : o \in \mathcal{O}_{c_{i_t}}^f \right\}$ 

$$1 \qquad \mathcal{O}_{c_{i_t}}^{pro} \coloneqq \left\{ o' \in \mathcal{O}_{c_{i_t}}^f : o' = \operatorname*{argmax}_{o \in \mathcal{O}_{c_{i_t}}^f} \left\{ r_{c_{i_t}} - \Delta \widetilde{V}_t(s_{t-1}, c_{i_t}, o) \right\}, r_{c_{i_t}} - \Delta \widetilde{V}_t(s_{t-1}, c_{i_t}, o') \ge 0 \right\}$$

$$2 \qquad \mathcal{O}_{c_{i_t}}^{clo} \coloneqq \left\{ o' \in \mathcal{O}_{c_{i_t}}^{pro} : o' = \operatorname*{argmin}_{o \in \mathcal{O}_{c_{i_t}}^{pro}} \left\{ \left| o - t_{c_{i_t}} \right| \right\} \right\}$$

$$3 \qquad \mathcal{O}_{j_t} \coloneqq \min_{o \in \mathcal{O}_{c_{i_t}}^{clo}} o$$

Fig. 17 Availability control function

## **Appendix G: Routing Control**

<b>Input</b> : Actual route plan including new order $\phi_t^{act}(o_{j_t})$ , set of actual orders $C_t^{act}$			
1	$\phi_t^{act} \coloneqq \phi_t^{act}(o_{j_t})$		
2	$\mathbf{if}t=t^s-1\mathbf{do}$		
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	$\phi_t^{act} \coloneqq parallel\_insertion(\phi_t^{act}, C_t^{act})$		

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  $\phi_t^{sam}$  is initialized as a copy of the actual route plan  $\phi_t^{act}$ . Then, the sampled orders remaining at the preceding decision epoch  $C_{t-1}^{sam}$  are inserted into  $\phi_t^{sam}$  (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).

<b>Input</b> : Current actual route plan $\phi_t^{act}$ , current set of sampled orders $C_{t-1}^{sam}$		
1	$\mathbf{if}t\geq t^{s}\mathbf{do}$	
2	$\phi_t^{sam}\coloneqq\phi_t^{act}$	
3	$\phi_t^{sam} \coloneqq parallel\_insertion(\phi_t^{sam}, C_{t-1}^{sam})$	
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

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 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).

In	<b>Input</b> : Current sampled route plan $\phi_t^{sam}$ , current request $i_t$ , option <i>o</i>		
	1	$\phi_t^{sam}(o) \coloneqq parallel\_insertion(\phi_t^{sam}, (i_t, o))$	
	2	$\mathbf{if}\boldsymbol{\phi}_t^{sam}(o) = \boldsymbol{\phi}_t^{sam}\mathbf{do}$	
	3	$\Delta \widetilde{V}_t(s_{t-1}, c_{i_t}, o) \coloneqq \infty$	
	4	else do	
	5	$\Delta \widetilde{V}_t(s_{t-1}, c_{i_t}, o) \coloneqq cost(\phi_t^{sam}(o)) - cost(\phi_t^{sam})$	

Fig. 20 Opportunity cost estimation function (anticipatory control)

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.

<b>Input</b> : Sampled route plan including new order $\phi_t^{sam}(o_{j_t})$ , set of sampled orders $C_t^{sam}$		
1	$\mathbf{if} \ o_{j_t} \neq 0 \ \mathbf{do}$	
2	$\phi_t^{sam}\coloneqq\phi_t^{sam}(o_{j_t})$	
3	forall $j \in C_t^{sam}$ do	
4	$\phi_{t,j}^{sam} \coloneqq remove(\phi_t^{sam}, j)$	
5	$j^* \coloneqq \operatorname*{argmin}_{j \in \mathcal{C}_t^{sam}} \{ cost(\phi_{t,j}^{sam}) \}$	
6	$\phi_t^{sam}\coloneqq\phi_{t,j^*}^{sam}$	

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

# Appendix I: Service Area Krumbach Divided in Zones



Fig. 22 Service area of Krumbach divided in zones (FLEXIBUS, 2024)

## **Appendix J: Pricing Scheme FLEXIBUS**

Table 9 Pricing scheme FLEXIBUS (FLEXIBUS, 2024)



# Appendix K: Control Behavior of Non-Selective and Non-Time-Shifting Policies



**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 (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.

# Article A6: Sustainable Dynamic Pricing for Rural Shared Mobility-on-Demand Systems

Fabian Anzenhofer<sup>a</sup>, David Fleckenstein<sup>a</sup>, Robert Klein<sup>a</sup>, Claudius Steinhardt<sup>b</sup> (<sup>a</sup>University of Augsburg, <sup>b</sup>University of the Bundeswehr Munich) Working paper

#### Abstract

Shared mobility-on-demand (SMOD) systems allow customers to request customized rides that the provider bundles through ridepooling. Owing to their flexibility and efficiency in low-demand settings, SMOD systems are expected to become a cornerstone of rural public transport. Operationally, dynamic pricing can influence which rides customers order, thereby enhancing system performance. In the literature, dynamic pricing is exclusively applied to maximize profit. However, rural SMOD services primarily provide basic mobility as a public service and should contribute to reducing the greenhouse gas emissions of passenger transportation. With this work, we present the first multi-objective, sustainable dynamic pricing approach that considers all three dimensions of sustainability, i.e., social, environmental, and economic objectives. The core idea is to set prices that maximize served demand while adhering to a dynamic lower price bound ensuring that the price of each ride at least covers its marginal cost. Thereby, our approach carefully balances the partially conflicting sustainability objectives. To accurately evaluate demand displacement and marginal cost of rides, which is crucial for making dynamic pricing decisions, we propose a post-decision rollout algorithm that anticipates the future evolution of the booking process. We comprehensively evaluate our approach based on a real-world data set of a rural SMOD provider and derive managerial insights regarding the practical impact of sustainable dynamic pricing. Our findings show that sustainable dynamic pricing achieves substantially better performance than static pricing while preventing monopoly markups, which would result from profit-based dynamic pricing.

Key words: Mobility-on-Demand, Rural Areas, Sustainability, Dynamic Pricing, Routing

# 1 Introduction

Shared mobility-on-demand (SMOD) services, also known as demand-responsive transport services, are a flexible form of public transport that allows customers to request customized rides via mobile applications. Providers can bundle the rides of multiple, unrelated customers and transport them in one vehicle, which is called ridepooling. In rural areas, the main advantage of SMOD systems is that they can efficiently provide area-wide mobility coverage due to their flex-ibility paired with at least a low degree of demand consolidation (Mounce et al., 2020). They are expected to be part of future rural public transport systems by complementing or replacing existing scheduled public transport services in low-demand areas (Mortazavi et al., 2024, Sieber et al., 2020).

Due to low and dispersed demand, rural public transport is often financially unprofitable, and SMOD systems are no exception (e.g., Zwick et al., 2022). However, there is a broad socio-political consensus across many countries that governments must guarantee an adequate basic mobility provision as a public service, although state-level definitions of which specific objectives should be pursued to achieve "adequacy" are lacking to date (Mounce et al., 2020). Traditionally, adequacy has been seen as economic efficiency, but an increasingly popular perspective is to equate adequacy with sustainability (e.g., Poltimäe et al., 2022 or Schwedes, 2021, p. 45).

The most popular definition of sustainability distinguishes three major objectives: social, environmental, and economic sustainability (Purvis et al., 2019). We consider operational planning of an SMOD service to be sustainable if it pursues the following lower-level objectives derived from the three major objectives:

- **Social sustainability**: The SMOD system should provide a) area-wide basic mobility that should be b) adequately priced because it is offered as a public service (Lu et al., 2024).
- Environmental sustainability: The SMOD system should contribute to emission savings by a) promoting modal shift from motorized individual transport to public transport (Burghard and Scherrer, 2022) and b) by decreasing emissions per passenger km (Lotze et al., 2023).
- Economic sustainability: The economic perspective is also relevant because subsidy requirements must be minimized to ensure long-term viability beyond the pilot stage (De Jong et al., 2011).

Responsible for planning SMOD services are municipal authorities, who either act as the service provider themselves or commission a private company (Lu et al., 2024). For the latter, performance-based contracts are becoming increasingly popular to ensure that the private provider operates the system in line with the authority's requirements (Hensher, 2020). In any case, municipal authorities aim at social, environmental, and economic sustainability objectives that are partially conflicting. Thus, the provider requires multi-objective optimization methods allowing to find a suitable trade-off. Please note that we consistently refer to the sustainability objectives as the provider's objectives in the remainder of the paper.
In this work, we focus on the operational planning level of SMOD systems. Given the strategic and tactical decisions on the supply of logistical capacity (e.g., shift planning), the provider can actively manage demand at the operational planning level to improve performance (Anzenhofer et al., 2024, Haferkamp and Ehmke, 2022).

We investigate demand management through dynamic pricing, which means that the provider determines an individual price for each ride based on the current system state. Dynamic pricing for (rural) SMOD systems is not an entirely new approach (e.g., Arian et al., 2022, Sharif Azadeh et al., 2022), but, following the traditional economic-driven planning paradigm, academic literature proposes to apply dynamic pricing for profit maximization (Hörcher and Graham, 2020a). While the objective of profit maximization is to some degree congruent with the social and environmental sustainability objectives, there is a fundamental issue (Bahamonde-Birke et al., 2021): If the primary objective is profit-oriented, the provider will raise prices solely to exploit the customers' expected willingness-to-pay, which is known as charging a *monopoly markup*. As a result, the provider can capitalize on the consumer surplus, which comes at the cost of lost demand. In view of the social and environmental sustainability objectives).

As an alternative to profit-based dynamic pricing, we propose a sustainable dynamic pricing approach that draws on multi-objective optimization to balance all three dimensions of sustainability. We begin by using multi-attribute decision analysis (Keeney and Raiffa, 1993) to structure the provider's objectives derived from the all-inclusive objective of improving sustainability. Then, we translate these results into a multi-objective optimization model, specifically a constrained Markov decision process (MDP) model. Its objective function reflects the primary objectives, while the secondary objectives are represented as constraints:

- **Primary objective**: Our primary objective for pricing optimization is to maximize the demand served by the SMOD system, supporting social sustainability (area-wide basic mobility) and environmental sustainability (modal shift). As the primary objective is not profitoriented, we prevent monopoly markups and ensure price adequacy. However, determining the specific contribution of a ride to this objective is particularly challenging because serving the ride consumes logistical capacity, potentially causing rides that are requested later to become infeasible, which results in displaced demand.
- Secondary objectives: Maximizing demand served alone is insufficient to reflect the environmental and economic sustainability objectives, which we consider secondary objectives. To account for them, we draw on *marginal cost pricing*, a widely established concept in public transport planning (e.g., Hörcher and Tirachini, 2021), where the price is based on the marginal societal cost of a ride. In the case of rural SMOD systems, this cost includes both the internalized variable routing cost for the provider and external cost due to emissions. By setting a dynamic lower price bound equal to a ride's expected marginal cost, we address both secondary objectives, as unfavorable rides are priced higher. The price signal allows

customers to decide whether to book the desired ride, choose an alternative ride, opt for another mode of transport, or cancel the ride (Eliasson, 2021).

Due to the dynamic and stochastic nature of the sustainable dynamic pricing problem, two key features of rides, displaced demand and marginal cost, cannot be exactly determined. For the sake of brevity, we use the term *ride evaluation* for the task of approximating them. We propose an anticipatory post-decision rollout algorithm (Bertsekas et al., 1997) for this task. Based on simulation, it compares the future evolution of the booking process from the current state onward for two cases: a) the current customer orders the ride and b) the current customer abandons the booking process. With anticipatory ride evaluation, we can attribute marginal cost accurately to individual rides. Thereby, we solve a central issue that, up to now, hindered the application of marginal cost pricing in the context of dynamic pricing for SMOD (Andrejszki and Török, 2018).

Overall, our work has the following scientific contributions:

- To our knowledge, we are the first to propose a sustainable dynamic pricing approach for rural SMOD systems. To this end, we formulate a multi-objective optimization model that is derived from an extensive multi-attribute decision analysis and incorporates marginal cost pricing.
- We present an anticipatory post-decision rollout algorithm to approximate displaced demand and marginal costs. It enables the practical application of marginal cost pricing for SMOD services, which has been an issue unsolved by existing research.
- To validate our approach, we apply it to a one-year, real-world data set provided by our industry partner FLEXIBUS who operate one of the most mature rural SMOD systems in Germany. By conducting an extensive computational study including sensitivity analyses, we explain the performance differences compared to benchmark approaches. Further, we derive managerial insights and recommendations for practice.

The remainder of this work is organized as follows: In Section 2, we give an overview of the related literature. Section 3 comprises an in-depth description of the considered dynamic pricing problem and the underlying assumptions. Further, we show how the entire operational planning problem can be formalized as a constrained MDP. In Section 4, we introduce the solution algorithm. In Section 5, we present the design and the results of our computational study. Finally, in Section 6, we discuss managerial insights, recommendations, and future research opportunities resulting from our work.

## 2 Literature Review

Related literature to the work at hand originates in two distinct disciplines. On the one hand, there is a body of literature from (public) transport economics on differentiated pricing, reviewed in Section 2.1. On the other hand, there are many operations research works on integrated demand management and vehicle routing, including those that consider dynamic pricing (Section 2.2). Finally, we summarize the research gap in Section 2.3.

### 2.1 Differentiated Pricing in Public Transport

In this section, we review the literature on dynamic pricing from the perspective of transport economics. Publications in this field typically use stylized models or simulations of multi-modal transport systems to analyze fundamental pricing concepts and derive policy implications. While our focus is narrower, directed at the provider's operational planning, this research provides the theoretical foundation for our work. For a deeper discussion, we refer to surveys by Hörcher and Tirachini (2021), Saharan et al. (2020), and Vickerman (2024).

Price differentiation, the basic principle behind dynamic pricing, is analyzed across different transport modes. Its justification is twofold: First, demand can be actively steered to align better with fixed supply of transport capacity. Second, it allows the internalization of external effects based on the marginal cost of transport demand. Bimpikis et al. (2019) and Ma et al. (2022) examine ride-hailing platforms. While the former emphasize the profit benefits of spatio-temporal demand balancing, the latter apply price differentiation for welfare maximization. Kaddoura et al. (2015) explore customer-specific price differentiation in bus networks. Outside of public transport, Eliasson (2021) considers congestion pricing, providing general insights into the theoretical justification, the impact, and the practical implementation of marginal cost pricing.

Some works consider the transfer of marginal cost pricing to mobility-as-a-service systems. Hörcher and Graham (2020a) show that the widely aspired subscription models are economically inefficient compared to differentiated pricing, which confirms the potential of marginal cost pricing for SMOD services. However, theoretical analysis by Hörcher and Graham (2020b) highlights the critical role of the provider's objectives and identifies monopoly markups as a major challenge for the application of dynamic pricing. Bahamonde-Birke et al. (2021) further show that an unregulated, profit-maximizing provider inevitably applies aggressive price differentiation, suggesting regulatory enforcement of marginal cost pricing as a potential solution.

In addition to the generic analysis, two implementations of marginal cost pricing are proposed: First, in urban areas, congestion pricing is applied through price increases for rides traversing congested road segments (Kaddoura et al., 2020a) or subsidies for less congested ones (Ke and Qian, 2023). However, this implementation is not suitable for rural areas with negligible congestion (Heinitz, 2022). Second, the concept of ex-post pricing is proposed (Andrejszki and Török, 2018, Karaenke et al., 2023). Following this concept, each customer is initially charged the marginal cost of a fictive solitary ride. After service fulfillment, the customer is then (partly) reimbursed based on a cost-sharing rule if the ride could be successfully pooled. While this implementation ensures marginal cost coverage, it does not actively manage demand because customers cannot adapt their ride choice according to the price signal when placing their order.

#### 2.2 Operational Demand Management for SMOD Services and Related Services

The operational decision problem considered in the work at hand belongs to the family of integrated demand management and vehicle routing problems (i-DMVRPs). To provide an overview of the related literature, we first list corresponding surveys (Section 2.2.1). Second, we extensively review the stream of i-DMVRP literature that considers (S)MOD systems, distinguishing it from less closely related work in demand-responsive passenger transportation (Section 2.2.2). Third, we review i-DMVRP literature and dynamic vehicle routing literature that considers other application areas but shares communalities with our work (Section 2.2.3).

# 2.2.1 Survey Literature

According to Fleckenstein et al. (2024b), an i-DMVRP arises if a logistical service provider both plans customized offers to customers requesting service dynamically and solves a vehicle routing problem to plan order fulfillment. A cross-application survey of existing i-DMVRP literature is provided by Fleckenstein et al. (2023). i-DMVRPs integrate two optimization problems: revenue management/dynamic pricing, and dynamic vehicle routing. For the former problems, recent surveys by Klein et al. (2020) and Strauss et al. (2018), as well as textbooks by Talluri and van Ryzin (2004) and Gallego and Topaloglu (2019) provide thorough coverage. Dynamic vehicle routing is surveyed by Pillac et al. (2013), Psaraftis et al. (2016), Rios et al. (2021), and Soeffker et al. (2022). Finally, for application-specific insights on operational planning for (S)MOD systems, readers can draw on the surveys by Rammohan et al. (2024), Vansteenwegen et al. (2022), and Zwick et al. (2022).

# 2.2.2 Demand Management for SMOD Services

In the following, we review the literature most closely related to our work. To suitably narrow the scope, we first delineate this stream from broader literature that also investigates flexible passenger transportation. More precisely, we exclude literature with any of the following characteristics:

- A large body of research considers the static, deterministic dial-a-ride problem. Although there are major differences between the dynamic and the static perspective, it is important to note that there are some related works applying multi-objective optimization (see Ho et al. (2018) for a detailed discussion), including the formulation of constraints representing secondary objectives (e.g., Anzenhofer et al., 2025).
- In pure dynamic vehicle routing, SMOD systems are frequently considered. Without demand management, requests are collected on a first-come-first-served basis (e.g., Hungerländer et al., 2021).
- Early works on SMOD often focus on services for specific target groups (e.g., Schilde et al., 2011).
- Ride-hailing literature investigates taxi-like services without ridepooling, often processing requests in batches (e.g., Alonso-Mora et al., 2017). We only consider such works if their focus is on dynamic pricing.

- There is a body of literature on dynamic pricing for managing both demand and supply in two-sided markets (e.g., Wang and Yang, 2019). This setting occurs if the provider does not operate the fleet but offers a platform matching self-employed drivers and customers.
- Finally, there is some literature on pricing for classical ride-sharing systems (e.g., Zhang et al., 2020).

Next, we compare existing publications on demand management for SMOD systems with our work, grouping the literature into three categories: rural SMOD systems, urban SMOD systems, and urban ride-hailing systems. Within each category, we focus on *problem definition*, *modeling approach*, and *solution approaches* for the dynamic pricing subproblem and for *ride evaluation*. An overview of the results is given in Table 1.

**Rural SMOD systems**: Our work shares the most similarities with Arian et al. (2022), who are the first to investigate dynamic pricing for rural SMOD systems, albeit with a much different *problem definition*. First, they focus solely on profit maximization. Second, they assume that customers place only ad-hoc requests, whereas we also account for advance requests. Further differences are in the *solution method*: For solving the dynamic pricing subproblem, they propose an efficient algorithm relying on a specific customer choice model, namely the multinomial logit model. For *ride evaluation*, they use value function approximation.

The only other work on demand management in the rural context is by Anzenhofer et al. (2024). Compared to our work, their *problem definition* mainly differs in its focus on profit maximization and the use of availability control for demand management. However, we draw on their *modeling approach* as we adapt their MDP model. Extending prior work by Haferkamp and Ehmke (2022), who also focus on the impact of vehicle routing, their main aim is to support the strategic decision on which demand control policy to apply. Despite their strategic focus, they present a practical, sampling-based lookahead algorithm suitable for operational *ride evaluation*, which is closest to ours in the literature. Its main shortcoming is its inability to capture demand displacement, which is why we introduce a post-decision rollout algorithm as a more sophisticated sampling-based approach.

**Urban SMOD systems**: In the urban setting, Qiu et al. (2018) and Sharif Azadeh et al. (2022) consider dynamic pricing problems that closely resemble our *problem definition*, aside from their focus on profit maximization. Both studies consider a setting with combined ride-hailing and ride-pooling. Sharif Azadeh et al. (2022) also account for variably sized pick-up time windows. Regarding the *solution approach*, Qiu et al. (2018) solve the dynamic pricing problem with a dedicated algorithm. Even though they employ a learning-based concept for *ride evaluation*, there are parallels to our sampling-based approach. They also use a post-decision rollout algorithm to learn the parameters of a linear function approximating demand displacement and marginal cost. However, a key difference is that they consider vehicle routing only in an aggregated way, which is not suitable for rural settings with much lower, dispersed demand. Sharif Azadeh et al. (2022)

solve a linear program to determine prices and use myopic *ride evaluation*, which does not capture demand displacement.

Other authors explore availability control with multiple alternative rides in urban areas: Atasoy et al. (2015) propose an assortment optimization model for the availability control problem. We use the same *modeling approach*. However, key differences include their focus on profit maximization, a different ride definition, and a myopic solution concept. Haferkamp (2024) considers a similar *problem definition* to ours, one of the few that does not aim at profit maximization, maximizing the number of orders instead. For *ride evaluation*, a cost function approximation is used that does not yield an explicit approximation of displaced demand and marginal cost.

In the urban context, we find the only two publications on *multi-objective optimization* in the dynamic setting. Lu et al. (2024) aim at maximizing the number of orders without a good alternative mode of transport and at minimizing the total travel time. To account for the former objective, they formulate a constraint which enforces the rejection of requests that can travel faster by an alternative mode compared to the SMOD service. Wu et al. (2024) model two objectives, namely the minimization of cost and unfairness regarding waiting times. Besides routing cost, there are penalties for request rejection and time window violation, which de facto results in profit maximization. As a *modeling approach*, they use scalarization to obtain a single-objective MDP model, and they apply multi-agent reinforcement learning as the *solution approach*. Despite the similarities, there are also decisive differences regarding our work: Both works consider only a simple accept/reject control without alternative ride options. Further, they do not account for all three sustainability dimensions.

Finally, some publications consider availability control without alternative ride options and a single objective (Heitmann et al., 2023, Heitmann et al., 2024, Jung et al., 2016, and Lotfi and Abdelghany, 2022).

**Urban ride-hailing systems**: A less closely related *problem definition* can be found in the context of urban ride-hailing systems (Al-Kanj et al., 2020, Chen et al., 2019, Haliem et al., 2021, Ni et al., 2021, Wang et al., 2021b). In ride-hailing, requests are often processed in batches since, without pooling, the vehicle routing problem can be formulated as a matching problem. The objective is exclusively revenue or profit maximization. Despite these differences, there are some similarities with our work. E.g., Chen et al. (2019) also combine pre-defined base prices and price multipliers. Al-Kanj et al. (2020), Haliem et al. (2021), and Ni et al. (2021) propose learning-based *ride evaluation* approaches that capture both demand displacement and marginal cost.

# 2.2.3 Related Applications

Finally, we turn toward literature from other applications using similar methodology. Each of these publications is only related to our work regarding the demand management subproblem or ride evaluation.

<b>Table 1</b> Literature overview; Abbreviations: Demand management: AR (accept/reject), AV(availability control), PR
(dynamic pricing); Objectives: NO (number of orders), RE (revenue), PR (profit), TT (travel time), FA (fairness);
Request processing: RT (real-time), BA (batched); Solution concept: M (myopic), S (anticipatory, sampling), L (antic-
ipatory, learning); Demand management subproblem: FE (full enumeration), DA (dedicated algorithm), LP (linear
programming), ML (machine learning)

Authors	Demand manage- ment	Objec- tives	Request pro- cessing	Pooling	Rural setting	Solution concept	Dis- place ment	Mar- ginal cost	Demand management subproblem
Al-Kanj et al. (2020)	PR	PR	BA	Х	Х	L	$\checkmark$	$\checkmark$	ML
Anzenhofer et al. (2024)	AV	PR	RT	$\checkmark$	$\checkmark$	(S)	х	$\checkmark$	FE
Arian et al. (2022)	PR	PR	RT	$\checkmark$	$\checkmark$	L	$\checkmark$	$\checkmark$	DA
Atasoy et al. (2015)	AV	PR	RT	$\checkmark$	Х	М	Х	$\checkmark$	LP
Chen et al. (2019)	PR	RE	RT	Х	Х	L	$\checkmark$	Х	ML
Haferkamp (2024)	AV	NO	RT	$\checkmark$	Х	М	Х	Х	FE
Haferkamp and Ehmke (2022)	AR	NO	RT	$\checkmark$	Х	(S)	$\checkmark$	Х	FE
Haliem et al. (2021)	PR	PR	BA	Х	Х	L	$\checkmark$	$\checkmark$	FE
Heitmann et al. (2023)	AR	RE	RT	$\checkmark$	Х	L	$\checkmark$	Х	FE
Heitmann et al. (2024)	AR	RE	RT	$\checkmark$	Х	L	$\checkmark$	Х	FE
Jung et al. (2016)	AR	PR/TT	RT	$\checkmark$	Х	М	Х	$\checkmark$	FE
Lotfi and Ab- delghany (2022)	AR	PR	RT	$\checkmark$	Х	М	Х	$\checkmark$	FE
Lu et al. (2024)	AR	NO/TT	RT	$\checkmark$	Х	М	Х	Х	FE
Ni et al. (2021)	PR	PR	BA	Х	Х	L	$\checkmark$	$\checkmark$	LP
Qiu et al. (2018)	PR	PR	RT	$\checkmark$	Х	L	$\checkmark$	$\checkmark$	DA
Sharif Azadeh et al. (2022)	PR	PR	RT	$\checkmark$	Х	М	Х	$\checkmark$	LP
Wang et al. (2021b)	PR	PR	BA	$\checkmark$	Х	М	Х	Х	DA
Wu et al. (2024)	AR	PR/FA	RT	$\checkmark$	Х	L	$\checkmark$	$\checkmark$	FE
Our work	PR	SO	RT	$\checkmark$	$\checkmark$	S	$\checkmark$	$\checkmark$	FE

**Demand management subproblem**: The most closely related non-SMOD demand management problem can be found in Lang et al. (2021). They consider multi-objective availability control in attended home delivery, accounting for the economic objectives of maximizing profit, social influence of customers, and visibility of delivery vehicles. For model development, they adopt a similar approach to ours in that they include each objective either in the objective function or as a hard constraint to suitably reflect the provider's preferences.

Dynamic pricing is also applied in attended home delivery (Koch and Klein, 2020) and same-day delivery (Klein and Steinhardt, 2023). The latter application is more strongly related to SMOD systems since, due to the low number of delivery options (the equivalent to rides in SMOD), the dynamic pricing subproblem becomes tractable by full enumeration. Particularly close parallels can be drawn between our approach and the one proposed by Ulmer (2020a) for same-day delivery. It also uses base prices for all fulfillment options and sets prices equal to displacement cost if they exceed the base price. This idea is very similar to the notion of marginal cost pricing that

our modeling approach is based on. While the dynamic pricing problems mentioned in this paragraph all aim at profit maximization, there is recent work in the generic dynamic pricing literature on incorporating fairness as a "non-profit" objective (Cohen et al., 2024). Finally, since we model the dynamic pricing subproblem as an assortment optimization problem (Davis et al., 2013), there is also a connection to assortment optimization literature (see Heger and Klein (2024) for a recent survey). From this generic view, Chen et al. (2024) investigate assortment optimization with multiple objectives. However, their modeling framework is not applicable to our problem since all objectives are represented in the objective function, which leads to monopoly markups.

**Ride evaluation**: Sustainable dynamic pricing decisions are optimized based on their expected future impact regarding displaced demand and marginal cost. In the profit-maximizing setting, this future impact is captured by opportunity cost. In model development, we exploit theoretical results by Fleckenstein et al. (2024a) who analytically show that opportunity cost in the context of profit-maximizing i-DMVRPs can be decomposed into displacement cost and marginal cost-to-serve. For sustainable dynamic pricing, displaced booked passenger km and marginal cost can be viewed as the direct equivalents of these opportunity cost components. For the ride evaluation task, we develop a post-decision rollout algorithm (Goodson et al., 2017). First developed by Bertsekas et al. (1997), such algorithms have not yet been applied to SMOD systems (see Table 1). However, there are existing applications for other i-DMVRPs (Ulmer et al., 2016, Ulmer et al., 2019, and Ulmer, 2020b).

# 2.3 Research Gap

To conclude the literature review, this section summarizes the research gap as identified:

- The transport economics literature provides strong evidence for the need to transfer marginal cost pricing to SMOD but highlights the challenge of preventing monopoly markups by profit-maximizing providers. Existing proposals of practical implementations are not suitable for rural areas.
- The operations research literature suggests sustainable dynamic pricing as a future research area (e.g., Qiu et al., 2018, Lang et al., 2021). To date, dynamic pricing is exclusively applied for profit maximization or closely related objectives (see Table 1 and Hörcher and Graham, 2020a). Further, except for two recent publications (Lu et al., 2024 and Wu et al., 2024), there is no work on multi-objective demand management for SMOD systems. Moreover, research on rural SMOD systems is scarce, despite their high potential.

Overall, our work is the first to address the research gap of developing a practical sustainable dynamic pricing approach for rural SMOD systems. Our contribution is comprehensive, including the definition of objectives based on multi-attribute decision analysis, the development of a multi-objective optimization model, the design of a solution algorithm, and an in-depth computational analysis based on real-world data.

# 3 Problem Definition and Modeling

In this section, we formalize the problem definition and generalize the MDP model by Anzenhofer et al. (2024). Section 3.1 outlines the problem definition by focusing on booking process, objectives, and the basic structure of the associated multi-objective optimization model. Section 3.2 formulates the model.

# 3.1 Problem Definition

This section describes the booking process, which handles both ad-hoc and advance requests (Section 3.1.1). It then outlines the objectives and attributes driving decision-making (Section 3.1.2) and concludes with the basic structure of the multi-objective optimization model (Section 3.1.3).

# 3.1.1 Booking Process

For a given service day, which we refer to as a service horizon, the provider receives customer requests dynamically via a smartphone app during the booking horizon. The booking horizon begins one or more days before the service horizon, allowing advance requests. Both horizons end simultaneously, permitting same-day requests (hours before the desired pick-up or drop-off) and ad-hoc requests (immediate pick-up).



Fig. 1 Booking process for an exemplary request

For a single request, the booking process consists of four steps, with two involving the customer and two involving the provider (see Fig. 1 for an exemplary overview).

• *Request arrival*: The customer specifies their origin, destination, the number of passengers, and either the desired pick-up time or drop-off time. In Fig. 1, we observe the request arrival on 3<sup>rd</sup> of November at 12:00 (time of request) six days in advance (advance request) of the desired pick-up time on 9<sup>th</sup> of November at 8:15. The customer also specifies the origin-destination (OD) pair from pick-up location A to drop-off location B for two passengers.

- Dynamic pricing decision: The provider determines fulfillment options, i.e., rides with alternative pick-up times or drop-off times and different prices, they offer to the customer. The provider may also opt to make no offer at all. This first decision is referred to as the dynamic pricing subproblem. Fig. 1 again exemplifies the provider's dynamic pricing decision. Here, the provider decides to offer three rides (7: 45, 8: 15, 8: 45), with different prices (4 €, 9 €, 5 €), including alternative pick-up times (+/-30 minutes). The prices indicate that the pick-up earlier than desired is the most favorable for the provider.
- *Order confirmation*: If an offer was made, the customer either chooses one of the offered rides, confirming the order, or abandons the booking process if unsatisfied. In the example, the customer confirms an order for the earliest pick-up time of 7:45, as highlighted in green in Fig. 1. Hence, the provider could successfully shift demand to a pick-up time 30 minutes earlier than the original desired time of 8:15.
- *Vehicle routing decision*: After an order is confirmed, the provider assigns the ride to a vehicle. This second decision is referred to as the vehicle routing subproblem. As shown in the example, the provider updates the current route plan by inserting the newly confirmed order (Fig. 1).

# 3.1.2 Structuring of Objectives and Specification of Attributes

Dynamic pricing enables providers to manage demand by influencing customer choices, thereby improving the system performance according to specific objectives. To prepare for building a multi-objective optimization model, we use the multi-attribute decision analysis framework by Keeney and Raiffa (1993) to define and structure objectives, ensuring alignment with the provider's preferences.

The all-inclusive objective of sustainably operating the SMOD service can be broken down into three major objectives: social sustainability, environmental sustainability, and social sustainability (Purvis et al., 2019). From these major objectives, we derive five lower-level objectives, each associated with a (proxy) attribute. Since we focus on operational planning, only objectives directly impacted by operational decisions are included. Below, we define and structure them into *primary* and *secondary objectives*, as summarized in Table 2.

**Social sustainability – a) Maximize basic mobility provision**: Municipal authorities are often legally required to provide mobility as a basic public service (Lu et al., 2024). Given a fixed fleet of vehicles, the objective is to maximize mobility provision, i.e., serving as much demand as possible (Hörcher and Tirachini, 2021). Since ensuring basic mobility provision is the very purpose of publicly funded rural SMOD systems, we consider maximizing it a *primary objective* (Dauer et al., 2024, Sörensen et al., 2021).

The attribute used to measure this objective is *total booked passenger km*. For an individual order, it equals the direct distance from pick-up stop to drop-off stop, multiplied by the number of

passengers (Zwick et al., 2022). Hence, the attribute excludes detours and captures only the mobility provision as originally demanded.

**Social sustainability** – **b) Prevent monopoly markup**: Price discrimination is widely seen as effective to enhance system performance (Section 2.1). However, a publicly funded provider is a natural monopolist. Hence, applying dynamic pricing with economic objectives leads to an undesirable monopoly markup (Hörcher and Graham, 2020a) since the provider excessively exploits the consumer surplus (Bahamonde-Birke et al., 2021).

To ensure social adequacy, the approach must be explicitly designed to prevent monopoly markups. Hence, this *primary objective* is binary, with the corresponding attribute being *whether monopoly markups are prevented*.

**Environmental sustainability** – **a) Maximize modal shift**: Increasing the share of public transport by shifting demand from motorized individual transport is a key strategy to reduce emissions (Carroll et al., 2019), which yields substantial emission savings per (booked) passenger km (Byrne et al., 2021). In rural areas, SMOD systems are particularly well-suited to facilitate this shift as they provide competitive service quality compared to scheduled public transport (Schasché et al., 2023). Additionally, indirect, long-term emission savings may occur if SMOD usage leads to decreased private car ownership, as suggested by evidence from related shared mobility services, such as ride-hailing (e.g., Wang et al., 2021a) and car-sharing (e.g., Jochem et al., 2020).

While maximizing modal shift was less relevant for early SMOD providers (Laws et al., 2008), it has sharply gained importance in recent years (Schasché et al., 2022), which is why we consider it a *primary objective*.

Since we cannot determine which mode a customer would have alternatively chosen, actual modal shift remains unobservable. However, *booked passenger km* can serve as a reasonable proxy attribute. In rural areas, where motorized individual transport dominates (Nobis and Kuhnimhof, 2018), most demand shifted to SMOD services likely originates from this mode (Sörensen et al., 2021).

Environmental sustainability - b) Minimize specific emissions: Besides modal shift, the SMOD system's environmental sustainability also depends on its own emissions. The efficiency of a ride (greenhouse gas emissions per booked passenger km) depends on its consolidation with other rides (Lotze et al., 2023).

We consider minimizing the system's emissions as a *secondary objective* assuming that the indirect emission savings of modal shift are sufficiently large. Since emissions depend heavily on vehicle powertrains and energy sources (García-Afonso, 2023), the *booked pooling rate* serves as a proxy attribute for this objective. This rate is defined as the quotient of booked passenger km and total vehicle km travelled (e.g., Anzenhofer et al., 2025).

**Economic sustainability** – **Minimize subsidy requirements**: As it typical for rural public transport (Imhof and Mayer, 2024), rural SMOD systems are financially unprofitable and depend

on public subsidies (e.g., Sörensen et al., 2021, or Zwick et al., 2022). Economic sustainability aims to minimize these subsidies by increasing fare revenues or reducing variable routing costs. Then, the service can scale up and avoid withdrawal after the pilot stage (Currie and Fournier, 2020, De Jong et al., 2011). This long-term establishment is critical as changing mobility behavior depends on trust in the persistence of improved public transport services (Hahn et al., 2023).

If long-term viability is not acutely at risk, minimizing subsidy requirements can be viewed as a *secondary objective* (Hörcher and Graham, 2020b). At the operational level, we use *profit after fulfillment*, i.e., fare revenue net of variable routing cost (such as fuel, maintenance etc.), as a proxy attribute. Fixed cost (e.g., wages), which also impact subsidy requirements, are excluded from this consideration as they are not decision-relevant.

Table 2 Overview of objectives

Major objective	Lower-level objective	Importance	(Proxy) attribute	
	Maximize basic mobility provision	Primary	Booked passenger km	
Social sustainability	Prevent monopoly markup	Primary	Analytical property	
Environmental quatainability	Maximize modal shift	Primary	Booked passenger km	
Environmental sustainaointy	Minimize specific emissions	Secondary	Pooling rate	
Economic sustainability	Minimize subsidy requirements	Secondary	Profit	

# 3.1.3 Model Structure

To enable mathematical optimization, the multi-attribute decision problem defined in the previous section must be translated into a multi-objective optimization model. Given the stochastic and dynamic nature of the problem, we model the problem as an MDP (Puterman, 2014), the standard modeling approach in i-DMVRPs research (Fleckenstein et al., 2023, Ulmer et al., 2020). Two types of MDP models can incorporate multiple objectives:

- Multi-objective MDPs: This approach scalarizes different rewards into a single reward function (Roijers et al., 2013). In the context of SMOD systems, it is applied by Wu et al. (2024).
- Constrained MDPs: Here, one objective is represented in the objective function, while others are modeled as constraints (Altman, 1999). Applications to i-DMVRPs include Lang et al. (2021) and Lu et al. (2024).

Considering the objectives defined in Section 3.1.2, a multi-objective MDP is clearly not suitable for our problem because profit would be part of the objective function causing monopoly markups. Instead, we formulate a constrained MDP, where only the *primary objectives* are directly included in the objective function. *Secondary objectives* are reflected by constraints on pricing decisions, which prevents monopoly markups by design. Below, we define the objective function and explain how secondary objectives are integrated as constraints:

**Objective function – Maximize booked passenger km**: The objective function is derived from the *primary objectives* (Social a), Environmental a)), which are congruent. Both maximizing basic mobility provision and maximizing modal shift lead to the same conclusion: The SMOD system

should aim to serve as much demand as possible. Hence, the objective function maximizes total booked passenger km.

**Constraint – Marginal cost coverage at request-level**: The objective function, based on the *primary objectives*, implies that the SMOD system should accept a request if its expected displacement of future booked passenger km does not exceed its own booked passenger km. However, during off-peak periods, this approach may lower prices excessively to maximize order conversion probability, creating conflicts with *secondary objectives* (Environmental b), Economic). From these perspectives, serving demand inefficiently, whether in terms of variable routing cost or emissions, is undesirable (Schasché et al., 2022).

To account for the *secondary objectives* while prioritizing the *primary objectives*, we transfer *marginal cost pricing* to SMOD systems (Hörcher and Tirachini, 2021) and present its first practical application in this context. The concept states that the price of a public transport ride must reflect the marginal societal cost directly attributable to it, including the provider's variable cost and potential externalities. For SMOD systems without congestion, the main externalities are variable carbon emissions (Heinitz, 2022, Schasché et al., 2022). To apply marginal cost pricing to SMOD, we impose a lower price bound, ensuring that each ride's price covers at least its marginal cost, including both the external cost of carbon emissions and the variable routing costs. This approach also aligns with the concept of internal carbon pricing (Bento and Gianfrate, 2020).

Both components of marginal cost are proportional to the marginal vehicle km caused by serving a ride. Assuming that the marginal vehicle km can be (approximately) determined, marginal cost is calculated by multiplying cost per vehicle km with marginal vehicle km. To ensure a minimum price level, we assume a pre-defined base price, e.g., depending on the pick-up and drop-off stop of a ride.

It is important to note that the constraint cannot guarantee a certain pooling rate or profit over the booking horizon. Instead, the pricing mechanism differentiates between "favorable rides" (low marginal cost), which are offered at the base price, and "unfavorable rides" (high marginal cost), which are offered at a higher price.

When offered an increased price, the customer is more likely to abandon the booking process, reducing booked passenger km (*primary objectives*) but improving the pooling rate and profit (*secondary objectives*). This risk of abandonment is acceptable, as the unfavorable ride would have incurred relatively high cost, and the customer does not value its utility enough to pay for this cost.

Of course, the customer may still order the ride despite the higher price. Then, there is a positive impact on booked passenger km and profit, but a negative impact on the pooling rate. This is reasonable, as the customer's individual utility for the ride is high enough to justify the negative impact on Environmental b). Overall, the demand is managed in a way that unfavorable, low-utility rides are avoided by allowing customers to assess themselves if a ride is worth its full

societal cost (Eliasson, 2021). This problem structure entirely prevents monopoly markups, thereby fulfilling Social b), given the marginal vehicle km value is accurately calculated.

#### 3.2 Mathematical Modeling

This section translates the structural model definition from Section 3.1.3 into a mathematical formulation. We begin with an MDP model formalizing the complete i-DMVRP (Section 3.2.1). Then, the MDP model is complemented by a model for the dynamic pricing subproblem, which we present in Section 3.2.2.

#### 3.2.1 Modeling the Sequential Decision Problem

The provider's operational planning problem (see Section 3.1) is a sequential decision problem. To formalize it, we generalize the MDP model by Anzenhofer et al. (2024) in two ways: First, we incorporate the sustainable dynamic pricing subproblem. Second, we include the primary objective of maximizing booked passenger km instead of profit. In the following, we state the model by first introducing some basic notation, then defining the MDP's components, i.e., decision epochs, states, actions, transitions, and rewards, and finally formulating the corresponding Bellman equation. Thereby, we only briefly restate the common elements of both models.

- *Planning horizon*: The booking horizon is divided into stages *T* = {1, ..., t<sup>s</sup>, ..., T}, where t<sup>s</sup> marks the start of the service horizon. Each stage is denoted by t ∈ *T*.
- *Requests*: Customers request rides between a pair of stops from a pre-defined set *H*. Each request *i* ∈ *J*, with *J* denoting the set of all requests, is defined by the following attributes:
  - Request type  $c_i$ : Specifies the pick-up stop  $p_{c_i} \in \mathcal{H}$ , drop-off stop  $d_{c_i} \in \mathcal{H}$ , desired time  $t_{c_i}$ , number of passengers  $m_{c_i}$ , and pre-defined base price  $b_{c_i}$ .
  - Booked passenger km  $\rho_{c_i}$ : Calculated as the direct distance from  $p_{c_i}$  to  $d_{c_i}$  multiplied with  $m_{c_i}$ .
  - Time of request τ<sub>i</sub> ∈ T: Denotes the stage when the request arrives. The relation between desired time and time of request yields a classification into three kinds of requests: 1) advance requests (τ<sub>i</sub> < t<sup>s</sup>), 2) same-day requests (τ<sub>i</sub> ≥ t<sup>s</sup>), and 3) ad-hoc requests (τ<sub>i</sub> = t<sub>ci</sub>).

Note that a dummy request type c = 0 indicates the case of no request arrival.

*Rides and products*: The set O<sub>c</sub> stores all rides that can potentially be offered in response to a request of type c. A ride o ∈ O<sub>c</sub> corresponds to a pick-up time or drop-off time that can be offered to a request of type c with desired time t<sub>c</sub>, potentially deviating from the desired time t<sub>c</sub>. A product k ∈ K<sub>c</sub>, with K<sub>c</sub> denoting the set of product indices for request type c, is defined as a combination of a ride o<sub>k</sub> ∈ O<sub>c</sub> and a price multiplier n<sub>k</sub> ∈ N, with N denoting the pre-defined price multipliers.

In less technical terms, it is possible to offer each ride at each of the pre-defined discrete price points. Thus, the price  $r_{c,k}$  of a product k equals the base price of request type c times the

price multiplier of product k. Further, we introduce a dummy product  $k = 0 \in \mathcal{K}_c$  to denote the no-purchase option.

- Order confirmation: If a customer chooses a product k ≠ 0, their request i ∈ J becomes an order j ∈ J with i = j and J denoting the set of all orders. Due to stochastic choice behavior, P<sub>k</sub>(g) encodes the likelihood that product k is chosen from an offer set g ⊆ K<sub>c</sub>. Hence, each order j features a chosen product k<sub>i</sub> ∈ K<sub>c</sub>.
- Order fulfillment: The provider assigns vehicles  $v \in \mathcal{V}$  from a fleet  $\mathcal{V}$  with seat capacity  $Q_v$  to fulfill orders.

Now, we formulate the MDP model using the basic notation defined above.

- Decision epochs: At the beginning of each stage t ∈ T, which is defined as a micro-period of sufficiently small duration to practically exclude more than one request arrival, the provider must make decisions.
- *States:* The post-decision state  $s_t = (C_t, \phi_t)$  includes all data required for demand control and vehicle routing decisions for the newly received request at the subsequent decision epoch t + 1:
  - $C_t$ : The set of all pending orders  $j \in \mathcal{J}$ .
  - $\phi_t$ : The current route plan.
- Actions: At decision epoch t, the action a<sub>t</sub> = (g<sub>t</sub>, (φ<sub>t</sub>(o<sub>k</sub>))<sub>k∈g<sub>t</sub></sub>) comprises two nested decisions:
  - Demand control decision g<sub>t</sub>: Selecting an offer set g<sub>t</sub> ∈ G(s<sub>t-1</sub>, c) ⊆ 2<sup>K<sub>c</sub></sup> \ Ø from the set of feasible offer sets G(s<sub>t-1</sub>, c), which may contain only the no-purchase option. To optimize the composition of g<sub>t</sub>, the provider must solve the dynamic pricing subproblem (see Section 3.2.2).
  - Vehicle routing decision (φ<sub>t</sub>(o<sub>k</sub>))<sub>k∈gt</sub>: Determining a route plan for each ride o<sub>k</sub> associated with an offered product k ∈ g<sub>t</sub>. These plans ensure that all pending orders from s<sub>t-1</sub> and the potential new order from the customer choosing product k are served. The action space for this decision is ∏<sub>k∈gt</sub> Φ(s<sub>t-1</sub>, c, o<sub>k</sub>), representing all combinations of feasible route plans.

Note that only a small part of the vehicle routing decision, i.e., the planned vehicle movements in  $\phi_t(o_k)$  that start until the next decision epoch t + 1, are definitive.

• *Transitions:* The transition from post-decision state  $s_{t-1}$  to successor state  $s_t$  is influenced by the provider's actions and realizations of the exogeneous information, such as request arrivals and customer choice behavior. First, the system transitions stochastically to pre-decision state  $s_t^{\text{pre}}$  as it is revealed which request  $i_t \in \mathcal{I}$  arrives, according to arrival rates  $\lambda_c^t$ . If a request arrives ( $c_{i_t} \neq 0$ ), the provider's decision is followed first by a deterministic transition and then by another stochastic transition to interim state  $s_t'$ . As part of these transitions, the provider presents offer set  $g_t$  and the customer confirms their order  $j_t$  (or abandons the booking process) by choosing a product  $k_{j_t} \in g_t$ , and therewith a ride  $o_{j_t}$ , according to choice probabilities  $P_{c_{i_t},k}(g_t)$ . If  $k_{j_t} \neq 0$ , the newly confirmed order  $j_t$  is added to  $C_{t-1}$ . Finally, the process deterministically reaches the succeeding post-decision state  $s_t$ . As part of this transition, the route plan  $\phi_t(o_{j_t})$  replaces the route plan  $\phi_{t-1}$  in the system state and is potentially partly executed.

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:

$$C_t = (C_{t-1} \cup \{j_t\})$$
(1)

$$\phi_t = \phi_t(o_{j_t}) \tag{2}$$

If no request arrives ( $c_{i_t} = 0$ ), the provider's decision is directly followed by the deterministic transition according to route plan  $\phi_t(0)$ .

- *Rewards:* The provider receives the booked passenger km ρ<sub>cj</sub> associated with a newly confirmed order j.
- Bellman equation: The optimality condition for the MDP model at hand is given by the following Bellman equation, which defines state values  $V_t(s_t)$ :

$$V_{t-1}(s_{t-1}) = \sum_{c \in \mathcal{C}} \lambda_c^t \max_{g_t \in \mathcal{G}(s_{t-1},c)} \left( \sum_{k \in g_t} P_{c,k}(g_t) \left[ \rho_c \cdot \mathbf{1}_{o_k \neq 0} - \Delta V_t(s_{t-1},c,o_k) \right] \right) + V_t'(s_t' | s_{t-1}, 0, 0)$$
(3)

with 
$$V'_t(s'_t|s_{t-1}, c, o_k) = \max_{\phi_t(o_k) \in \Phi(s_{t-1}, c, o_k)} \left( V_t(s_t|s_{t-1}, c, \phi_t(o_k)) \right)$$
 (4)

and 
$$\Delta V_t(s_{t-1}, c, o_k) = V'_t(s'_t|s_{t-1}, c, 0) - V'_t(s'_t|s_{t-1}, c, o_k).$$
 (5)

Formulation (3) of the Bellman equation represents the formal basis for the popular decomposition-based solution concepts for i-DMVRPs (Fleckenstein et al., 2024a), which we also rely on (Section 4). The maximum operator  $\max_{g_t \in \mathcal{G}(s_{t-1},c)}(\cdot)$  corresponds to the dynamic pricing subproblem and yields the offer set  $g_t$ . To evaluate each potential offer set, the expected gain in booked passenger km must be known for each product  $k \in g_t$ . It results from the probability  $P_k(g_t)$  that product k is chosen multiplied with the booked passenger km  $\rho_c$  of the ordered ride net of the expected displacement of booked passenger km  $\Delta V_t(s_{t-1}, c, o_k)$ . The latter result from future rides, which cannot be served when serving the requested ride (Strauss et al., 2018).

## 3.2.2 Modeling the Dynamic Pricing Subproblem

In this section, we provide a detailed model for the dynamic pricing subproblem, which is given by the maximum operator  $\max_{g_t \in \mathcal{G}(s_{t-1},c)} (\cdot)$  in Bellman equation (3). The problem can be modeled as a generalization of the classical assortment optimization problem, and we adopt the respective standard model (Heger and Klein, 2024). Instead of encoding the dynamic pricing decision as a set of offered products  $g_t$ , the model relies on binary decision variables  $x_{o,n}$ , with  $x_{o,n} = 1$  if ride *o* is offered with price multiplier *n* and  $x_{o,n} = 0$  otherwise. Since each product *k* is fully specified by  $o_k$  and  $n_k$ ,  $x_{o_k,n_k}$  indicates whether product *k* is included in the offer set. In addition, the following parameters are part of the model:

- Set O<sup>f</sup> ⊆ O<sub>c</sub> comprises all feasible rides. For a ride o to be feasible (o ∈ O<sup>f</sup>), there must be at least one feasible route plan φ<sub>t</sub>(o) for fulfilling all orders in C<sub>t-1</sub> and the new potential order given ride o is chosen. Mathematically, this is equivalent to Φ(s<sub>t-1</sub>, c, o) ≠ Ø. For determining whether Φ(s<sub>t-1</sub>, c, o) ≠ Ø holds, a constraint satisfaction problem must be solved (see Anzenhofer et al. (2024) for a MIP formulation).
- Probabilities P<sub>on</sub>(x) are equivalent to probabilities P<sub>k</sub>(g<sub>t</sub>) in Bellman equation (3) adapted to the binary encoding of the offer set. To define their values, an arbitrary discrete choice model must be specified (Heger and Klein, 2024). In the work at hand, we draw on the consider-then-choose model proposed by Anzenhofer et al. (2024). Also see Section 5.1.1 for a further description of the specific model used in the work at hand.
- Set  $\mathcal{N}$  comprises all pre-defined price multipliers.
- Parameter  $\rho_c$  denotes the booked passenger km associated with the arriving request of type c.
- Parameter ΔV<sub>t</sub>(s<sub>t-1</sub>, c, o) equals the expected displacement of booked passenger km due to selling ride o to a customer submitting a request of type c in state s<sub>t-1</sub>.
- Parameter  $r_{con}$  denotes the price, i.e., the fare revenue received, when selling ride *o* at price multiplier *n* to a customer submitting a request of type *c*.
- $\epsilon$  is the cost parameter consisting of variable routing cost and external cost (emissions) per vehicle km.
- Parameter  $\Delta F_t(s_{t-1}, c, o)$  denotes the expected marginal vehicle km caused by selling ride *o* to a customer submitting a request of type *c* in state  $s_{t-1}$ .

In the following, we state the model formulation:

$$\max Z(\mathbf{x}) = \sum_{o \in \mathcal{O}^f} \sum_{n \in \mathcal{N}} P_{on}(\mathbf{x}) \cdot [\rho_c - \Delta V_t(s_{t-1}, c, o)] \cdot x_{on}$$
(6)  
s.t.

$$\sum_{n \in \mathcal{N}} x_{on} \le 1 \qquad \qquad \forall o \in \mathcal{O}^f \tag{7}$$

$$\sum_{n \in \mathcal{N}} r_{con} \cdot x_{on} \ge \epsilon \cdot \Delta F_t(s_{t-1}, c, o) \qquad \qquad \forall o \in \mathcal{O}^f$$
(8)

$$x_{on} \in \{0,1\} \qquad \qquad \forall o \in \mathcal{O}^f, n \in \mathcal{N} \qquad (9)$$

The objective function (6) is equivalent to the maximum operator in Bellman equation (3). Thus, it maximizes the expected net contribution of the arriving request in terms of booked passenger km. Although the provider makes dynamic pricing decisions, the price  $r_{con}$  itself is not directly represented in the objective function. However, it still influences the objective function indirectly via choice probabilities  $P_{on}(x)$ , which depend on  $r_{con}$ .

Constraints (7) ensure that each feasible ride  $o \in O^f$  is offered with at most one price multiplier (Davis et al., 2013). Formulating these constraints as inequalities results in the selective variant of the dynamic pricing problem, which allows not offering feasible rides (e.g., Anzenhofer et al., 2024). If every feasible ride has a negative contribution ( $\rho_c - \Delta V_t(s_{t-1}, c, o) < 0 \forall o \in O^f$ ), the provider can reject the request by not offering any ride. Replacing Constraints (7) by equations leads to the non-selective variant. Then, all feasible rides must be offered at some price multiplier. We compare both variants in our computational study (Section 5).

By Constraints (8), the secondary objectives are incorporated into the model. More precisely, the constraints require that the price of each feasible ride  $o \in O^f$  must be at least equal to its marginal cost in terms of variable routing cost and external cost of emissions.

In the selective variant of the model, Constraints (8) may lead to potential request rejections if the highest price multiplier in  $\mathcal{N}$  is insufficient relative to the cost parameter  $\epsilon$ . Then,  $x_{on} = 0 \forall o \in \mathcal{O}^f$ ,  $n \in \mathcal{N}$ , i.e., no feasible rides are offered. To prevent infeasibility in the non-selective variant, either  $\mathcal{N}$  must include sufficiently large price multipliers, or a post-processing step is needed to offer infeasible rides at the highest price multiplier.

## 4 Solution Method

To solve the provider's sequential decision problem (Section 3.2.1) including the dynamic pricing subproblem (Section 3.2.2), we propose a decomposition-based solution concept. We divide the explanation in three parts: First, we provide an overview of the basic algorithmic structure (Section 4.1). Second, we present the post-decision rollout algorithm for ride evaluation (Section 4.2). Third, we show how the dynamic pricing subproblem can be solved by full enumeration (Section 4.3).

#### 4.1 Overview of the Decomposition-based Solution Concept

The basic idea behind the decomposition-based solution concept is as follows: Upon receiving a customer request, ride evaluation is performed for all feasible rides, determined by a feasibility check. Then, the dynamic pricing subproblem is solved based on the results of feasibility check and ride evaluation. The basic algorithmic structure of our solution method reflects this idea (see Fig. 2). In the initial state  $s_0$ , the set of pending orders  $C_0$  and the current route plan  $\phi_0$  are empty (lines 1-2). At each decision epoch, the first step (line 4) is to update  $C_t$  and  $\phi_t$  based on the legs of the route plan completed since the previous decision epoch. The feasibility check (lines 6-9) is performed by based on the current route plan  $\phi_t$  using a parallel insertion heuristic (Anzenhofer et al., 2024). If a feasible insertion position could be found ( $\phi_t(o) \neq \emptyset$ ), ride o is added to the set of feasible rides  $\mathcal{O}_{c_{i_t}}^f$ . Then, ride evaluation is performed for each feasible ride  $o \in \mathcal{O}_{c_{i_t}}^f$  (line 10) to approximate displaced booked passenger km  $\Delta \tilde{V}_t(s_{t-1}, c_{i_t}, o)$  and marginal vehicle km  $\Delta \tilde{F}_t(s_{t-1}, c_{i_t}, o)$ . Given the results of feasibility check and ride evaluation, the dynamic pricing

subproblem is solved to determine the offer set  $g_t^*$  (line 11). Next, the exogeneous information about the customer's purchase choice realizes in the form of the chosen product  $k_{j_t} \in g_t^*$  (line 12). If  $k_{j_t} \neq 0$ , the newly collected order is added to  $C_t$  (lines 13-14) and a vehicle routing decision is made using the route plan  $\phi_t(o_{j_t})$  resulting from the feasibility check for the ride  $o_{j_t}$ chosen by the customer. Note that in case of a no-purchase, the route plan is not changed.

1	$C_0 := \emptyset$
2	Initialize $\phi_0$ as empty
3	forall $t \in \mathcal{T}$ do
4	$\phi_t, C_t \coloneqq execute\_route\_plan(\phi_{t-1}, C_{t-1}, \tau_{i_t})$
5	$\mathcal{O}^f_{c_{i_t}} \coloneqq \emptyset$
6	forall $o \in \mathcal{O}_{c_{i_t}}$ do
7	$\phi_t(o) \coloneqq feasibility\_check\left(\phi_t, \mathcal{O}^f_{c_{i_t}}, i_t, o\right)$
8	if $\phi_t(o) \neq \emptyset$ :
9	$\mathcal{O}^{f}_{c_{i_{t}}} \coloneqq \mathcal{O}^{f}_{c_{i_{t}}} \cup \{o\}$
10	$\left\{ \left( \Delta \widetilde{V_t}(s_{t-1}, c_{i_t}, o), \Delta \widetilde{F_t}(s_{t-1}, c_{i_t}, o) \right) : o \in \mathcal{O}_{c_{i_t}}^f \right\} \coloneqq ride\_evaluation\left( \mathcal{O}_{c_{i_t}}^f, \phi_t, \left\{ \phi_t(o) : o \in \mathcal{O}_{c_{i_t}}^f \right\} \right)$
11	$g_t^* \coloneqq demand\_control\left(i_t, \mathcal{O}_{c_{i_t}}^f \left\{ \left( \Delta \widetilde{V}_t(s_{t-1}, c_{i_t}, o), \Delta \widetilde{F}_t(s_{t-1}, c_{i_t}, o) \right) : o \in \mathcal{O}_{c_{i_t}}^f \right\} \right)$
12	Observe $k_{j_t} \in g_t^*$
13	if $k_{j_t} \neq 0$ then
14	$C_t \coloneqq C_t \cup \{j_t\}$
15	$\phi_t \coloneqq routing\_control(\phi_t(o_{j_t}), \mathcal{C}_t)$

Fig. 2 Basic solution algorithm

#### 4.2 Ride Evaluation

Ride evaluation is equivalent to approximating opportunity cost (displacement cost and marginal cost-to-serve), in profit-based i-DMVRPs (Fleckenstein et al., 2024a). For sustainable dynamic pricing, displaced booked passenger km  $\Delta \tilde{V}_t(s_{t-1}, c_{i_t}, o)$  are needed to formulate objective function (6). Marginal vehicle km  $\Delta \tilde{F}_t(s_{t-1}, c_{i_t}, o)$  enter Constraints (8) as a parameter. Their true values depend on the final set of orders (see Lang et al. (2021) who introduce the term "set-related"). For this reason, authors propose ex-post pricing (Andrejszki and Török, 2018, Karaenke et al., 2023) as a solution. It allows calculating  $\Delta \tilde{V}_t(s_{t-1}, c_{i_t}, o)$  and  $\Delta \tilde{F}_t(s_{t-1}, c_{i_t}, o)$  at the end of the booking horizon. However, for dynamic pricing, both values must be calculated in real-time, and the final set of orders must be anticipated to compute accurate approximations  $\Delta \tilde{V}_t(s_{t-1}, c_{i_t}, o)$ .

To this end, we propose a post-decision rollout algorithm (Bertsekas et al., 1997). The main reason is that its properties are well-suited to the characteristics of the rural SMOD dynamic pricing problem. First, it requires no offline training and hardly any parameter tuning, which makes it

readily applicable in practice. Second, as shown by Ulmer et al. (2016), it outperforms learningbased approaches if the explicit consideration of customer locations is advantageous. This is indeed the case in our problem since requests occur in small villages scattered across the service area with no demand at all in the areas in-between. Third, a post-decision rollout algorithm can identify detailed structure in the MDP (Ulmer et al., 2019), which is crucial in the rural problem setting with small instances and generally few pooling opportunities (Anzenhofer et al., 2024).

The basic idea underlying post-decision rollout algorithms is to explicitly explore multiple branches of the decision tree starting from the possible post-decision states at the current decision epoch. In less technical terms, the algorithm simulates the possible future evolution of the booking horizon resulting from making a certain decision now. Applied to the problem at hand, we can compare the observed total booked passenger km and total vehicle km in case a certain order is confirmed to the case that no order is confirmed. The respective difference yields an approximation of  $\Delta \tilde{V}_t(s_{t-1}, c_{i_t}, o)$  and  $\Delta \tilde{F}_t(s_{t-1}, c_{i_t}, o)$ .

The key design elements of the algorithm are the set of sample paths  $\Omega := \{\omega_q : q = 1, ..., |\Omega|\}$  of future customer requests and the base policy. In the following, we briefly describe the design we propose:

Sample paths: Each sample path ω ∈ Ω corresponds to a sequence of request arrivals. In our algorithm, this sequence covers the entire remaining booking horizon. This is necessary for an SMOD system with advance requests because the desired time of arriving requests has no natural ordering (Elting and Ehmke, 2021). Hence, the impact of a decision at the current decision epoch can extend to any point in the future and does not systematically decline, which would call for a limited sampling horizon as it is the case for similar i-DMVRPs, e.g., arising in same-day delivery (Klein and Steinhardt, 2023).

For generating the sample paths  $\Omega$ , we use instance-based sampling, which performs superior to distribution-based sampling techniques (Köhler et al., 2024). Each instance-based sample path corresponds to the request arrivals (with their original features) of a specific historical booking horizon. To ensure that the sample path starts at the correct time, we remove all requests placed before the current decision epoch.

Finally, we determine the number of sample paths  $|\Omega| = \min\{B^{sta}, B^{dyn}\}$  as the minimum of a static upper bound  $B^{sta} \coloneqq \alpha$  and a dynamic upper bound  $B^{dyn} \coloneqq \left\lfloor \frac{\beta}{|o_{c_{l_l}}^{\ell}| + 1} \right\rfloor$ . The latter

is computed as follows: We measure the runtime w of the first simulated sample path. Using a pre-defined time budget  $\beta$  for the rollout algorithm per decision epoch, we then estimate the total number of sample paths that can be simulated within the time budget. Finally, we divide this value by the number of feasible rides (including the no-purchase option) to obtain an upper bound on the sample paths that can be simulated for all feasible rides within the time budget. With this definition of  $|\Omega|$ , we avoid two undesirable effects: At the beginning of the booking horizon, even a rather small number of sample paths can consume much runtime, so  $B^{dyn}$  prevents excessive loading times of the smartphone app. At the end of the booking horizon, sample paths are very short and yield limited information. Hence,  $B^{sta}$  prevents the simulation of too many sample paths.

Base policy: We offer each feasible ride at the base price g<sup>base</sup> := {k ∈ K<sub>c</sub>: o<sub>k</sub> ∈ O<sup>f</sup><sub>cit</sub>, n<sub>k</sub> = 1} ∪ {0}, which ensures a short runtime of the simulation. While this would also be true for, e.g., rule-based dynamic pricing policies, there is a strong argument for using static pricing: A dynamic pricing base policy would tend to price out requests that, based on the known orders, cannot be pooled well and would instead "cherry-pick" requests that can be pooled well. Hence, the anticipation would cause a "self-fulfilling prophecy": If it is anticipated that a certain subset of requests that currently cannot be pooled well will be priced out in the future, a currently arriving request of this subset will actually be priced out due to its high marginal vehicle km. With static pricing, requests receive comparable offers, which ensures that actual decisions are also more balanced, especially in the critical early phase of the booking process (Fleckenstein et al., 2024b)

Fig. 3 depicts a pseudocode of the ride evaluation algorithm, which is initiated in line 10 of the basic solution algorithm (Fig. 2). In the initialization, the set of sample paths  $\Omega$  is generated (line 1). Then, starting with the no-purchase option o = 0, each sample path is simulated for each ride starting from the post-decision state  $(C^{o,q}, \phi^{o,q})$  that results from collecting a hypothetical order  $j'_t(o)$  for this ride (lines 4-6). To perform the simulation, the dynamic pricing subproblem is solved (Fig. 4) using the base policy (line 11 in Fig. 2). Finally, for all rides  $o \in \mathcal{O}_{c_{it}}^f$ , the realized total booked passenger km  $\hat{V}_{o,q}$  is compared to that of the no-purchase option  $\hat{V}_{0,q}$ . The average difference over all sample paths yields an approximation of displaced booked passenger km  $\Delta \tilde{V}_t(s_{t-1}, c_{i_t}, o)$  (line 8). Likewise, we approximate the marginal vehicle km  $\Delta \tilde{F}_t(s_{t-1}, c_{i_t}, o)$  (line 9).

<b>Input</b> : Set of feasible rides $\mathcal{O}_{c_{i_t}}^f$ , current route plan $\phi_t$ , set of route plans including potential orders $\{\phi_t(o): o \in \mathcal{O}_{c_{i_t}}^f\}$
1 Draw set of sample paths $\Omega \coloneqq \{\omega_q : q = 1,,  \Omega \}$ of future customer requests
2 <b>forall</b> $o \in \mathcal{O}_{c_{i_t}}^f \cup \{0\}$ <b>do</b>
3 forall $\omega_q \in \Omega$ do
4 $C^{o,q} \coloneqq C_t \cup \{j'_t(o)\}$
5 $\phi^{o,q} \coloneqq \phi_t(o)$
6 $\hat{V}_{o,q}, \hat{F}_{o,q} \coloneqq simulation_with\_base\_policy(C^{o,q}, \phi^{o,q}, \omega_q)$
7 <b>if</b> $o \neq 0$ then
8 $\Delta \widetilde{V}_t(s_{t-1}, c_{i_t}, o) \coloneqq \frac{\sum_{q=1}^{ \Omega } (\widetilde{V}_{0,q} - \widetilde{V}_{o,q})}{ \Omega }$
9 $\Delta \widetilde{F}_t(s_{t-1}, c_{i_t}, o) \coloneqq \frac{\sum_{q=1}^{ \Omega } (\widehat{F}_{o,q} - \widehat{F}_{0,q})}{ \Omega }$
Eig 2 Dart darising will not all and then from it a combination

Fig. 3 Post-decision rollout algorithm for ride evaluation

#### 4.3 Solving the Dynamic Pricing Problem

In this section, we consider the dynamic pricing subproblem for an individual request  $i_t$  at decision epoch t. Given the set of feasible rides, the approximations of displaced booked passenger km and marginal vehicle km, and the choice probabilities, Model (6)-(9) must be solved to determine the offer set  $g_t^*$ . Since prices are discrete, Model (6)-(9) corresponds to a constrained assortment optimization problem. As an alternative to dedicated solution algorithms (Heger and Klein, 2024), full enumeration can also be applied if the number of possible offer sets is not prohibitively large. For the problem at hand, this is a valid assumption:

- *Low number of price multipliers*: Naturally, the price multiplier has a lower bound (1) and an upper bound well below taxi fares. Also, for simplicity, the provider likely selects only a few price multipliers.
- *Low number of potential rides*: Similarly, the number of potential rides is likely low because of the display capacity of smartphones (Arian et al., 2022, Haferkamp, 2024).
- *Preprocessing opportunities*: Constraints (7) and Constraints (8) can be exploited in pre-processing steps, such that infeasible offer sets can be excluded from the enumeration process apriori.

**Input**: Current customer request  $i_t$ , set of feasible rides  $\mathcal{O}_{c_{i_t}}^f$ , set of estimated displaced booked passenger km and marginal vehicle km  $\left\{ \left( \Delta \tilde{V}_t(s_{t-1}, c_{i_t}, o), \Delta \tilde{F}_t(s_{t-1}, c_{i_t}, o) \right) : o \in \mathcal{O}_{c_{i_t}}^f \right\}$ 

1 
$$\mathcal{K}_{c_{i_t}} \coloneqq \left\{ k \in \mathcal{K}_{c_{i_t}} : o_k \in \mathcal{O}_{c_{i_t}}^f, \rho_{c_{i_t}} - \Delta \widetilde{V}_t(s_{t-1}, c_{i_t}, o_k) \ge 0, r_{c_{i_t}, k} \ge \epsilon \cdot \Delta \widetilde{F}_t(s_{t-1}, c_{i_t}, o_k) \right\}$$

2 
$$\mathcal{G}(s_{t-1}, c_{i_t}) \coloneqq \left\{ g \in 2^{\mathcal{K}'_{c_{i_t}}} : o_k \neq o_{k'} \forall k, k' \in g, k \neq k' \right\}$$

3 
$$g_t^* \coloneqq \operatorname*{argmax}_{g_t \in \mathcal{G}(s_{t-1},c_{i_t})} \left( \sum_{k \in g_t} P_k(g_t) \left[ \rho_{c_{i_t}} \cdot \mathbf{1}_{o_k \neq 0} - \Delta \widetilde{V}_t(s_{t-1},c_{i_t},o) \right] \right)$$

Accordingly, the algorithm we propose for solving Model (6)-(9) consists of two pre-processing steps before full enumeration is performed to determine  $g_t^*$  (Fig. 4). The first pre-processing step computes the set of feasible products  $\mathcal{K}'_{c_{i_t}} \subseteq \mathcal{K}_{c_{i_t}}$  (line 1). A feasible product  $k \in \mathcal{K}_{c_{i_t}}$  must satisfy three conditions:

- First, the associated ride  $o_k$  must be feasible based on the result of the feasibility check.
- Second, the displaced booked passenger km  $\tilde{V}_t(s_{t-1}, c_{i_t}, o_k)$  attributable to the associated ride  $o_k$  must not exceed the booked passenger km  $\rho_{c_{i_t}}$  gained with the request.
- Third, the price  $r_{c_{i_t},k}$  of product k must at least equal its marginal cost  $\epsilon \cdot \Delta \widetilde{F}_t(s_{t-1}, c_{i_t}, o_k)$ .

Next, the enumeration is performed including another pre-processing step (line 2). Based on the first pre-processing step, any element of the power set of the set of feasible products  $2^{\mathcal{K}'_{c_{i_t}}}$  must be enumerated. However, due to Constraints (7), we can exclude any offer set that contains at least one ride at multiple price points. Thus, line 2 yields the action space  $\mathcal{G}(s_{t-1}, c_{i_t})$  for the

Fig. 4 Solution algorithm for dynamic pricing subproblem

dynamic pricing decision. In line 3, we solve the maximum operator in Bellman equation (3) drawing on the approximated displaced booked passenger km  $\Delta \tilde{V}_t(s_{t-1}, c_{i_t}, o)$ .

# 5 Computational Results

This section presents the results of a computational study evaluating sustainable dynamic pricing against benchmarks from practice and literature. After introducing the experimental setup (Section 5.1), we analyze the base scenario (Section 5.2). Then, we conduct sensitivity analyses regarding the cost parameter (Section 5.3). Finally, we provide insights into the benefits of anticipatory ride evaluation (Section 5.4).

## 5.1 Experimental Setup

All experiments are based on a real-world data set provided by our industry partner FLEXIBUS and are conducted on an Intel<sup>©</sup> Core<sup>©</sup> i7-8700 processor with 6 cores, 3.20 GHz, and 32 GB RAM. The algorithms were implemented in PYTHON (Version 3.11). In the following, we first describe how we generate problem instances from the real-world data set (Section 5.1.1) and which benchmark policies we consider (Section 5.1.2).

#### 5.1.1 Generation of Problem Instances

FLEXIBUS operates an SMOD service across multiple areas, each requiring separate operational planning. For this study, we use data from the Krumbach service area, where FLEXIBUS has been active since 2009. For instance generation, we draw on data for 200 service horizons (working days) during an observation period between February 2022 and February 2023, excluding weekends and holidays. Each instance matches one of the 200 historical days and is defined by two subsets of parameters: request parameters and scenario parameters.

Request parameters are derived directly from historical data. Request arrival sequence and request parameters match the original data set. Scenario parameters represent higher planning level decisions (system parameters), customer choice behavior (choice parameters), and the configuration of the solution method (hyperparameters). These parameters are assumed to be given and introduced below, along with their values in the base scenario.

System parameters: In the considered service area, FLEXIBUS offers transportation between  $|\mathcal{H}| = 563$  stops, including a depot. We determine travel distances and travel times (including a service time of one minute) using Open Source Routing Machine (OSRM, n.d.). The values defining fleet plan, ride options, and time windows are based on service quality targets set by the Association of German Transport Companies (VDV, 2023):

• *Vehicle fleet*: We assume  $|\mathcal{V}| = 2$  vehicles: one throughout the service day, the second operates an additional eight-hour shift (10 a.m. to 4 p.m.). This setup ensures request acceptance rates within the rural service quality target (between 70% and 90%).

- *Ride options*: For any type of request, the provider can offer  $|\mathcal{O}_c| = 5$  potential rides allowing pick-up and drop-off times 30 (60) minutes earlier or later than the desired time, which itself is also a potential ride.
- *Time windows*: Pick-up and drop-off time windows are constructed using a waiting time of 10 minutes and a maximum added ride time factor of 0.5.

The base prices follow the pricing scheme applied by FLEXIBUS during the observation period, which derives prices from the number of fare zones a direct ride from pick-up stop to drop-off stop traverses and the number of passengers. Hence, prices roughly depend on the booked passenger km and range from 2.4€ (one zone) to 9.9€ (eight zones) per passenger. Further, we define  $|\mathcal{N}| = 5$  price multipliers (1, 1.4, 1.8, 2.2, and 2.6). It is worth noting that even the highest price multiplier still only equals about 50% of the local taxi price. The cost parameter is set to  $\epsilon = 0.3 \frac{\epsilon}{km} + 0.2 \frac{\epsilon}{km} = 0.5 \frac{\epsilon}{km}$ , combining FLEXIBUS's operational experience with variable routing costs and recent estimates of external cost from the literature (Lethmate and Paegert, 2024).

Choice parameters: To define choice probabilities, we generalize the consider-then-choose discrete choice model proposed in Anzenhofer et al. (2024). It assumes that the customer population can be divided into customer segments  $l \in \mathcal{L}$  with a unique consideration set of products  $S_l \subseteq \mathcal{K}$ . In addition to the quality cutoff that restricts the deviation from the desired time  $t_c$  that the customer is willing to accept in both directions, we also define a price cutoff, i.e., a maximum acceptable price increase beyond the base price. Combining both cutoffs results in  $|\mathcal{L}| = 15$  customer segments. To determine the segment shares  $\gamma_l$ , we rely on rough assumptions due to the lack of empirical data specific to rural SMOD systems. On the one hand, we assume a price sensitivity of -0.5, which is in line with the empirical findings for public transport in general (Hansson et al., 2019, Holmgren, 2007, Hörcher and Tirachini, 2021). On the other hand, we assume that 25% of the customers are inflexible (cutoff at deviation of 0 minutes), 50% have a low flexibility (cutoff at deviation of 30 minutes), and 25% have a high flexibility (cutoff at deviation of 60 minutes). Within their consideration set  $S_l$ , customers from all segments  $l \in \mathcal{L}$  rank the products according to a uniform ranking function  $\zeta$ . We assume that the price serves as the primary ranking criterion, and the deviation from the desired time serves as the secondary ranking criterion, with earlier rides ranked higher than later rides with the same deviation.

Table 3 provides an overview of the choice model. To apply the different policies to solve the instances, we use a simulator that replays the historical sequence of requests and determines a realization of the customers' purchase choice according to the consider-then-choose model. **Table 3** Discrete choice model

l	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Quality cutoff	0	0	0	0	0	30	30	30	30	30	60	60	60	60	60
Price cutoff	1	1.4	1.8	2.2	2.6	1	1.4	1.8	2.2	2.6	1	1.4	1.8	2.2	2.6
$\gamma_l$	0.05	0.05	0.05	0.05	0.05	0.10	0.10	0.10	0.10	0.10	0.05	0.05	0.05	0.05	0.05

**Hyperparameters**: The anticipatory post-decision rollout algorithm has two hyperparameters defining how the set of sample paths is computed. First, parameter  $\alpha$  sets a fixed upper bound for the number of sample paths. In the base scenario, we set  $\alpha = 15$ . Second, parameter  $\beta$  defines the time budget available for executing the post-decision rollout algorithm at a decision epoch. For this parameter, we select a value of  $\beta = 10$  seconds. A sensitivity analysis regarding the values of both hyperparameters can be found in Appendix C.

#### 5.1.2 Benchmark Policies

We compare the performance of the proposed method to several benchmarks. In the following, we introduce them by explaining their modeling differences and algorithmic differences compared to the MDP model presented in Section (3.2.1) and the solution method presented in Section 4, i.e., **sustainable dynamic pricing (S)**.

- Base price policy (B): This policy represents the status-quo at our industry partner FLEXI-BUS. In response to any type of request c ∈ C, it offers the full set of feasible rides O<sup>f</sup> at the base price b<sub>c</sub>. Hence, from a modeling perspective, prices are static and the dynamic pricing subproblem is entirely omitted. Accordingly, in Fig. 2, line 10 is removed because no ride evaluation is necessary. Further, line 11 is replaced by g<sub>t</sub><sup>\*</sup> := {k ∈ K<sub>c</sub>: o<sub>k</sub> ∈ O<sup>f</sup><sub>cit</sub>, r<sub>c,k</sub> = b<sub>c</sub>} ∪ {0}.
- Profit-based dynamic pricing policy (P): To compare our proposed approach to the state-of-the-art in the scientific literature, we consider a profit-based benchmark policy. It results from changing the reward definition of the MDP model such that the provider receives a positive reward for selling a product, which equals its price, and a negative reward for the variable routing cost caused by routing decisions (see Anzenhofer et al. (2024) for a formal definition). Analogously, Bellman equation (3) and the objective function (6) of the model for the dynamic pricing subproblem must be adapted. Further, Constraints (8) must be removed. For the ride evaluation algorithm, this means that Δ*V*<sub>t</sub>(s<sub>t-1</sub>, c<sub>it</sub>, o) encodes the monetary opportunity cost. Other than that, we make no changes to the ride evaluation algorithm (Fig. 3). In the algorithm for solving the dynamic pricing subproblem (Fig. 4), the third condition in line 1 is removed.
- Sustainable dynamic pricing policy with primary objectives (S-P): To analyze the impact of considering the secondary objectives by Constraints (8), we apply a policy without these constraints. This means that, in Fig. 4, the third condition in line 1 is removed and the policy can offer rides at any price multiplier.
- Non-selective sustainable dynamic pricing policy (S-N): In practice, providers may not want to deny customers the booking of rides that are feasible but have a negative contribution in terms of booked passenger km or have marginal cost exceeding even the highest possible price. Hence, we also test a non-selective benchmark policy that offers every feasible ride. From a modeling point of view, this means that Constraints (7) are replaced by the

corresponding equations. To adapt the solution method, we insert an additional pre-processing step between line 1 and line 2 (Fig. 4). It (re-)inserts the product with the highest price point for each ride into the feasible products  $(\mathcal{K}'_{c_{i_t}} \coloneqq \mathcal{K}'_{c_{i_t}} \cup \{k \in \mathcal{K}_{c_{i_t}} : o_k \in \mathcal{O}^f_{c_{i_t}}, n_k = |\mathcal{N}|\}$ ). Also, we add a condition in line 2, ensuring that feasible offer sets contain all feasible rides  $(|g| = |\mathcal{O}^f_{c_{i_t}}|)$ .

- Sustainable dynamic pricing policy with myopic marginal vehicle km (S-MM): To analyze the benefit of the anticipatory approximation of marginal vehicle km, we consider a partly myopic benchmark policy. It draws on the marginal increase in vehicle km due to inserting the hypothetical ride into the tentative route plan. Hence, we replace line 9 in Fig. 3 by Δ*F*<sub>t</sub>(s<sub>t-1</sub>, c<sub>it</sub>, o) := distance(φ<sub>t</sub>(o)) distance(φ<sub>t</sub>(0)). Note that the policy still approximates the displaced booked passenger km (no change to line 8 in Fig. 3).
- Sustainable dynamic pricing policy without displacement (S-MD): To analyze the benefit of considering displaced booked passenger km, we apply a policy that is partly myopic since it does not capture displaced demand. It results from replacing line 8 by ΔV<sub>t</sub>(s<sub>t-1</sub>, c<sub>it</sub>, o) := 0 in Fig. 3. Again, the policy is not fully myopic due to the anticipatory approximation for the marginal vehicle km (no change to line 9 in Fig. 3).

#### 5.2 Base Scenario

In this section, we analyze performance of the policies in the base scenario as defined in Section 5.1.1. All reported values are arithmetic means over 200 instances generated from the FLEXIBUS data set. For now, we only consider the non-myopic benchmark policies B, P, S-P, and S-NS. We begin with analyzing the solution quality regarding the (proxy) attributes defined in Section 3.1.2, i.e., booked passenger km, pooling rate, and profit (Section 5.2.1). Then, we analyze the offers made to different types of requests (Section 5.2.2).

#### 5.2.1 Solution Quality Analysis

Table 4 compares the solution quality of the newly introduced Policy S to benchmark policies. *B vs. S-P*: We observe an increase in the primary objective (3%) and both secondary objectives (pooling rate: 7%, profit: 15%). This result implies that dynamic pricing, even if based purely on the primary objective, raises the general solution quality including the secondary objectives. The relatively small gain in terms of the primary objective can be explained by weak displacement effects because of the low vehicle utilization of 65%.

*S vs. S-P*: Policy S leads to a re-distribution of solution quality from the primary objective to the secondary objectives due to Constraints (8). The improvement compared to Policy S-P by 4% (pooling rate) and 25% (profit) comes at a loss of -7% in terms of booked passenger km.

*S vs. S-N*: Policy S-N shows a very similar performance compared to Policy S. Seemingly, most demand with a negative contribution in booked passenger km or very high marginal cost can be

successfully priced out by Policy S-N. If this is the case, withholding feasible fulfillment options from these customers is not beneficial.

*P vs. B or S*: Finally, for Policy P, we observe a drastic growth in profit (138%) and a small growth in the pooling rate (8%) compared to Policy B, which is associated with a considerable loss in booked passenger km (-29%). Policy S largely avoids this loss and still improves the secondary objectives considerably.

Table 4 Overview of solution quality

Policy	В	Р	S-P	S-N	S
Booked passenger km	375.28	265.15	387.76	360.35	360.03
Pooling rate	0.6414	0.6911	0.6859	0.7102	0.7137
Profit [€]	89.69	213.72	103.18	135.33	128.74

## 5.2.2 Characteristics of Offers and Orders

Now, we additionally analyze the request level to better understand the root causes of the performance differences and find out how customers would perceive the behavior of the different policies. To this end, we first define suitable metrics and analyze the entire customer population. Then, we consider groups of similar requests and search for disparities in terms of the offers they receive. Table 5 shows the metrics characterizing the offers.

## Entire customer population:

In terms of acceptance and rejection rates, all policies show similar results, which are at the upper end of the target corridor of 70% to 90% (VDV, 2023). Static pricing yields the lowest acceptance rate, which implies that with any dynamic pricing policy, customers have a higher chance of receiving an offer. For selective sustainable dynamic pricing (Policy S-P and Policy S), only a small share of requests is rejected despite feasible rides exist, which again points toward weak displacement effects in general.

Greater differences are observable regarding the offers made.

- Policy S-P gently steers demand: Compared to Policy B, the abandonment rate increases slightly since offer sets become smaller. The base offer rate, and thus, also the lower bound offer rate, remains at 100%, i.e., no monopoly markup is charged. Hence, prices are increased only slightly to steer customers toward rides with a higher gain in booked passenger km.
- For Policy S, changes are greater because the consideration of secondary objectives causes more steering of demand. The average price level increases, and the base price is only available to 80% of accepted requests. Still, the lower bound offer rate is at 100%, which shows that no monopoly markup is charged.
- The corresponding non-selective policy (Policy S-N) makes more offers with larger offer sets, but the additional offers are accompanied by a higher average price and a higher abandonment rate. The 5% of offers without a base price correspond to requests that the selective Policy S would reject and that are priced out.

• For profit-based dynamic pricing (Policy P), demand management is aggressive: Despite offering many rides overall, hardly any customer is offered a ride at the base price. Only 27% of all offers contain at least one option at the lowest price above marginal cost, so the provider charges a monopoly markup with 73% of all offers. Together with the high average price, this causes a no-purchase for almost half of the accepted requests. This explains the loss in booked passenger km observed in Section 5.2.1. Even though more customers receive offers with, on average, more rides, demand is lost due to the high abandonment rate.

Metric	Description	В	Р	S-P	S-N	S
Acceptance rate	Share of requests with at least one offered ride.	0.83	0.93	0.84	0.89	0.84
Rejection rate	Share of requests without an offer.	0.17	0.07	0.16	0.11	0.16
Feasibility	Share of requests for which no feasible ride is found.	0.17	0.07	0.11	0.11	0.11
Displacement	Share of requests for which no feasible ride with a positive contribution in booked passenger km is found.	0	0	0.05	0	0.05
Marginal cost	Share of requests for which all feasible rides have a higher marginal cost than the highest possible price.	0	0	0	0	0
Abandonment rate	Share of accepted requests that did not convert into an order.	0.12	0.43	0.17	0.25	0.22
Base price offer rate	Share of accepted requests whose offer includes at least one ride at the base price.	1	0.02	1	0.75	0.80
Lower bound offer rate	Share of accepted requests whose offer includes at least one ride at the lowest possible price above the marginal cost.	1	0.27	1	0.95	1
Average offer set size	Average number of rides in the offer set.	3.31	3.63	2.81	3.46	2.83
Average price multiplier	Average price multiplier of all rides in the offer set.	1	1.79	1.06	1.52	1.25

Table 5 Results of metrics characterizing offers and orders

Groups of requests - Price paid, marginal vehicle km, and deviation from desired time

Now, we analyze the relationship between the price paid by a customer and two key factors: a) the marginal vehicle km of their request and b) the deviation from the desired time. The former relation reveals how policies S-P, S, and P set prices for requests that are progressively unfavorable in terms of the secondary objectives. The latter relation shows if customers can get cheaper prices if they are more time flexible.

We visualize the results in the form of heatmaps in Fig. 5, where the vertical axis plots the price multiplier of the product the customer ordered (including no-purchases). The horizontal axis plots the marginal vehicle km of the ordered ride in 2.5 km bins. In case of a no-purchase (np), we calculate the average over all feasible rides. Each pixel value represents the percentage share of offers resulting in the same outcome (price multiplier/np) among all offers for rides of a marginal vehicle km bin.

• *Policy S-P*: We observe that most customers order a ride at the base price, irrespective of its marginal vehicle km. Price multipliers above 1.4 are practically never chosen. Hence, the

policy only gently steers customers toward more favorable rides in terms of displaced demand, and prices do not reflect marginal vehicle km.

- *Policy S*: For requests with low marginal vehicle km, Policy S behaves similarly since marginal cost is still below the base price. With increasing marginal vehicle km, we observe a stepped increase of the paid price multiplier, which is very "sharp", indicating a strong dependency between both features. The increasing share of no-purchases shows that customers not willing to pay the marginal cost of their ride are priced out.
- *Policy P*: We observe a general price increase, which is independent from marginal vehicle km. This is further evidence for the monopoly markup charged by a profit-maximizing provider. We still find that the extent of the price increase depends on the marginal vehicle km, albeit much more "loosely" compared to Policy S. The reason is that higher marginal cost requires a higher revenue to make a request profitable.





Fig. 6 Dependency of price and deviation from desired time over all offers

In Fig. 6, the horizontal axis plots the deviation of the ordered pick-up (drop-off) time from the desired time. Pixel values are percentage shares among all orders with the same deviation from

the desired time. For all three policies, we find that customers who are flexible regarding their pick-up (drop-off) time can get cheaper rides. For Policy S-P, a flexibility of 60 minutes guarantees a ride at the base price. While it is consistently observable, this pattern is generally weaker compared to the relation between price and marginal vehicle km (Fig. 5).

# Groups of requests – Time of request, desired time, direct distance between pick-up and drop-off

Next, we analyze selected metrics for different groups of requests. Therewith, we can find disparities in the offers made based on time of request, desired time, or direct distance between pickup and drop-off. To this end, we compare the Policy B and Policy P to the sustainable dynamic pricing by Policy S and Policy S-P. Since they are quite extensive, we refer the reader to Appendix B for the detailed results. In the following, we provide a high-level presentation (Table 6) complemented by an explanation of the most prominent observations.

The columns of Table 6 correspond to the combination of a certain request attribute and a benchmark policy. A row lists a certain metric, in terms of which disparities can potentially occur. The table entries indicate whether sustainable dynamic pricing (Policy S and Policy S-P), exacerbates (-/--), alleviates (+/++) or does not alter (o) the disparities regarding the respective metric (row) and request parameter compared to the respective benchmark policy (column). Below, we elaborate on the disparities that are exacerbated and alleviated.

## Exacerbated:

- Regarding the acceptance rate, we observe discrimination against very short requests by both sustainable dynamic pricing policies. In comparison, Policy P achieves a much more balanced acceptance rate.
- For Policy S, the base price offer rate deteriorates slightly for same-day requests, requests with very early or very late desired times, and more strongly for long requests.
- Policy S and Policy S-P offer markedly less rides to short requests. The benchmarks' offers are balanced.
- Compared to both benchmarks, we observe slightly more unbalanced average price multipliers offered by Policy S. Prices are higher for long requests and requests with a very early or very late desired time.

Overall, we do not see strong indications of discriminatory behavior by the sustainable dynamic pricing policies. Even the strongest inequalities regarding request length are not systematic: Short requests face a lower acceptance rate and offer set size yet also lower price multipliers, and vice versa for long requests.

#### Alleviated:

- Due to its first-come-first-served nature, Policy B favors early advance requests over later requests. Policy S and Policy S-P behave in a more balanced way, also in terms of the abandonment rate.
- Policy B achieves a progressively lower acceptance rate the greater the longer a request is. In contrast, Policies S and S-P accept fewer of the very short requests but more of the longer requests.
- Policy P causes long requests to abandon exceedingly. Policy S and Policy S-P show more balanced rates.
- With a lower bound offer rate of 1, both Policy S and Policy S-P alleviate disparities observed for Policy P, which discriminates against requests based on their desired time and direct distance.

In summary, we find that sustainable dynamic pricing alleviates crucial imbalances compared to the benchmark policies regarding time of request and request length.

Request parameter	Time of	request	Desire	ed time	Direct distance		
Benchmark	Policy B	Policy P	Policy B	Policy P	Policy B	Policy P	
Acceptance rate	++	0	0	0	++		
Abandonment rate	+	0	0	0	0	++	
Base price offer rate	-	-	-	-		0	
Lower bound offer rate	0	0	0	++	0	++	
Average offer set size	0	0	0	0			
Average price multiplier	0	0	-	-		0	

Table 6 Impact of sustainable dynamic pricing on disparities in terms of offers compared to benchmarks

## 5.3 Sensitivity Regarding Cost Parameter

In this section, we vary the cost parameter value  $\epsilon$  with the aim of analyzing the shift of solution quality between primary objective and secondary objectives by adjusting the cost parameter  $\epsilon$ . In the underlying computational experiment, we maintain the parameter setting of the base scenario and vary the cost parameter starting with  $\epsilon = 0 \frac{\epsilon}{km}$  and increasing it successively by  $0.25 \frac{\epsilon}{km}$  until we reach  $\epsilon = 1.5 \frac{\epsilon}{km}$ . To allow for a visualization given the three objectives of interest, we separately investigate the trade-off between the primary objective, and each of the secondary objectives. Hence, Fig. 7a (Fig. 7b) plots the booked passenger km against the pooling rate (profit) achieved. Each data point corresponds to the result of applying a certain policy with a certain cost parameter value  $\epsilon$ . For now, we only consider Policy S (blue points). The different shades of blue encode the value of  $\epsilon$ , with the lightest indicating  $\epsilon = 0 \frac{\epsilon}{km}$  and the darkest indicating  $\epsilon = 1.5 \frac{\epsilon}{km}$ . Starting with  $\epsilon = 0 \frac{\epsilon}{km}$ , we observe that increasing the cost parameter sacrifices booked passenger km but steadily improves the pooling rate. In contrast, the relation with profit is non-monotonous. Steady improvement can only be observed for small cost parameter values, with  $\epsilon = 0.75 \frac{\epsilon}{km}$  yielding maximum profit. From this value onward, profit deteriorates. Underlying the observed parabolic profit trend are two inverse effects: On the one hand, increasing  $\epsilon$  leads to higher revenue per order and lower cost per order since prices for rides with high marginal vehicle km are higher such that less of these rides are ordered and, those that are, yield a higher revenue. On the other hand, these improvements in a relative way are accompanied by a declining absolute number of orders. For low  $\epsilon$  values, the positive relative effect prevails causing revenues to increase. For higher values, it is increasingly outweighed by the negative absolute effect.



Fig. 7 Solution quality depending on cost parameter

#### 5.4 Benefit of Anticipation

To measure the benefit of computing anticipatory approximations for the displaced booked passenger km  $\Delta \tilde{V}(s, c, o)$  and the marginal vehicle km  $\Delta \tilde{F}(s, c, o)$ , we benchmark the (fully) anticipatory policy against the two partly myopic policies (see Section 5.1.2). The results are also visualized in Fig. 7.

To isolate the benefit of the anticipatory approximation of displaced booked passenger km  $\Delta \tilde{V}(s, c, o)$ , we consider Policy S-MD, which assumes  $\Delta \tilde{V}(s, c, o) = 0$  in general but still features the same anticipatory approximation of  $\Delta \tilde{F}(s, c, o)$  as Policy S. Comparing both policies, we find that Policy S-MD does not yield any pareto-efficient solution. E.g., for  $\epsilon = 0.5 \frac{\epsilon}{km}$ , Policy S outperforms Policy S-MD by 3%, 6%, and 14% in terms of booked passenger km, pooling rate, and profit, respectively. This shows that an anticipatory approximation of  $\Delta \tilde{V}(s, c, o)$  raises the solution quality in general, i.e., independent from the value of parameter  $\epsilon$ . The reason is the price penalty for those rides that cause high displacement. Hence, if a customer still orders such a ride, the revenue is high, which positively impacts profit. If the customer orders an alternative ride that causes less displacement because the ride is shifted into a period with less demand, only the booked passenger km improve. However, if less displacement occurs due to better pooling, the pooling rate also improves.

Analogously, we can analyze the impact of using an anticipatory approximation of the marginal vehicle km  $\Delta \tilde{F}(s, c, o)$  (Policy S) compared to the myopic approximation, which draws on the

insertion cost into the tentative route plan (Policy S-MM). With increasing cost parameter values  $\epsilon$ , we observe that the shift of solution quality away from booked passenger km toward pooling rate and profit is much more pronounced than for Policy S. E.g., for  $\epsilon = 5$ , Policy S-MM yields 20% less booked passenger km but profit increases by 27% and the pooling rate improves by 4%.



Fig. 8 Impact of fully anticipatory ride evaluation on the average price multiplier

The reason behind this performance is that the myopic approximation systematically overestimates the true marginal vehicle km. This effect is fundamental to i-DMVRPs and has been first described extensively in Fleckenstein et al. (2024b). The overestimation error arises because the myopic approximation does not account for consolidation with future rides and, hence, is strongest in the beginning of the booking horizon. Fig. 8, which plots the average price multiplier over the booking horizon, shows that this overestimation error translates into higher prices, especially for early requests. Interestingly, Policy S-MM still yields pareto-efficient solutions. However, its price increases compared to Policy S can be viewed as an undesirable markup since prices rise due to a systematic overestimation of the true marginal cost. This is especially problematic since it creates disparities between early requests and late requests (Andrejszki and Török, 2018).

# 6 Managerial Insights and Future Research Directions

In this work, we presented the first approach to sustainable dynamic pricing for rural SMOD systems. Relative to existing research, we bridge the gap between transportation economics, where the transfer of marginal cost pricing to SMOD systems is proposed but not practically implemented, and operations research, where existing dynamic pricing approaches are purely profit-based and multi-objective problems are hardly considered. Having thoroughly derived the relevant objectives in practice by means of multi-attribute decision analysis, we modeled the providers' sequential decision problem as a constrained MDP, with a particular focus on the dynamic pricing subproblem. The model maximizes the demand served while ensuring that any ride's price is at least equal to its marginal cost, i.e., the provider's variable routing cost and external cost due

to emissions. With this structure, undesirable monopoly markups are avoided. To enable this model's application, we also proposed an anticipatory post-decision rollout algorithm, which evaluates rides regarding their expected displacement of future demand and expected marginal vehicle km. We comprehensively validated our approach using a real-world data set provided by our industry partner FLEXIBUS.

In this section, we summarize the key managerial insights (Section 6.1), derive recommendations for practice (Section 6.2), address limitations (Section 6.3), and sketch future research directions (Section 6.4).

## 6.1 Managerial Insights

- Favorable performance from the provider's perspective: Sustainable dynamic pricing shows superior performance in balancing the relevant objectives in line with the preferences of municipal authorities. We find the following results in the base scenario: Compared to static pricing, which is the real-world status-quo, a small loss in booked passenger km (-4%) leads to an 11% improvement in the pooling rate and a 44% improvement in profit. Compared to profit-based dynamic pricing, which is the state-of-the-art in the operations research literature, the provider obtains less profit (-40%) but collects 36% more booked passenger km, and the pooling rate is 3% higher. This confirms earlier theoretical findings about the potential severity of monopoly markups resulting from profit-based dynamic pricing (Bahamonde-Birke et al., 2021). In turn, sustainable dynamic pricing guarantees that any offer includes a least one ride at the lowest possible price point above its marginal cost. Thereby, given an accurate approximation of marginal vehicle km, monopoly markups are entirely avoided.
- Small impact of selectiveness: We observe hardly any performance difference between the non-selective variant and the selective variant of sustainable dynamic pricing. This means that providers can choose whether they prefer rejecting unfavorable requests (selective) or pricing them out (non-selective).
- Weak demand displacement: Our results point toward demand displacement being rather weak in general. The likely root cause for this is that for reaching an acceptance rate above 80%, the required supply-demand ratio is high enough that supply is not restrictive for most requests.
- Favorable performance from the customer's perspective: Considering the base scenario, sustainable dynamic pricing achieves a similar request acceptance rate compared to static pricing. Further, abandonment rate and average price multiplier increase moderately, and 80% of offers include at least one ride at the base price. Customers will likely perceive these changes in offers due to sustainable dynamic pricing as moderate, while the behavior of profit-based dynamic pricing would probably be seen as overly aggressive.
- Mitigation of undesirable disparities between different groups of requests: We do not observe indications for an increase in discriminatory behavior of sustainable dynamic pricing

compared to the benchmarks. Most prominently, short requests face a lower acceptance rate and offer set size yet also lower price multipliers (vice versa for long requests). Some of the disparities shown by the benchmarks regarding same-day requests and long requests are even mitigated by sustainable dynamic pricing.

- Shaping of served demand: Sustainable dynamic pricing successfully prices out demand with insufficient willingness-to-pay relative to marginal cost. Also, time flexible customers can order cheaper rides.
- Redistribution of performance due to the cost parameter: We find that higher cost parameter values shift performance from the primary objective to the secondary objectives. Hence, they lead to monotonically decreasing booked passenger km and an increasing pooling rate. However, the relation with profit is non-monotonous because profit first increases for small cost parameter values driven by a higher revenue and lower cost per order and then decreases again due to the lost demand.
- Benefits of anticipatory ride evaluation: Anticipation is essential for both the approximation of demand displacement and marginal vehicle km. The consideration of displaced booked passenger km yields performance benefits of around 10%. An anticipatory marginal vehicle km approximation also prevents overestimation errors. These are problematic because they cause undesirable markups and a systematic discrimination against advance requests, which then face higher prices.

## 6.2 Recommendations

- Adoption of sustainable dynamic pricing with anticipatory ride evaluation: Our results provide strong evidence for the superiority of sustainable dynamic pricing over static pricing and profit-based dynamic pricing. Therefore, we recommend its adoption in practice for publicly funded rural SMOD services together with anticipatory ride evaluation, which is superior compared to myopic algorithms.
- Integration into existing static pricing schemes: By design, our sustainable dynamic pricing approach can be applied together with an arbitrary scheme for the base prices. Thus, as demonstrated in our computational study, we encourage providers to set base prices equal to existing static pricing schemes. Thereby, they can ensure that the change in the pricing system is perceived as small as possible, especially given that, for realistic cost parameters, a vast majority of offers still includes at least one ride at the base price.
- Stressing the advantages for customers: Customers may have a skeptical stance toward dynamic pricing in general (e.g., Friesen et al., 2024, Schlereth et al., 2018). Therefore, its introduction should be accompanied by a communication strategy. We recommend highlighting the benefits for the individual customer:
  - Groups of customers with compatible OD pairs (families, neighbors, company employees) can jointly place a single request, increasing the chance of riding at the base price

compared to individual requests. A similar indirect effect occurs when customers convince their social environment to become SMOD service users because this reduces the marginal vehicle km of their own rides.

- Customers can still book cheap rides if they are flexible regarding their desired pick-up or drop-off time. Conversely, unfavorable requests that would be rejected under static pricing are more likely to receive an offer, albeit at a higher price.
- Compared to static pricing, the acceptance rate of same-day requests and ad-hoc requests increases, which makes the system more reliable if customers need a ride at short notice.
- The avoidance of monopoly markups and the adherence to marginal cost pricing ensures that every individual offer is adequate and fair.
- Contractual or regulatory enforcement: Since sustainable dynamic pricing leads to smaller profit compared to profit-based dynamic pricing, the provider should be obliged to apply it, including an accurate ride evaluation algorithm to avoid "hidden" markups by erroneous marginal vehicle km approximation. Similar price regulation is commonly applied in markets with natural monopolies, e.g., the district heating market (Billerbeck et al., 2023). In the context of SMOD, it may appear unusual to enforce a certain price calculation method instead of directly setting certain system performance requirements. In fact, current legislation in Germany states that municipal authorities can impose such a requirement in the form of a minimum pooling rate (Deutscher Bundestag, 2021). However, Anzenhofer et al. (2025) show that such regulation is problematic for both providers and authorities since it is difficult to adhere to it and requirements must be adjusted based on extensive computational experiments under potentially limited data availability.

# 6.3 Limitations

While we find very promising results, there are some limitations in both our approach for sustainable dynamic pricing and its analysis in the computational study. We discuss them in the following:

## • Approach:

- No consideration of alternative modes: In our model of the dynamic pricing subproblem, prices depend only on the future impact of the different rides within the SMOD system itself. This means that alternative modes and the associated side-impact are not considered explicitly. E.g., short rides tend to be offered at lower price multipliers, which may encourage the cannibalization of active modes, such as walking or cycling (Rich, 2024). Similarly, existing scheduled public transport services could be overly cannibalized if they are not explicitly considered in the pricing decisions (Lu et al., 2024).
- *Scalability of ride evaluation algorithm*: Since our ride evaluation algorithm is tailored to rural instances, its scalability is limited due to the high computational effort. While
sustainable dynamic pricing is generally transferable to urban settings, this would require a scalable ride evaluation algorithm.

- Computational study:
  - Limited generalizability to other systems: For validating our approach, we use a realworld data set provided by FLEXIBUS. While the high-level findings should generalize well, some of the observations, especially regarding the characteristics of offers and orders, may be specific to our data set.
  - Lack of reliable data on customers' time preferences and price sensitivity: Since FLEX-IBUS applies a first-come-first-served policy with static prices, the data set does not provide reliable information on price sensitivities and customers' time preferences. Hence, some observations could result from assumptions about the choice model.
  - No consideration of long-term impact: In our computational study, we do not consider the long-term side-effects of sustainable dynamic pricing. These include possible changes in the demand structure such as long-term modal shift (Kaddoura et al., 2021, Kaddoura et al., 2020b) or strategic behavior by customers (Anzenhofer and Fleckenstein, 2024).

### 6.4 Future Research Directions

Finally, we discuss potential avenues for future research resulting from the findings and limitations of our work:

- Exploration of customer preferences and price sensitivity: Since rural SMOD services are still an emerging mode of transport, there is a general lack of demand data. Moreover, most existing systems use static pricing, which yields no data at all about price sensitivities. Hence, we suggest actively and systematically exploring customer choice behavior to obtain this data (Jain et al., 2024, Te Morsche et al., 2019). This would allow the data-based estimation of discrete choice models and improve the accuracy of computational analyses. Likewise, the long-term behavior of customers should be thoroughly investigated to allow for the consideration of side-effects, especially regarding modal shift.
- Consideration of alternative modes: Since we present the first approach for sustainable dynamic pricing, there is potential for extending it in different ways. We believe that the most important extension is the consideration of alternative modes. E.g., recent work by Lu et al. (2024) presents promising results regarding the rejection of requests with a good alternative transport option (walking, cycling, or scheduled public transport). Combining this approach with sustainable dynamic pricing would result in prices also reflecting the sustainability of available alternative modes.
- **Ride evaluation for large systems**: The most important limitation of our approach lies in its scalability to larger instances due to the computationally intensive ride evaluation. Hence, future research could develop faster ride evaluation algorithms that allow the approach to be transferred to (sub-)urban systems.

## Acknowledgment

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.

#### References

- Al-Kanj, L., Nascimento, J., & Powell, W. B. (2020). Approximate dynamic programming for planning a ride-hailing system using autonomous fleets of electric vehicles. *European Jour*nal of Operational Research, 284(3), 1088-1106. https://doi.org/10.1016/j.ejor.2020.01.033.
- Alonso-Mora, J., Samaranayake, S., Wallar, A., Frazzoli, E., & Rus, D. (2017). On-demand highcapacity ride-sharing via dynamic trip-vehicle assignment. *Proceedings of the National Academy of Sciences of the United States of America*, 114(3), 462-467. https://doi.org/10.1073/pnas.1611675114.
- Altman, E. (1999). Constrained Markov Decision Processes. New York, NY: Routledge.
- Andrejszki, T., & Török, Á. (2018). New pricing theory of intelligent flexible transportation. *Transport*, 33(1), 69-76. https://doi.org/10.3846/16484142.2015.1056828.
- Anzenhofer, F., & Fleckenstein, D. (2024). Extended booking horizons in rural shared mobilityon-demand systems: Insights and implications for demand management. Working paper, University of Augsburg.
- Anzenhofer, F., Fleckenstein, D., Klein, R., & Steinhardt, C. (2024). Analyzing the impact of demand management in rural shared mobility-on-demand systems. Working paper, University of Augsburg. http://dx.doi.org/10.2139/ssrn.4682056.
- Anzenhofer, F., Schmidbaur, S., Klein, R., & Steinhardt, C. (2025). The potential of governmental regulation on shared mobility-on-demand systems. *Transportation Research Part A: Policy* and Practice, 192, 104360. https://doi.org/10.1016/j.tra.2024.104360.
- 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.
- Atasoy, B., Ikeda, T., Song, X., & Ben-Akiva, M. E. (2015). The concept and impact analysis of a flexible mobility on demand system. *Transportation Research Part C: Emerging Technol*ogies, 56, 373-392. https://doi.org/10.1016/j.trc.2015.04.009.
- Bahamonde-Birke, F. J., Goletz, M., & Ettema, D. (2021). The provision of mobility as a service with autonomous vehicles. The necessity of regulatory schemes for a natural monopoly. *Research in Transportation Economics*, 90, 100993. https://doi.org/10.1016/j.retrec.2020.100993.
- Bento, N., & Gianfrate, G. (2020). Determinants of internal carbon pricing. *Energy Policy*, *143*, 111499. https://doi.org/10.1016/j.enpol.2020.111499.
- Bertsekas, D. P., Tsitsiklis, J. N., & Wu, C. (1997). Rollout algorithms for combinatorial optimization. *Journal of Heuristics*, *3*, 245-262. https://doi.org/10.1023/A:1009635226865.
- Billerbeck, A., Breitschopf, B., Winkler, J., Bürger, V., Köhler, B., Bacquet, A., ... & Ragwitz, M. (2023). Policy frameworks for district heating: A comprehensive overview and analysis of regulations and support measures across Europe. *Energy Policy*, 173, 113377. https://doi.org/10.1016/j.enpol.2022.113377.
- Bimpikis, K., Candogan, O., & Saban, D. (2019). Spatial pricing in ride-sharing networks. Operations Research, 67(3), 744-769. https://doi.org/10.1287/opre.2018.1800.
- Burghard, U., & Scherrer, A. (2022). Sharing vehicles or sharing rides Psychological factors influencing the acceptance of carsharing and ridepooling in Germany. *Energy Policy*, 164, 112874. https://doi.org/10.1016/j.enpol.2022.112874.
- Byrne, L., Bach, V., & Finkbeiner, M. (2021). Urban transport assessment of emissions and resource demand of climate protection scenarios. *Cleaner Environmental Systems*, 2, 100019. https://doi.org/10.1016/j.cesys.2021.100019.
- Carroll, P., Caulfield, B., & Ahern, A. (2019). Measuring the potential emission reductions from a shift towards public transport. *Transportation Research Part D: Transport and Environment*, 73, 338-351. https://doi.org/10.1016/j.trd.2019.07.010.
- Chen, H., Jiao, Y., Qin, Z., Tang, X., Li, H., An, B., Hongtu, Z., & Ye, J. (2019). InBEDE: Integrating contextual bandit with TD learning for joint pricing and dispatch of ride-hailing

platforms. In 2019 IEEE International Conference on Data Mining (ICDM), 61-70. https://doi.org/10.1109/ICDM.2019.00016.

- Chen, Z., Zhang, H., Li, H., & Webster, S. (2024). Multi-Objective Assortment Optimization: Profit, Risk, Customer Utility, and Beyond. *Foundations and Trends® in Technology, Information and Operations Management*, 18(1), 103-115. http://dx.doi.org/10.1561/0200000114-5.
- Cohen, M. C., Miao, S., & Wang, Y. (2024). Dynamic pricing with fairness constraints. Working paper, McGill University. http://dx.doi.org/10.2139/ssrn.3930622.
- Currie, G., & Fournier, N. (2020). Why most DRT/Micro-Transits fail What the survivors tell us about progress. *Research in Transportation Economics*, *83*, 100895. https://doi.org/10.1016/j.retrec.2020.100895.
- Dauer, A., Dias, T. G., de Sousa, J. P., & de Athayde Prata, B. (2024). Configurations and features of demand responsive transports. *Transportation Research Procedia*, 78, 71-78. https://doi.org/10.1016/j.trpro.2024.02.010.
- Davis, J., Gallego, G., & Topaloglu, H. (2013). Assortment planning under the multinomial logit model with totally unimodular constraint structures. Working paper, Cornell University.
- De Jong, W., Vogels, J., van Wijk, K., & Cazemier, O. (2011). The key factors for providing successful public transport in low-density areas in The Netherlands. *Research in Transportation Business & Management*, *2*, 65-73. https://doi.org/10.1016/j.rtbm.2011.07.002.
- Deutscher Bundestag (2021). Personenbeförderungsgesetz § 50 Gebündelter Bedarfsverkehr. Retrieved from https://www.gesetze-im-internet.de/pbefg/\_\_50.html. Last accessed: 22/11/2024.
- Eliasson, J. (2021). Efficient transport pricing Why, what, and when?. *Communications in Transportation Research*, *1*, 100006. https://doi.org/10.1016/j.commtr.2021.100006.
- Elting, S., & Ehmke, J. F. (2021). Potential of shared taxi services in rural areas A case study. *Transportation Research Procedia*, *52*, 661-668. https://doi.org/10.1016/j.trpro.2021.01.079.
- Fleckenstein, D., Klein, R., Klein, V., & Steinhardt, C. (2024a). On the concept of opportunity cost in integrated demand management and vehicle routing. *Transportation Science*, online *first*. https://doi.org/10.1287/trsc.2024.0644.
- Fleckenstein, D., Klein, R., Klein, V., & Steinhardt, C. (2024b). From approximation error to optimality gap – Explaining the performance impact of opportunity cost approximation in integrated demand management and vehicle routing. Working paper, University of Augsburg. https://doi.org/10.48550/arXiv.2412.13851.
- Fleckenstein, D., Klein, R., & Steinhardt, C. (2023). Recent advances in integrating demand management and vehicle routing: A methodological review. *European Journal of Operational Research*, 306(2), 499-518. https://doi.org/10.1016/j.ejor.2022.04.032.
- Friesen, M., Albers, S., Wajerski, J., & Laeremans, J. (2024). The diffusion of dynamic pricing in the European parking industry. *Journal of Revenue and Pricing Management*, 23(1), 3-13. https://doi.org/10.1057/s41272-023-00439-5.
- Gallego, G., & Topaloglu, H. (2019). Revenue management and pricing analytics. New York, NY: Springer. https://doi.org/10.1007/978-1-4939-9606-3.
- García-Afonso, Ó. (2023). Impact of powertrain electrification on the overall CO2 emissions of intercity public bus transport: Tenerife Island test case. *Journal of Cleaner Production*, 412, 137365. https://doi.org/10.1016/j.jclepro.2023.137365.
- Goodson, J. C., Thomas, B. W., & Ohlmann, J. W. (2017). A rollout algorithm framework for heuristic solutions to finite-horizon stochastic dynamic programs. *European Journal of Operational Research*, 258(1), 216-229. https://doi.org/10.1016/j.ejor.2016.09.040.
- Haferkamp, J. (2024). Design of multi-optional pickup time offers in ride-sharing systems. *EURO Journal on Transportation and Logistics*, 100134. https://doi.org/10.1016/j.ejtl.2024.100134.

- Haferkamp, J., & Ehmke, J. F. (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.
- Hahn, A., Pakusch, C., & Stevens, G. (2023). The impact of service expansion on modal shift from private car to public transport. A quantitative analysis in the Bonn/Rhein-Sieg area, Germany. *Journal of Urban Mobility*, 4, 100064. https://doi.org/10.1016/j.urbmob.2023.100064.
- Haliem, M., Mani, G., Aggarwal, V., & Bhargava, B. (2021). A distributed model-free ride-sharing approach for joint matching, pricing, and dispatching using deep reinforcement learning. *IEEE Transactions on Intelligent Transportation Systems*, 22(12), 7931-7942. https://doi.org/10.1109/TITS.2021.3096537.
- Hansson, J., Pettersson, F., Svensson, H., & Wretstrand, A. (2019). Preferences in regional public transport: a literature review. *European Transport Research Review*, 11(1), 1-16. https://doi.org/10.1186/s12544-019-0374-4.
- Heger, J., & Klein, R. (2024). Assortment optimization: A systematic literature review. OR Spectrum, 46, 1099-1161. https://doi.org/10.1007/s00291-024-00752-4.
- Heinitz, F. (2022). Sustainable development assessment of incentive-driven shared on-demand mobility systems in rural settings. *European Transport Research Review*, 14(1), 1-16. https://doi.org/10.1186/s12544-022-00565-y.
- Heitmann, R. J. O., Soeffker, N., Klawonn, F., Ulmer, M. W., & Mattfeld, D. C. (2024). Accelerating value function approximations for dynamic dial-a-ride problems via dimensionality reductions. *Computers & Operations Research*, 167, 106639. https://doi.org/10.1016/j.cor.2024.106639.
- Heitmann, R. J. O., Soeffker, N., Ulmer, M. W., & Mattfeld, D. C. (2023). Combining value function approximation and multiple scenario approach for the effective management of ridehailing services. *EURO Journal on Transportation and Logistics*, 12, 100104. https://doi.org/10.1016/j.ejtl.2023.100104.
- Hensher, D. A. (2020). Bus transport: Demand, economics, contracting, and policy. Amsterdam: Elsevier.
- Ho, S. C., Szeto, W. Y., Kuo, Y. H., Leung, J. M., Petering, M., & Tou, T. W. (2018). A survey of dial-a-ride problems: Literature review and recent developments. *Transportation Research Part B: Methodological*, 111, 395-421. https://doi.org/10.1016/j.trb.2018.02.001.
- Hörcher, D., & Graham, D. (2020a). MaaS economics: Should we fight car ownership with subscriptions to alternative modes?. *Economics of Transportation*, 22, 100167. https://doi.org/10.1016/j.ecotra.2020.100167.
- Hörcher, D., & Graham, D. (2020b). Pricing and efficient public transport supply in a mobility as a service context. Working paper, Imperial College London.
- Hörcher, D., & Tirachini, A. (2021). A review of public transport economics. *Economics of Transportation*, 25, 100196. https://doi.org/10.1016/j.ecotra.2021.100196.
- Holmgren, J. (2007). Meta-analysis of public transport demand. *Transportation Research Part A: Policy and Practice*, *41*(10), 1021-1035. https://doi.org/10.1016/j.tra.2007.06.003.
- Hungerländer, P., Maier, K., Pachatz, V., & Truden, C. (2021). Improving sharing rates of a diala-ride problem implemented for an Austrian mobility provider. *Transportation Research Procedia*, 52, 525-532. https://doi.org/10.1016/j.trpro.2021.01.062.
- Imhof, S., & Mayer, H. (2024). How social innovations emerge in a rigid regulatory context: the case of demand responsive transport in Switzerland. *Transportation Planning and Technol*ogy, 1-24. https://doi.org/10.1080/03081060.2024.2350535.
- Jain, L., Li, Z., Loghmani, E., Mason, B., & Yoganarasimhan, H. (2024). Effective adaptive exploration of prices and promotions in choice-based demand models. *Marketing Science*, 43(5), 1002-1030. https://doi.org/10.1287/mksc.2023.0322.

- Jochem, P., Frankenhauser, D., Ewald, L., Ensslen, A., & Fromm, H. (2020). Does free-floating carsharing reduce private vehicle ownership? The case of SHARE NOW in European cities. *Transportation Research Part A: Policy and Practice*, 141, 373-395. https://doi.org/10.1016/j.tra.2020.09.016.
- Jung, J., Jayakrishnan, R., & Park, J. Y. (2016). Dynamic shared-taxi dispatch algorithm with hybrid-simulated annealing. *Computer-Aided Civil and Infrastructure Engineering*, *31*(4), 275-291. https://doi.org/10.1111/mice.12157.
- Kaddoura, I., Bischoff, J., & Nagel, K. (2020a). Towards welfare optimal operation of innovative mobility concepts: External cost pricing in a world of shared autonomous vehicles. *Transportation Research Part A: Policy and Practice*, 136, 48-63. https://doi.org/10.1016/j.tra.2020.03.032.
- Kaddoura, I., Kickhöfer, B., Neumann, A., & Tirachini, A. (2015). Optimal public transport pricing: Towards an agent-based marginal social cost approach. *Journal of Transport Economics and Policy*, 49(2), 200-218.
- Kaddoura, I., Leich, G., & Nagel, K. (2020b). The impact of pricing and service area design on the modal shift towards demand responsive transit. *Procedia Computer Science*, 170, 807-812. https://doi.org/10.1016/j.procs.2020.03.152.
- Kaddoura, I., Leich, G., Neumann, A., & Nagel, K. (2021). From today's ride-sharing services to future mobility concepts: A simulation study for urban and rural areas. Working paper, Technische Universität Berlin.
- Karaenke, P., Schiffer, M., & Waldherr, S. (2023). On the benefits of ex-post pricing for ridepooling. *Transportation Research Part C: Emerging Technologies*, 155, 104290. https://doi.org/10.1016/j.trc.2023.104290.
- Ke, Z., & Qian, S. (2023). Leveraging ride-hailing services for social good: Fleet optimal routing and system optimal pricing. *Transportation Research Part C: Emerging Technologies*, 155, 104284. https://doi.org/10.1016/j.trc.2023.104284.
- Keeney, R. L., & Raiffa, H. (1993). Decisions with multiple objectives: Preferences and value tradeoffs. Cambridge: Cambridge University Press.
- Klein, R., Koch, S., Steinhardt, C., & Strauss, A. K. (2020). A review of revenue management: Recent generalizations and advances in industry applications. *European Journal of Operational Research*, 284(2), 397-412. https://doi.org/10.1016/j.ejor.2019.06.034.
- Klein, V., & Steinhardt, C. (2023). Dynamic demand management and online tour planning for same-day delivery. *European Journal of Operational Research*, 307(2), 860-886. https://doi.org/10.1016/j.ejor.2022.09.011.
- Koch, S., & Klein, R. (2020). Route-based approximate dynamic programming for dynamic pricing in attended home delivery. *European Journal of Operational Research*, 287(2), 633-652. https://doi.org/10.1016/j.ejor.2020.04.002.
- Köhler, C., Campbell, A. M., & Ehmke, J. F. (2024). Data-driven customer acceptance for attended home delivery. OR Spectrum, 46(2), 295-330. https://doi.org/10.1007/s00291-023-00712-4.
- Lang, M. A. K., Cleophas, C., & Ehmke, J. F. (2021). Multi-criteria decision making in dynamic slotting for attended home deliveries. *Omega*, 102, 102305. https://doi.org/10.1016/j.omega.2020.102305.
- Laws, R. (2008). DRT schemes in England and Wales and considerations for their future. In *87th Annual Meeting of the Transportation Research Board*, 08-0892.
- Letmathe, P., & Paegert, M. (2024). External effects of urban automated vehicles on sustainability. *Journal of Cleaner Production*, 434, 140257. https://doi.org/10.1016/j.jclepro.2023.140257.
- Lotfi, S., & Abdelghany, K. (2022). Ride matching and vehicle routing for on-demand mobility services. *Journal of 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. https://doi.org/10.48550/arXiv.2306.05851.
- Lu, C., Tiwari, S., Nassir, N., & Nagel, K. (2024). Efficient operation of demand-responsive transport (DRT) systems: Active requests rejection. *Procedia Computer Science*, 238, 81-90. https://doi.org/10.1016/j.procs.2024.06.001.
- Ma, H., Fang, F., & Parkes, D. C. (2022). Spatio-temporal pricing for ridesharing platforms. *Operations Research*, *70*(2), 1025-1041. https://doi.org/10.1287/opre.2021.2178.
- Mortazavi, A., Ghasri, M., & Ray, T. (2024). Integrated demand responsive transport in lowdemand areas: A case study of Canberra, Australia. *Transportation Research Part D: Transport and Environment*, 127, 104036. https://doi.org/10.1016/j.trd.2023.104036.
- Mounce, R., Beecroft, M., & Nelson, J. D. (2020). On the role of frameworks and smart mobility in addressing the rural mobility problem. *Research in Transportation Economics*, *83*, 100956. https://doi.org/10.1016/j.retrec.2020.100956.
- Ni, L., Sun, B., Wang, S., & Tsang, D. H. (2021). Dynamic pricing mechanism design for electric mobility-on-demand systems. *IEEE Transactions on Intelligent Transportation Systems*, 23(8), 11361-11375. https://doi.org/10.1109/TITS.2021.3103199.
- Nobis, C. & Kuhnimhof, T. (2018). Mobilität in Deutschland Ergebnisbericht. Retrieved from https://www.mobilitaet-in-deutschland.de/archive/pdf/MiD2017\_Ergebnisbericht.pdf. Last accessed: 23/12/2024.
- OSRM (n.d.). Open source routing machine. https://project-osrm.org/. Last accessed: 23/12/2024.
- Pillac, V., Gendreau, M., Guéret, C., & Medaglia, A. L. (2013). A review of dynamic vehicle routing problems. *European Journal of Operational Research*, 225(1), 1-11. https://doi.org/10.1016/j.ejor.2012.08.015.
- Poltimäe, H., Rehema, M., Raun, J., & Poom, A. (2022). In search of sustainable and inclusive mobility solutions for rural areas. *European Transport Research Review*, 14(1), 13. https://doi.org/10.1186/s12544-022-00536-3.
- Psaraftis, H. N., Wen, M., & Kontovas, C. A. (2016). Dynamic vehicle routing problems: Three decades and counting. *Networks*, 67(1), 3-31. https://doi.org/10.1002/net.21628.
- Purvis, B., Mao, Y., & Robinson, D. (2019). Three pillars of sustainability: In search of conceptual origins. *Sustainability Science*, 14, 681-695. https://doi.org/10.1007/s11625-018-0627-5.
- Puterman, M. L. (2014). Markov decision processes: Discrete stochastic dynamic programming. New York, NY: John Wiley & Sons. https://doi.org/10.1002/9780470316887.
- Qiu, H., Li, R., & Zhao, J. (2018). Dynamic pricing in shared mobility on demand service. Working paper, Massachusetts Institute of Technology.
- Rammohan, S., Marathe, R. R., & Sudarsanam, N. (2024). Recent advancements in revenue management of taxi services: A systematic review and research agenda. *Management Review Quarterly*, 74(2), 1029-1055. https://doi.org/10.1007/s11301-023-00326-8.
- Rich, J. (2024). Let's walk! The fallacy of urban first-and last-mile public transport. *Transportation, online first.* https://doi.org/10.1007/s11116-024-10505-5.
- Rios, B. H. O., Xavier, E. C., Miyazawa, F. K., Amorim, P., Curcio, E., & Santos, M. J. (2021). Recent dynamic vehicle routing problems: A survey. *Computers & Industrial Engineering*, 160, 107604. https://doi.org/10.1016/j.cie.2021.107604.
- Roijers, D. M., Vamplew, P., Whiteson, S., & Dazeley, R. (2013). A survey of multi-objective sequential decision-making. *Journal of Artificial Intelligence Research*, 48, 67-113. https://doi.org/10.1613/jair.3987.
- Saharan, S., Bawa, S., & Kumar, N. (2020). Dynamic pricing techniques for intelligent transportation system in smart cities: A systematic review. *Computer Communications*, 150, 603-625. https://doi.org/10.1016/j.comcom.2019.12.003.

- Schasché, S. E., Sposato, R. G., & Hampl, N. (2022). The dilemma of demand-responsive transport services in rural areas: Conflicting expectations and weak user acceptance. *Transport Policy*, 126, 43-54. https://doi.org/10.1016/j.tranpol.2022.06.015.
- Schasché, S. E., Wankmüller, C., & Hampl, N. (2023). Understanding the behavioral intention of the rural population to use demand-responsive transport services. *Transportation Research Interdisciplinary Perspectives*, 22, 100984. https://doi.org/10.1016/j.trip.2023.100984.
- Schilde, M., Doerner, K. F., & Hartl, R. F. (2011). Metaheuristics for the dynamic stochastic diala-ride problem with expected return transports. *Computers & Operations Research*, 38(12), 1719-1730. https://doi.org/10.1016/j.cor.2011.02.006.
- Schlereth, C., Skiera, B., & Schulz, F. (2018). Why do consumers prefer static instead of dynamic pricing plans? An empirical study for a better understanding of the low preferences for timevariant pricing plans. *European Journal of Operational Research*, 269(3), 1165-1179. https://doi.org/10.1016/j.ejor.2018.03.033.
- Schwedes, O. (2021). Öffentliche Mobilität. Voraussetzungen für eine menschengerechte Verkehrsplanung. Wiesbaden: Springer. https://doi.org/10.1007/978-3-658-32106-2.
- Sharif Azadeh, S., Atasoy, B., Ben-Akiva, M. E., Bierlaire, M., & Maknoon, M. Y. (2022). Choice-driven dial-a-ride problem for demand responsive mobility service. *Transportation Research Part B: Methodological*, 161, 128-149. https://doi.org/10.1016/j.trb.2022.04.008.
- Sieber, L., Ruch, C., Hörl, S., Axhausen, K. W., & Frazzoli, E. (2020). Improved public transportation in rural areas with self-driving cars: A study on the operation of Swiss train lines. *Transportation Research Part A: Policy and Practice*, 134, 35-51. https://doi.org/10.1016/j.tra.2020.01.020.
- Soeffker, N., Ulmer, M. W., & Mattfeld, D. C. (2022). Stochastic dynamic vehicle routing in the light of prescriptive analytics: A review. *European Journal of Operational Research*, 298(3), 801-820. https://doi.org/10.1016/j.ejor.2021.07.014.
- Sörensen, L., Bossert, A., Jokinen, J. P., & Schlüter, J. (2021). How much flexibility does rural public transport need? – Implications from a fully flexible DRT system. *Transport Policy*, 100, 5-20. https://doi.org/10.1016/j.tranpol.2020.09.005.
- Strauss, A. K., Klein, R., & Steinhardt, C. (2018). A review of choice-based revenue management: Theory and methods. *European Journal of Operational Research*, 271(2), 375-387. https://doi.org/10.1016/j.ejor.2018.01.011.
- Talluri, K. T., & Van Ryzin, G. J. (2004). The theory and practice of revenue management. New York, NY: Springer. https://doi.org/10.1007/b139000.
- Te Morsche, W., Puello, L. L. P., & Geurs, K. T. (2019). Potential uptake of adaptive transport services: An exploration of service attributes and attitudes. *Transport Policy*, 84, 1-11. https://doi.org/10.1016/j.tranpol.2019.09.001.
- Ulmer, M. W. (2020a). Dynamic pricing and routing for same-day delivery. *Transportation Science*, 54(4), 1016-1033. https://doi.org/10.1287/trsc.2019.0958.
- Ulmer, M. W. (2020b). Horizontal combinations of online and offline approximate dynamic programming for stochastic dynamic vehicle routing. *Central European Journal of Operations Research*, 28(1), 279-308. https://doi.org/10.1007/s10100-018-0588-x.
- Ulmer, M. W., Goodson, J. C., Mattfeld, D. C., & Hennig, M. (2019). Offline-online approximate dynamic programming for dynamic vehicle routing with stochastic requests. *Transportation Science*, 53(1), 185-202. https://doi.org/10.1287/trsc.2017.0767.
- Ulmer, M. W., Goodson, J. C., Mattfeld, D. C., & Thomas, B. W. (2020). On modeling stochastic dynamic vehicle routing problems. *EURO Journal on Transportation and Logistics*, 9(2), 100008. https://doi.org/10.1016/j.ejtl.2020.100008.
- Ulmer, M. W., Mattfeld, D. C., Hennig, M., & Goodson, J. C. (2016). A rollout algorithm for vehicle routing with stochastic customer requests. In *Logistics Management: Contributions* of the Section Logistics of the German Academic Association for Business Research, 2015, 217-227. https://doi.org/10.1007/978-3-319-20863-3\_16.

- Vansteenwegen, P., Melis, L., Aktaş, D., Montenegro, B. D. G., Vieira, F. S., & Sörensen, K. (2022). A survey on demand-responsive public bus systems. *Transportation Research Part C: Emerging Technologies*, 137, 103573. https://doi.org/10.1016/j.trc.2022.103573.
- VDV (2023). Linienbedarfsverkehr: zukunftsgerecht, integriert und nachfragegesteuert. Retrieved from https://www.vdv.de/positionspapier-linienbedarfsverkehr-maerz23.pdfx. Last accessed: 15/11/2024.
- Vickerman, R. (2024). The transport problem: The need for consistent policies on pricing and investment. *Transport Policy*, *149*, 49-58. https://doi.org/10.1016/j.tranpol.2024.02.009.
- Wang, Y., Shi, W., & Chen, Z. (2021a). Impact of ride-hailing usage on vehicle ownership in the United States. *Transportation Research Part D: Transport and Environment*, 101, 103085. https://doi.org/10.1016/j.trd.2021.103085.
- Wang, H., & Yang, H. (2019). Ridesourcing systems: A framework and review. Transportation Research Part B: Methodological, 129, 122-155. https://doi.org/10.1016/j.trb.2019.07.009.
- Wang, L., Zeng, L., Ma, W., & Guo, Y. (2021b). Integrating passenger incentives to optimize routing for demand-responsive customized bus systems. *IEEE Access*, 9, 21507-21521. https://doi.org/10.1109/ACCESS.2021.3055855.
- Wu, W., Zhu, Y., & Liu, R. (2024). Dynamic scheduling of flexible bus services with hybrid requests and fairness: Heuristics-guided multi-agent reinforcement learning with imitation learning. *Transportation Research Part B: Methodological*, 190, 103069. https://doi.org/10.1016/j.trb.2024.103069.
- Zhang, C., Xie, J., Wu, F., Gao, X., & Chen, G. (2020). Pricing and allocation algorithm designs in dynamic ridesharing system. *Theoretical Computer Science*, 803, 94-104. https://doi.org/10.1016/j.tcs.2019.05.045.
- Zwick, F., Kuehnel, N., & Axhausen, K. W. (2022). Review on theoretical assessments and practical implementations of ride-pooling. In 22nd Swiss Transport Research Conference (STRC 2022). https://doi.org/10.3929/ethz-b-000548675.

## **Appendix A: Notation**

$t \in \mathcal{T} = \{1, \dots, t^s, \dots, T\}$	Decision epoch
$t^s$	Start of the service horizon
$p_c, d_c \in \mathcal{H}$	Pick-up (drop-off) stop of request type <i>c</i>
$c \in C$	Customer request type
$i \in \mathcal{I}$	Request
$t_c \in \{t^s, \dots, T\}$	Desired time of request type <i>c</i>
m <sub>c</sub>	Number of passengers of request type <i>c</i>
b <sub>c</sub>	Base price of request type <i>c</i>
ρ <sub>c</sub>	Booked passenger km of request type <i>c</i>
$\tau_i \in \mathcal{T}$	Time of request for request <i>i</i>
$o\in \mathcal{O}_c$	Ride defined for request type <i>c</i>
$k\in \mathcal{K}_c$	Product defined for request type <i>c</i>
$n\in \mathcal{N}$	Price multiplier
r <sub>c,k</sub>	Price of product k for request type c
$j \in \mathcal{J}$	Order
$g \subseteq \mathcal{O}_c$	Offer set that can be presented to customer type <i>c</i>
$P_{c,k}(g)$	Probability of customer placing a request of type $c$ choosing product $k$ when presented offer set $g$
$v \in \mathcal{V}$	Vehicle
$Q_{v}$	Seat capacity of vehicle <i>v</i>
s <sub>t</sub>	Post-decision state at decision epoch t
C <sub>t</sub>	Set of confirmed but not yet fulfilled orders at decision epoch t
$\phi_t$	Route plan at decision epoch t
$a_t$	Action at decision epoch t
$g_t \in \mathcal{G}(s_{t-1},c)$	Demand 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 $o$
$ \begin{pmatrix} \phi_t(o_k) \end{pmatrix}_{k \in g_t} \in \\ \prod_{k \in g_t} \Phi(s_{t-1}, c, o_k) $	Vehicle routing decision at decision epoch $t$ for request type $c$ demand control decision $g_t$
$s_t^{\rm pre}$	Pre-decision state at decision epoch t
$\lambda_c^t$	Arrival rate of request type $c$ in stage $t$
s't	Interim state at decision epoch t
k <sub>jt</sub>	Product chosen by the customer placing order $j_t$
o <sub>jt</sub>	Ride chosen by the customer placing order $j_t$
$V_t(s_t)$	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 $o$

Table 7 Notation Markov decision process model

<i>x</i> <sub><i>o</i>,<i>n</i></sub>	Decision variable encoding whether ride $o$ is offered with price multiplier $n$
$\mathcal{O}^f \subseteq \mathcal{O}_c$	Set of feasible rides for request type <i>c</i>
E	Cost parameter consisting of variable routing cost and carbon cost per vehicle km
$\Delta F_t(s_{t-1},c,o)$	Expected marginal vehicle km c caused by selling ride $o$ to a customer placing a request of type $c$ in state $s_{t-1}$

### Table 8 Notation of dynamic pricing subproblem model

#### Table 9 Notation of solution algorithms

$\Delta \widetilde{V}_t(s_{t-1},c_{i_t},o)$	Approximation of displaced booked passenger km of an order by request type $c$ with fulfillment option $o$
$\Delta \widetilde{F}_t(s_{t-1},c_{i_t},o)$	Approximation of marginal vehicle km of an order by request type $c$ with fulfillment option $o$
$\omega\in\Omega$	Sample path
$B^{sta} = \alpha$	Static upper bound for the number of sample paths
$B^{dyn}$	Dynamic upper bound for the number of sample paths
β	Time budget for rollout algorithm per decision epoch
w	Estimate of runtime for simulating a single sample path
$j_t'(o)$	Hypothetical order for ride o
$(\mathcal{C}^{o,q},\phi^{o,q})$	Post-decision state resulting from a hypothetical order $j'_t(o)$ for ride $o$
$\widehat{V}_{o,q}$	Realized total booked passenger km for ride $o$ and the $q$ -th sample path
$\widehat{F}_{o,q}$	Realized total vehicle km for ride o and the q-th sample path
$\mathcal{K}_{c_{i_t}}'$	Set of feasible products for request $i_t$

#### Table 10 Notation of customer choice model

$l \in \mathcal{L}$	Customer segment
$S_l$	Consideration set of customer segment <i>l</i>
$\Delta_l^+ + \Delta_l^-$	Total time flexibility provided by segment <i>l</i>
ζ	Ranking function over fulfillment options
$\gamma_l$	Share of segment $l$ in the customer population



## Appendix B: Detailed Results Regarding the Characteristics of Offers and Orders

Fig. 9 Characteristics of offers and orders based on the arrival time of the request



Fig. 10 Characteristics of offers and orders based on the desired time of the request



Fig. 11 Characteristics of offers and orders based on the OD pair length of the request

## Appendix C: Sensitivity Regarding Hyperparameters

The post-decision rollout algorithm for ride evaluation consumes around 99,8% of the average runtime per decision epoch, which equals around 4.7 seconds in the base scenario. Due to this high computational effort, we now investigate the sensitivity of solution quality and runtime regarding the hyperparameters  $\alpha$  (maximum number of sample paths) and  $\beta$  (time budget for rollout per decision epoch). Considering Policy S, we conduct a grid search around the values (15,10) used in the base scenario. We test  $\alpha \in \{10,15,20\}$  and  $\beta \in \{10,15,20\}$ . Fig. 12a (Fig. 12b) shows the solution quality (runtime) for the 9 resulting settings relative to the base scenario.

Both on the runtime and on the solution quality, the impact of  $\alpha$  is smaller for small values of  $\beta$  since, in this case, the time budget is restrictive over a larger portion of the booking horizon, and vice versa. Regarding solution quality, there is a clear pattern for both  $\alpha$  and  $\beta$ :

For all objectives, a higher value of  $\alpha$  generally tends to improve the solution quality. This can be explained by a reduction in the stochastic approximation error. However, the marginal change in solution quality decreases, and may even become negative, especially for the secondary objectives. A potential explanation is that raising the static upper bound on the number of sample paths only leads to more sample paths being considered at the end of the booking horizon, where anticipation is hardly beneficial. Considering, the impact of  $\beta$ , we find that it particularly affects the booked passenger km. The reason is that a higher time budget allows more sample paths especially in the early phase of the booking horizon, where displacement effects are the strongest. As expected, this benefit seems to converge with higher values of  $\beta$ . In comparison, the secondary objectives are hardly affected because marginal vehicle km are not time-dependent.

Overall, we can conclude that the time budget must be sufficiently large. This is especially true for approximating displaced booked passenger km. Hence, a rough parameter tuning is required before applying our approach.



Fig. 12 Solution quality and runtime dependent on hyperparameter values

### **III** Conclusion

This dissertation covers theoretical, high-level considerations around i-DMVRPs in its first part (articles A1-A3) and application-oriented topics in rural SMOD in its second part (articles A4-A6). Over the course of the dissertation, i-DMVRPs are addressed with a focus on the demand management subproblem and the opportunity cost approximation. The latter is a key component of the popular decomposition-based solution concept.

Both parts differ in their target audience: The theoretical insights presented in the first part provide a foundation and guidance for algorithm selection and development. On the one hand, we consolidate research that up to now has been application-centric. This fosters knowledge transfer between application areas of i-DMVRPs. On the other hand, we provide new impetus for the further development of solution algorithms.

In the second part of the dissertation, the results are particularly relevant for practitioners in public transport who are involved in the introduction or the enhancement of demand management approaches in a rural SMOD system. Articles A4-A6 provide a readily applicable and innovative methodology for the preparatory analysis of the demand structure, for the sustainability-oriented selection of demand control policies at the strategic planning level, and for sustainable dynamic pricing at the operational planning level. By applying the methodology to real-world data form FLEXIBUS, an experienced German industry partner, we gain manifold insights and formulate recommendations for practice.

The remainder of this conclusion critically discusses the main results from an overarching perspective including their implications, limitations, and directions for future research. The discussion is organized according to the two main parts of the dissertation at hand, i.e., articles A1-A3 (first part) and articles A4-A6 (second part).

#### First part of the dissertation

The first part of this dissertation yields results that are valid for the entire family of i-DMVRPs and concern problem definition, modeling, and solution approaches.

• **Problem definition**: Article A1 shows how large the family of i-DMVRPs has become, comprising a variety of problem definitions and applications with the same basic structure. Despite this, there is a surprising homogeneity in the objective: The vast majority of works consider single-objective i-DMVRPs with the maximization of profit after fulfillment or closely related objectives like maximizing the number of orders. Future research should more frequently address the multi-objective nature of many i-DMVRPs in practice. This applies to AHD and SDD, where maximizing long-term market share (Lang et al., 2021) is also relevant. It is especially valid for MOD, where the all-inclusive objective in practice is often maximizing the different pillars of sustainability (Article A6). Another trend that can be observed across application areas is that fulfillment option definitions become more complex. In connection, a shift from simple accept/reject demand control toward availability control and

dynamic pricing occurs. While this results in more complex i-DMVRPs, it also makes services more customer-centric and allows for more sophisticated demand management.

- **Modeling**: Regarding modeling, we observe an increasing standardization with more and more authors using MDP models to formalize the considered i-DMVRP. In general, this is a positive development since a mathematically accurate documentation of the problem facilitates reproducibility, e.g., when implementing a simulator of the provider's operations and customers' behavior. Further, the analytical analysis or numerical analysis of MDP models yields actionable insights for the development of solution approaches. Articles A2 and A3 provide various contributions in this regard, namely generic model formulations and transformations, the interim state as a novel modeling instrument, proofs of opportunity cost properties, and explainability techniques. Still, the downside of MDP models is their complexity, which may pose an entry barrier for researchers and practitioners. To overcome this issue, the establishment of a few standard models may be helpful, as in, e.g., static, deterministic vehicle routing.
- Solution approaches: As detailed in Article A1, there has been considerable progress in the development of solution approaches. As a result, researchers and practitioners can now draw on an extensive toolbox of potential solution algorithms. In contrast to modeling, however, the design of solution algorithms is far less standardized. Hence, there is a need for bringing guidance and structure into algorithm selection and development for specific i-DMVRPs. Article A3 presents such contributions in the form of an application-independent explainability technique. It can directly support the development process and provides a characterization of fundamental opportunity cost approximation errors. A valuable contribution that could be provided by future research is a large-scale computational study comparing the performance of the most popular solution algorithms.

Next to standardization, the targeted expansion of the toolbox of solution approaches is still important. Exploring entirely new approaches and revisiting myopic approaches appears particularly interesting. E.g., in Article A2, we propose hybrid reward approximations that take advantage of the decomposability of opportunity cost. Article A3 suggests that advanced myopic approaches, which are commonly used only as benchmarks, deserve a more thorough investigation.

Finally, in view of the rather limited performance gains that can be achieved at the operational level, decision support at the tactical level and the strategic level should receive more attention (Waßmuth et al., 2023). However, this does not mean a loss of relevance of the operational planning level. Since the operational consequences of potential decisions at higher planning levels must be evaluated, diligently simplified methodology is required that is reasonably accurate and resembles actual operational decision-making as closely as possible. We present an example in Article A5 and demonstrate that our methodology is suitable for

providing decision support not only for demand management but also for other strategic decisions such as shift planning.

#### Second part of the dissertation

In the second part of the dissertation at hand, we investigate the application of demand management for improving the sustainability of rural SMOD services. The main findings can be grouped into four topics:

 Benefits of demand management: The overall impact of applying demand management in rural SMOD is positive. In the computational results of both Article A5 and Article A6, a consistent improvement in solution quality can be observed compared to first-come-firstserved decision-making. Thereby, offering multiple fulfillment options is crucial: If a request is only unfavorable due to the desired time, the customer can be steered toward choosing a more favorable alternative time. The empirical results in Article A4 indicate that this is realistic since customers have substantial time flexibility. As a result, only entirely unfavorable requests must be rejected or priced out.

Anticipatory opportunity cost approximation (ride evaluation) also proves to be beneficial based on the results in articles A5 and A6. Although sustainability improvements are limited, it avoids overestimation errors for early requests. This mitigates disparities in the offer quality, which would make round trips more difficult and incentivize strategic reservation behavior (Article A4).

The central limitation of the analyses in Articles A4-A6 is the lack of data revealing customers' true time flexibility and price sensitivity. While we can at least determine a lower bound for the former, the latter remains unknown. We address this issue by drawing on supplementary data from public transport in general and by conducting sensitivity analyses. However, collecting and analyzing data specifically from rural SMOD services is an important topic for future research. Another limitation is that our approaches do not consider cancellations. In Article A4, we show that they can have a major impact on operational planning. Since we also find that cancellation probabilities can be estimated reliably from historical data, incorporating them into demand management approaches appears very promising for future research.

2. **Managing the trade-off between sustainability objectives**: Due to the multi-dimensional nature of the all-inclusive objective of maximizing sustainability, sustainable demand management must balance multiple conflicting objectives. In Article A5, we find that single-objective demand management based on profit maximization can be a viable option that is easy for the provider to implement. However, the prioritization of economic sustainability can lead to undesirable outcomes in terms of social sustainability and environmental sustainability. Further, Article A6 shows that this approach is clearly not suitable in connection with dynamic pricing since it can lead to aggressive demand management with severe monopoly markups. To resolve this issue, we propose a multi-objective approach with the primary

objective being the maximization of served demand under restricted prices according to secondary objectives reflecting environmental sustainability and economic sustainability. With our work, we show that explicitly considering all relevant objectives in a multi-objective approach is similarly easy to implement for the provider as a single-objective approach.

- 3. **Comparison of availability control and dynamic pricing**: Based on the results from Article A5 and Article A6, we can conclude that dynamic pricing has crucial advantages over availability control. By design, demand is steered persuasively and not coercively. This means that customers ultimately decide whether to order a ride based on a price signal instead of the provider deciding to withhold unfavorable rides. Thereby, unfavorable rides are not avoided sweepingly but only if the customer's willingness-to-pay is insufficient. We show that, from the customers' perspective, our sustainable dynamic pricing approach would likely be perceived as moderate. It can be seamlessly built on top of a pre-existing static pricing scheme and causes only a slight increase in the general price level.
- 4. **Absolute performance**: It is important to note that the computational studies conducted in articles A5 and A6 are designed to yield insights about the relative performance impact of different demand management approaches within the SMOD system in isolation. Hence, the absolute values of performance attributes must be interpreted with caution as we discuss in the following:

Regarding social sustainability, the absolute values indicating served demand strongly depend on the fleet size, which determines the supply-demand ratio. Hence, the strategic fleet sizing decision and the tactical shift planning decision have a strong influence on social sustainability. Providing decision support for these decisions is rudimentary discussed in Article A5 and certainly deserves more attention in future research.

The absolute values of attributes measuring environmental sustainability consistently indicate that the vehicle km driven by the SMOD fleet exceed the (booked) passenger km. This means that a hypothetical scenario, in which all customers are assumed to use a private car, yields less vehicle km in total. However, this observation does not allow any clear-cut conclusions since it ignores important effects within the SMOD system itself and the entire transport system. Among them are several effects with a positive sustainability impact:

- The SMOD system often (partially) replaces inefficient scheduled services such that their vehicle km are saved (Mortazavi et al., 2024, Sieber et al., 2020, Viergutz and Schmidt, 2019).
- The higher service level of public transport due to the SMOD service can lead to households reducing the number of private cars or switching to public transport entirely, which could cause knock-on emission savings (e.g., Wang et al., 2021 or Jochem et al., 2020).
- People who cannot travel by private car are often driven by friends or relatives, which generates only half a passenger km per vehicle km (Thao et al., 2023).

• Finally, the Mohring Effect (Mohring, 1972) is not only valid in scheduled public transport systems but also in SMOD systems. Hence, serving more demand by scaling up SMOD services improves ridepooling and thereby environmental sustainability (Kaddoura und Schlenther, 2021, Lotze et al., 2023).

Contrary to these positive effects, there are also negative ones, first and foremost induced demand and cannibalization of active modes and (remaining) scheduled public transport (Rich, 2024). Overall, we can conclude that more research is necessary to holistically assess the environmental sustainability impact of rural SMOD and its interrelation with other modes of future rural public transport.

Finally, the absolute results regarding economic sustainability indicate a very large gap to a positive operating result, meaning that subsidies will be necessary in the foreseeable future. However, as for environmental sustainability, it is important to consider that some of these subsidies can be shifted from previously existing scheduled transport that the SMOD system replaces.

### References

- Jochem, P., Frankenhauser, D., Ewald, L., Ensslen, A., & Fromm, H. (2020). Does free-floating carsharing reduce private vehicle ownership? The case of SHARE NOW in European cities. *Transportation Research Part A: Policy and Practice*, 141, 373-395. https://doi.org/10.1016/j.tra.2020.09.016.
- Kaddoura, I., & Schlenther, T. (2021). The impact of trip density on the fleet size and pooling rate of ride-hailing services: A simulation study. *Procedia Computer Science*, 184, 674-679. https://doi.org/10.1016/j.procs.2021.03.084.
- Lang, M. A. K., Cleophas, C., & Ehmke, J. F. (2021). Multi-criteria decision making in dynamic slotting for attended home deliveries. *Omega*, 102, 102305. https://doi.org/10.1016/j.omega.2020.102305.
- Lotze, C., Marszal, P., Jung, F., Manik, D., Timme, M., & Schröder, M. (2023). Identifying the threshold to sustainable ridepooling. Working paper, Technical University Dresden. https://doi.org/10.48550/arXiv.2306.05851.
- Mohring, H. (1972). Optimization and scale economies in urban bus transportation. *The American Economic Review*, 62(4), 591-604.
- Mortazavi, A., Ghasri, M., & Ray, T. (2024). Integrated demand responsive transport in lowdemand areas: A case study of Canberra, Australia. *Transportation Research Part D: Transport and Environment*, 127, 104036. https://doi.org/10.1016/j.trd.2023.104036.
- Rich, J. (2024). Let's walk! The fallacy of urban first-and last-mile public transport. *Transportation, online first.* https://doi.org/10.1007/s11116-024-10505-5.
- Sieber, L., Ruch, C., Hörl, S., Axhausen, K. W., & Frazzoli, E. (2020). Improved public transportation in rural areas with self-driving cars: A study on the operation of Swiss train lines. *Transportation Research Part A: Policy and Practice*, 134, 35-51. https://doi.org/10.1016/j.tra.2020.01.020.
- Thao, V. T., Imhof, S., & von Arx, W. (2023). Demand responsive transport: New insights from peri-urban experiences. *Travel Behaviour and Society*, *31*, 141-150. https://doi.org/10.1016/j.tbs.2022.11.012.
- Viergutz, K., & Schmidt, C. (2019). Demand responsive-vs. conventional public transportation: A MATSim study about the rural town of Colditz, Germany. *Procedia Computer Science*, 151, 69-76. https://doi.org/10.1016/j.procs.2019.04.013.
- Wang, Y., Shi, W., & Chen, Z. (2021). Impact of ride-hailing usage on vehicle ownership in the United States. *Transportation Research Part D: Transport and Environment*, 101, 103085. https://doi.org/10.1016/j.trd.2021.103085.
- Waßmuth, K., Köhler, C., Agatz, N., & Fleischmann, M. (2023). Demand management for attended home delivery – A literature review. *European Journal of Operational Research*, 311(3), 801-815. https://doi.org/10.1016/j.ejor.2023.01.056.