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

TRANSPORTATION SCIENCE
Articles in Advance, pp. 1–28
ISSN 0041-1655 (print), ISSN 1526-5447 (online)

The Competitive Pickup and Delivery Orienteering Problem for Balancing Car-Sharing Systems

Layla Martin,[a] Stefan Minner,[b,c] Diogo Poças,[d] Andreas S. Schulz[e]

[a] Operations, Planning, Accounting and Control, Eindhoven University of Technology, 5612 AZ Eindhoven, Netherlands; [b] Logistics and Supply Chain Management, Technical University of Munich, 80333 Munich, Germany; [c] Munich Data Science Institute (MDSI), Technical University of Munich, 85748 Garching, Germany; [d] LASIGE, Faculdade de Ciências, Universidade de Lisboa, 1649-004 Lisbon, Portugal; [e] Operations Research, Technical University of Munich, 80333 Munich, Germany

Contact: l.martin@tue.nl, https://orcid.org/0000-0003-4580-9498 (LM); stefan.minner@tum.de, https://orcid.org/0000-0001-6127-8223 (SM); dmpocas@fc.ul.pt, https://orcid.org/0000-0002-5474-3614 (DP); andreas.s.schulz@tum.de, https://orcid.org/0000-0002-9595-459X (ASS)

Received: November 16, 2019
Revised: May 25, 2020; October 15, 2020
Accepted: January 13, 2021
Published Online in Articles in Advance: October 5, 2021

https://doi.org/10.1287/trsc.2021.1041

Abstract. Competition between one-way car-sharing operators is currently increasing. Fleet relocation as a means to compensate demand imbalances constitutes a major cost factor in a business with low profit margins. Existing decision support models have so far ignored the aspect of a competitor when the fleet is rebalanced for better availability. We present mixed-integer linear programming formulations for a pickup and delivery orienteering problem under different business models with multiple (competing) operators. Structural solution properties, including existence of equilibria and bounds on losses as a result of competition, of the competitive pickup and delivery problem under the restrictions of unit-demand stations, homogeneous payoffs, and indifferent customers based on results for congestion games are derived. Two algorithms to find a Nash equilibrium for real-life instances are proposed. One can find equilibria in the most general case; the other can only be applied if the game can be represented as a congestion game, that is, under the restrictions of homogeneous payoffs, unit-demand stations, and indifferent customers. In a numerical study, we compare different business models for car-sharing operations, including a merger between operators and outsourcing relocation operations to a common service provider (coopetition). Gross profit improvements achieved by explicitly incorporating competitor decisions are substantial, and the presence of competition decreases gross profits for all operators (compared with a merger). Using a Munich, Germany, case study, we quantify the gross profit gains resulting from considering competition as approximately 35% (over assuming absence of competition) and 12% (over assuming that the competitor is omnipresence) and the losses because of the presence of competition to be approximately 10%.

Funding: The work of Layla Martin was supported by Deutsche Forschungsgemeinschaft as part of the Research Training Group 2201 (Advanced Optimization in a Networked Economy) and the Diversity Promotion Fund of the Technical University of Munich School of Management. The work of Diogo Poças and Andreas S. Schulz was supported by the Alexander von Humboldt Foundation with funds from the German Federal Ministry of Education and Research.
Supplemental Material: The e-companion is available at https://doi.org/10.1287/trsc.2021.1041.

Keywords: car-sharing • competition • mixed-integer linear programming

1. Introduction

Car-sharing is an economically and environmentally sustainable alternative to private car ownership and a supplement to public transportation (Bellos, Ferguson, and Toktay 2017). With a larger number of car-sharing offers around the world, competition increases. In several cities, more than one car-sharing operator offers its service to customers (Kortum et al. 2016), and mobility-as-a-service apps allow customers to book any vehicle regardless of the operator. Some operators are starting to merge their companies and fleets, most recently Car2Go and DriveNow (now called "ShareNow") (BMW Group 2018). Soon after, the former DriveNow shareholder Sixt launched a car-sharing service, "Sixt Share," that competes with ShareNow (Sixt 2019). Perboli et al. (2018) report frequent changes in the Turin market with BlueTurino entering the market and competing with Car2Go and Enjoy as IoGuido withdrew service. Although the body of literature on the optimization of car-sharing operations is growing, it mostly has ignored the choice of business and operational models under competition so far (with exceptions Albiński and Minner 2019, Balac et al. 2019). In addition to the merger (and, thus, a monopoly) as currently pursued by DriveNow and Car2Go and direct competition, such as between Sixt Share and ShareNow, operators can cooperate in parts of their operations. For example, they can hire a third party that relocates vehicles for them, or both operators relocate vehicles such that the overall gross profit is maximized (but still use different workers for rebalancing). We call these modes "coopetition" and "welfare maximization." Ghosh et al. (2017) report outsourced relocation operations in bike-sharing. Brook (2004) reports a collaboration between different car-sharing operators concerning other aspects, such as billing. In practice, operators frequently ignore competition with respect to relocation. Mostly, they do not even consider the current location of vehicles of the competitor, which can be accessed using web-scraping techniques or accessing publicly available application programming interfaces (Kortum et al. 2016, Trentini and Losacco 2017), let alone foresee how the competitor rebalances those vehicles (which could be derived from historic data). Incorporating the reaction of a competitor on one's own routing and servicing decision is paramount for profitable operations, in particular, if fleets are small and margins are low. Operators may service the same customer, thereby reducing the individual payoff but still incurring relocation costs.

A particular focus of this work is rebalancing, which shows significant potential for cost reduction in car-sharing systems (Jorge and Correia 2013). One-way car-sharing comes at the cost of unevenly distributed fleets as some locations (such as shopping malls) are more frequently the point of origin of a car-sharing trip than of its destination (or vice versa). To cause as little customer dissatisfaction as possible and, thus, as few lost sales as possible, vehicles are relocated (predominantly by pooling multiple rebalancing operations during the night; Weikl and Bogenberger 2015). Vasconcelos et al. (2017) report a positive impact of relocation on the profitability of a car-sharing service. They state that a car-sharing service is only profitable if relocations are performed. Consequently, nonoptimal car-sharing relocations (or relocations that do not consider competition) can easily result in negative profits.

We answer three research questions and contribute algorithmically to solution methods for carsharing relocation under competition.

1. How much can operators gain from considering the presence of competition in their rebalancing operations with regards to gross profits? Put differently, what is the price of ignoring the presence of competition?

2. How much is lost by competing in comparison with jointly optimizing fleet rebalancing with regard to gross profits, and how do alternative business models under competition compare with each other?

3. Which parameters drive the gains from considering competition and the losses resulting from the presence of competition?

We briefly address the research questions analytically under some technical assumptions (giving bounds on gains and losses) and study the full generality of the model in an extensive numerical study.

Methodologically, we present models and solution algorithms to tackle the aforementioned research questions. We give complexity as well as average case algorithmic performance results. Advantages and drawbacks that the solution algorithms and model formulations entail are studied. We devise a simplified model for car-sharing relocation that incorporates the movements of vehicles as well as the movements of the workers who relocate them. This model is called the pickup and delivery orienteering problem (PDOP). A feasible solution to the PDOP takes two decisions simultaneously: First, it determines which vehicles are moved and which stations are serviced with how many vehicles (customer demand is given in a previous step, e.g., by using a demand-prediction model as reviewed in Vosooghi et al. 2017). Second, it specifies which route relocation workers take to perform those relocations. We then create variations of this model to facilitate the special features of different business models under competition: operators can either compete directly (C-PDOP), jointly optimize their rebalancing (W-PDOP), merge their fleets completely (M-PDOP), or outsource their relocation operations to a third party (Coop-PDOP).

We choose pure strategy Nash equilibria as the solution concept (as opposed to mixed strategy Nash equilibria in which players randomize over a set of strategies), which is a single routing and servicing choice rather than a combination of multiple choices. Mixed strategy Nash equilibria provably exist in the C-PDOP because the number of pure strategies is finite (Nash 1951). Pure strategy Nash equilibria represent a more intuitive, easier-to-implement solution for the car-sharing operators. They represent a less controversial concept of competition and are more likely to be adopted by car-sharing operators. Pure strategy Nash equilibria provably exist subject to three assumptions: player homogeneous payoffs, unit demand stations, and indifferent customer choice. For the C-PDOP with these assumptions, we present structural properties and performance guarantees (algorithmic and with respect to profitability) and show average case performances in a numerical study. The numerical study also considers the influence of relaxing the assumptions.

To solve the C-PDOP, we present two different algorithms: iterated best response (IBR) and potential function optimizer (PFO). Although IBR is faster on most instances and can be used even for the general model in which pure strategy Nash equilibria are not guaranteed to exist, PFO returns higher profit Nash equilibria, on average, and has a higher degree of fairness with respect to how profits are distributed between operators. Both algorithms draw upon the fact that the C-PDOP with player homogeneous payoffs, unit demand stations, and indifferent customer choice belongs to the class of congestion or potential games introduced in Rosenthal (1973) and Monderer and Shapley (1996).

The remainder of this paper is structured as follows. In Section 2, we discuss the literature and how it relates to the (competitive) pickup and delivery orienteering problem. Section 3 presents the model for the PDOP and extensions for the various business models. It also investigates properties of these models. In Section 4, we present the IBR and the PFO to find a pure strategy Nash equilibrium of the C-PDOP if it exists. In Section 5, we test the aforementioned algorithms on several different, artificially generated instances and a Munich, Germany, case study before concluding the paper in Section 6. All proofs can be found in the appendix.

2. Literature Review

The (competitive) pickup and delivery orienteering problem is related to three different streams of literature. In recent years, there has been a growing body of literature in the field of optimal car-sharing operations; we focus on papers that present different rebalancing approaches. Second, a few papers already discuss routing or rebalancing problems under the presence of competition. Third, our work draws upon findings in the field of routing, namely pickup and delivery problems and traveling salesman problems (TSPs) or vehicle routing problems (VRPs) with profits.

2.1. Car-Sharing Operations

Research in car-sharing operations predominantly focuses on the tactical problem of setting the proper inventory levels or vehicle stock in a region (Nair and Miller-Hooks 2011; He, Hu, and Zhang 2020) or on incentivizing users toward rebalancing the system themselves by adapting their travel behavior (e.g., Ströhle, Flath, and Gärttner 2019). Some research also focuses on the underlying VRP (see Laporte, Meunier, and Calvo 2015, 2018 for overviews). Krumke et al. (2013) study the k-convoy pickup and delivery problem in which vehicles are relocated by k drivers. They study the problem in both static and dynamic settings, presenting approximation algorithms for both cases. Bruglieri, Colorni, and Luè (2014) study the electric vehicle relocation problem as a variant of the one-skip relocation problem and the roll on–roll off problem. For their Milan case study, they report that two workers are sufficient to fulfill, on average, 86% of 30 relocation requests. When maximizing the profit rather than a service level, they observe that a "fixed revenue component" (future revenue of satisfied customers) per served customer must be at least 15€ to ensure that enough rebalancing occurs to serve most of the customers (Bruglieri, Pezzella, and Pisacane 2017). Recent work addresses challenges in repositioning autonomous vehicles, such as having to consider matching demand and supply rather than finding an optimal route. Unlike worker-based relocation, autonomous relocations are mainly performed during the day and, thus, require online algorithms (i.e., Cepolina and Farina 2014; Fagnant, Kockelman, and Bansal 2015; Pavone 2015).

2.2. Competitive Routing and Rebalancing Models

Only few studies take a competitive view on rebalancing car-sharing operations, and those take a tactical rather than an operational point of view or consider pricing instead of routing. Albiński and Minner (2019) calculate the number of vehicles that are to be relocated to reach an equilibrium under demand uncertainty (inventory transshipment problem). They do not focus on the actual routes of workers who relocate the vehicles (operational routing problem). Balac et al. (2019) use an agent-based model to derive optimal prices for two car-sharing operators and other options, such as walking or public transport. They provide insights into whether it is advisable to additionally rebalance vehicles during the course of the day. They observe that charging the same (comparably high) price is most profitable for both operators. However, high prices are unstable as operators benefit from offering lower prices than their competitors. They state that relocations during the course of the day are unprofitable and further observe that relocation in the presence of competition primarily benefits the competitor who does not rebalance. The origin of this "free-rider" phenomenon, however, is demand during the relocation operations and the selection of a simple policy-based heuristic for relocation. Both simplifications do not apply in our model. The two aforementioned papers are restricted to the two-operator case. For the ridesharing sector with a variable number of operators, Pandey et al. (2019) argue that competition decreases the efficiency as well as the quality of service and show that even little cooperation can substantially increase the service quality. Their main focus, however, is not on rebalancing during the night, but on assigning customers to routes and the subsequent online rerouting of vehicles.

2.3. Routing Problems for Rebalancing

In the PDOP, the goal is to find optimal routes for operators who relocate their vehicles; thus, the PDOP is related to other routing problems, in particular pickup and delivery problems (see Berbeglia et al. 2007 for a comprehensive overview and classification) and TSPs with profits (see Feillet, Dejax, and Gendreau 2005; Vansteenwegen, Souffriau, and Van Oudheusden 2011; Archetti, Speranza, and Vigo 2014 for overviews). Following the classification by Berbeglia et al. (2007), we identify the PDOP as a one-commodity VRP. So far, there is neither a combination of the profitable tour problem with pickups and deliveries nor a bipartite profitable tour problem to be found in the literature. In addition to a formulation of TSPs with pickup and deliveries and with profits, our paper relates to the literature on lateral inventory transshipment (see Paterson et al. 2011 for an overview) focusing on operational decisions in a unidirectional inventory transshipment problem in a single echelon. Lien et al. (2011) compare different setups of the inventory transshipment graph with a particular focus on "chains," that is, circular setups in which facilities can only exchange inventory with their direct neighbors. They state that bidirectional chains are more efficient than unidirectional chains, which again are more efficient than other setups with fewer links within the same echelon. We observe similar efficiency gains when comparing competitive solutions in which every operator rebalances for the operator's self (thus, resulting in substantially fewer arcs) to the monopoly or coopetition solution.

3. Model

In the following, we first present a master model for car-sharing relocation. We then instantiate this model to a competitive variant (C-PDOP), a monopoly variant (M-PDOP), and a coopetition variant (Coop-PDOP) and discuss properties of these problems.

3.1. Pickup and Delivery Orienteering Problem

In our car-sharing model, each of N operators knows the operator's current fleet distribution and wants to rebalance it, assuming that no customers are requesting vehicles during the process. We use index $n = 1, 2, \ldots, N$ to refer to each of the operators. In car-sharing systems, such relocations are usually performed periodically during each night and when customer demand is low (Weikl and Bogenberger 2015). To rebalance the operator's fleet, operator n sends workers who drive the vehicles from their current location to another location where they are expected to incur a higher payoff. The overall goal of an operator is to maximize the operator's gross profit (profit of a given fleet), defined as payoffs minus operational costs. The payoffs depend on the presence of competition. Although the strategy of the competitors need not be known, all operators must be aware of the payoff functions of all competitors and assume that their competitors act rationally and payoff maximizing.

3.1.1. Stations, Locations, and Payoffs.

Each operator has a set of stations, D_n, at which the operator can place vehicles. The sets of stations of two different operators are not necessarily disjoint as operators may place stations very close to each other or even at the same location (e.g., at public transit stops or the trade fair). Each station is then referred to as $\iota \in D, D = \bigcup_n D_n$. We assume that the operators employ forecasting models and, thus, estimate the expected customer demand in each station, allowing them to predict if vehicles shall be moved to or removed from this station (but not both simultaneously). Each operator n chooses, for each station ι, the number q_n^ι of vehicles to place there. For every station ι, operator n has an estimate of the payoff function $\pi_n^\iota(\mathbf{q}^\iota)$, where $\mathbf{q}^\iota = (q_1^\iota, q_2^\iota, \ldots, q_N^\iota)$ is the vector that describes the number of vehicles placed by each operator n at station ι. The quantity $\pi_n^\iota(\mathbf{q}^\iota)$ captures the direct (expected) revenue from setting a vehicle at this location given the availability of competitors $\mathbf{q}^\iota$, minus the direct (expected) costs, such as fuel and wear, but it does not include rebalancing costs, which are modeled using a different term.

We use the standard game-theoretic notation of $\mathbf{q}_{-n}^\iota$ to denote the vector of the decisions of the players other than n so that we can write $\mathbf{q}^\iota = (q_n^\iota, \mathbf{q}_{-n}^\iota)$. The model allows for a very generic formulation and choice of these payoff functions, which adhere to the following:

- Keeping the values $\mathbf{q}_{-n}^\iota$ of the operators other than n fixed, the function $\pi_n^\iota(q_n^\iota, \mathbf{q}_{-n}^\iota)$ is *concave*, *nondecreasing* in q_n; the monotonicity suggests that we can only increase revenue by placing more vehicles at the same station, whereas concavity suggests that the marginal payoff (by placing one extra vehicle) decreases with the number of vehicles placed in the station.
- For $m \neq n$, keeping fixed the values $\mathbf{q}_{-m}^\iota$ of the operators other than m, the function $\pi_\iota^n(q_m^\iota, \mathbf{q}_{-m}^\iota)$ is *nonincreasing* in q_m; this suggests that the revenue of an operator can only decrease if competing operators place more vehicles in that station.

We deliberately do not restrict $\pi_n^\iota(\mathbf{q}^\iota)$ further. This allows us to formulate different types of customer choice. For example, consider a station ι at which a single customer is expected who strictly prefers operator 1 (but would take operator 2 if operator 1 is unavailable). This could be modeled by $\pi_1^\iota(1,1) = \pi_1^\iota(1,0) = \pi_2^\iota(0,1) > 0$ and $\pi_2^\iota(1,1) = 0$. For another example, consider a station ι at which a single customer is expected who is indifferent between operators 1 and 2. This could be modeled by $\pi_1^\iota(1,1) = \pi_2^\iota(1,1) = 0.5\pi_1^\iota(1,0) = 0.5\pi_2^\iota(0,1)$.

We further assume that there is a maximum number of vehicles that can be moved to a station profitably; that is, q_n^ι is upper-bounded by some constant $\hat{q}_n^\iota$ that represents the maximum demand at station ι. In practice, operators usually employ some "filtering" upfront (Weikl and Bogenberger 2015), resulting in a reasonably small number of vehicles that can be profitably placed in a station.

In this section, we optimize for an operator n, considering that the other operators are "frozen" in their strategies. Thus, the payoff values $\pi_n^\iota(q_n^\iota, \mathbf{q}_{-n}^\iota)$ vary in q_n^ι only. We formulate the problem as an integer program, and as such, "break" the nonlinearity of the payoff functions π_n^ι into a linear objective function. The simplest way to do so is to split a station ι into a set $\mathcal{Z}^\iota$ of *locations* labeled $i = 1, \ldots, \hat{q}_n^\iota$. Each station ι can be split into at most $\hat{q}_n^\iota$ locations. We separate the station-based payoff function into a sum of location-based payoff values, assigning to the i th location the marginal gain of placing the i th vehicle. These marginal gains can be computed by consecutive differences; that is, $\pi_n^i = \pi_n^\iota(i, \mathbf{q}_{-n}^\iota) - \pi_n^\iota(i-1, \mathbf{q}_{-n}^\iota)$. From our assumption of concavity, the marginal gains are nonincreasing, which results in an intrinsic ordering of locations; that is, without loss of generality, an operator would choose to visit the i th location associated to a station ι only after visiting all lower-indexed locations $1, \ldots, i-1$.

Locations are represented as nodes in a rebalancing graph. We denote the full set of nodes available to operator n as $\mathcal{Z}_n$. The operator already has vehicles present at some locations but not others; we refer to locations that do not have a vehicle as delivery locations ($\mathcal{Z}_n^-$) and to those with a vehicle present as pickup locations ($\mathcal{Z}_n^+$). Thus, we can write $\mathcal{Z}_n = \mathcal{Z}_n^+ \cup \mathcal{Z}_n^-$. Finally, we also include a depot node, which is modeled as both a pickup and delivery location.

In an analogy to Bruglieri, Pezzella, and Pisacane (2017), we formulate the problem as a variant of a TSP with profits in which not all locations need to be visited (similar to prize-collecting TSPs). This allows the integration of two decisions: the decision on which locations to service and the actual routing decision. Additionally, one can specify subsets of locations $S_n^- \subseteq \mathcal{Z}_n^-, S_n^+ \subseteq \mathcal{Z}_n^+$ that must be visited. In this way, we cater to mandatory customers (e.g., to fulfill a contract) and also enforce that vehicles that are parked illegally are picked up.

Importantly, we assume that nearby stations do not influence each other. This is arguably a simplifying assumption as users might walk to the next available vehicle in a nearby station even if this increases the walking distance and because vehicle returns depend on the availability of vehicles at the origin. This assumption is common in the literature on car-sharing rebalancing (Bruglieri, Pezzella, and Pisacane (2018), for example, randomly assign customers who belong to the catchment area of multiple stations), considering that interdependent demand processes would significantly increase the problem complexity and is beyond the scope of this paper.

3.1.2. Arcs and Costs. For each pair of locations $i \in \mathcal{Z}_n^-$, $j \in \mathcal{Z}_n^+$, we associate costs $c_{\langle ij \rangle}, c_{\langle ji \rangle}$ and travel times $\tau_{\langle ij \rangle}, \tau_{\langle ji \rangle}$ between these locations. The costs serve as an estimate for the costs of the amount of fuel as well as the payment for the time it takes to move between the locations (either proportionate wages or payments to a service provider). Note that the costs and travel times are not necessarily symmetric and are always equal for pairs of locations from the same station (e.g., $c_{ij} = c_{kl}$ if i and k are locations for the same station ι_1 and j and l are locations for the same station ι_2). Operators drive from pickup to delivery locations, but the travel from a delivery to a pickup location may be done in a lot of different ways, such as using a foldable bike, walking, public transit, or using a second car driven by a colleague ("double driving"). We can then describe the problem in terms of its underlying network or directed bipartite graph G, whose nodes correspond to the locations $\mathcal{Z}_n$, and with arcs A_n between each pickup location and each delivery location in either direction.

3.1.3. Workers and Tours. Operator n employs W_n workers who relocate vehicles by starting at a depot, visiting pickup and delivery locations in alternating order, and returning to the depot at the end of their shift of at most T units of time. A worker can only relocate one vehicle at a time. As already mentioned, we assume that there is a central depot at which workers can collect and return equipment, cleaning supplies, and lost items. Although assuming this rebalancing mode poses a simplified model, we later show that the game-theoretical formulation can be extended to more complex cost models, such as rebalancing using a truck (that can relocate multiple vehicles) or assistance using a minibus (that can transport multiple workers).

Our problem is, thus, specified by the bipartite graph G (i.e., the set of locations $\mathcal{Z}_n = \mathcal{Z}_n^+ \cup \mathcal{Z}_n^-$ including two copies of the depot $d_n \in \mathcal{Z}_n^-$, $p_n \in \mathcal{Z}_n^+$ and the set of arcs A_n); for each location i (associated to a station ι), an expected revenue π_n^i (corresponding to the marginal gains of placing an i th vehicle in the corresponding station); for each arc e, a travel cost c_e and travel time τ_e; sets of enforced visits $S_n^+ \subseteq \mathcal{Z}_n^+$, $S_n^- \subseteq \mathcal{Z}_n^-$; a number of workers W_n; and a maximum shift time T. We set up an integer programming formulation to search for a profit-maximizing tour. The decision variables are

• For each location i, a_i^n denotes whether operator n has a vehicle at location i after rebalancing; in other words, $a_i^n = 1$ iff either i is a pickup location at which n chooses not to remove the vehicle or i is a delivery location that n chooses to serve with a vehicle.

• For each arc e, x_e^n denotes whether operator n chooses arc e on one of the routes.

- For each location i, t_i denotes the point in time at which location i is visited.

$$\max \Pi_n = \sum_{i\in\mathcal{Z}_n\setminus\{d_n,p_n\}} \pi_i^n a_i^n - \sum_{e\in A_n} c_e x_e^n, \tag{1a}$$

$$\text{s.t. } a_i^n = \sum_{j\in\mathcal{Z}_n^+} x_{\langle ij\rangle}^n = \sum_{j\in\mathcal{Z}_n^+} x_{\langle ji\rangle}^n \quad \forall\, i\in\mathcal{Z}_n^-\setminus\{d_n\}, \tag{1b}$$

$$1-a_i^n = \sum_{j\in\mathcal{Z}_n^-} x_{\langle ij\rangle}^n = \sum_{j\in\mathcal{Z}_n^-} x_{\langle ji\rangle}^n \quad \forall\, i\in\mathcal{Z}_n^+\setminus\{p_n\}, \tag{1c}$$

$$\sum_{i\in\mathcal{Z}_n^+\setminus\{p_n\}} x_{\langle d_n i\rangle}^n \le W_n, \tag{1d}$$

$$\sum_{i\in\mathcal{Z}_n^-\setminus\{d_n\}} x_{\langle i p_n\rangle}^n \le W_n, \tag{1e}$$

$$t_i + x_{ij}(\tau_{ij}+T) \le t_j + T \quad \forall\, \langle I,j\rangle\in A_n, \tag{1f}$$

$$\sum_{e\in A(C)} x_e^n \le |C|-1 \quad \forall\, C\subseteq \mathcal{Z}_n\setminus\{d_n,p_n\}, \tag{1g}$$

$$a_i^n = 1 \quad \forall\, i\in S_n^-, \tag{1h}$$

$$a_i^n = 0 \quad \forall\, i\in S_n^+, \tag{1i}$$

$$x_e^n, a_i^n \in \{0,1\} \quad \forall\, e\in A_n, i\in\mathcal{Z}_n, \tag{1j}$$

$$0\le t_i \le T \quad \forall\, i\in\mathcal{Z}_n. \tag{1k}$$

This model is an alternative notation of a multi-vehicle profitable tour problem on a bipartite graph and differs from the two-index notation of the VRP in three key components: first, the operator maximizes gross profit; second, visiting a location is not mandatory—instead, there are only incoming and outgoing arcs if a location is visited; and third, the graph G is bipartite, which restricts the routing options and necessitates a split depot. The model can also be interpreted as a pickup and delivery problem or as a dial-a-ride problem with unit capacity (see Parragh, Doerner, and Hartl 2008 for reviews on these problem classes). Equation (1a) maximizes the gross profit given by payoffs minus costs (assuming fleet procurement costs are sunk). Equations (1b) and (1c) are assignment constraints that link availability and visits of locations, ensuring that each location has the same number of incoming and outgoing arcs (flow conservation) and at most one of each. Equations (1d) and (1e) guarantee that no more than W_n workers leave the depot and return there (either one would be sufficient and directly imply the other). Constraints (1f) ensure that all workers return to the depot within the shift length T. Constraints (1g) are used for subtour elimination; although formally they are redundant because of Constraints (1f), including these constraints helps improve the runtime in practice. Equations (1h) and (1i) specify that all locations in the sets S_n^- and S_n^+ must be visited; that is, a vehicle must either be left there or removed from there.

3.2. Competitive Pickup and Delivery Orienteering Problem

In the PDOP model, each operator implicitly assumes that the number of vehicles at a station is known. However, competitors also react to how many vehicles this operator deploys at some station. With the rise of mobility-as-a-service solutions, the number of customers who are registered with multiple operators grows, which increases the relevance of considering competition in the optimization models.

We define competitive stations as those that can be accessed by multiple players. For example, if shared customers are expected at station ι and if, before relocation, two or more operators have a vehicle there, it can be beneficial for one of them to pick up the car and service another station. In the C-PDOP, each operator relocates the operator's fleet with the goal of maximizing the gross profit while considering that the other operators are relocating their fleets and strategizing accordingly. Thus, we consider Nash equilibria as the desired solution concept; that is, we search for a strategy profile by which no operator can benefit from unilateral deviation. A strategy defines the number of vehicles operator n has at a station ι after rebalancing, that is, the vector $\mathbf{q}_n$. We refer to $\mathbf{q}_n$ as a compact strategy as it is sufficient to represent the entire solution. In other words, the routing decisions follow directly from the availability at the locations of player n by calculating an optimal solution to (1a)–(1k).

The model in Equations (1a)–(1k) is further extended by introducing the competitive profit functions Π_n for each player n. Let $\mathbf{q} = (\mathbf{q}_1, \ldots, \mathbf{q}_N)$ denote the joint profile of (compact) strategies; $\mathbf{q}_{-n}$ denotes the joint profile of all players except n. From $\mathbf{q}_{-n}$ and for a given player n, the competitive profit is defined as the difference between station-separable payoff and costs, $\Pi_n(\mathbf{q}) = R_n(\mathbf{q}) - C_n(\mathbf{q})$. The cost term can be written as $C_n(\mathbf{q}) = C_n(\mathbf{q}_n) = \sum_{e\in A_n} c_e x_e^n$, where x_e^n is an optimal (i.e., cost-minimizing) choice of routing decisions. Notice that the cost terms depend only on the strategy of operator n; henceforth, we write $C_n(\mathbf{q}_n)$ instead.

The competitive profit functions can be written as

$$\Pi_n(\mathbf{q}) = R_n(\mathbf{q}) - C_n(\mathbf{q}_n) = \sum_{\iota\in D_n} \pi_n^\iota(\mathbf{q}_{-n}^\iota) - C_n(\mathbf{q}_n). \tag{2}$$

Given the definition of the strategies and the associated gross profits, a strategy profile $\mathbf{q}$ forms a Nash equilibrium if and only if, for every player n,

$$\Pi_n(\mathbf{q}) = \max_{\mathbf{q}_n'} \Pi_n(\mathbf{q}_n', \mathbf{q}_{-n}); \tag{3}$$

that is, the Nash equilibrium strategies maximize the gross profit of the players given the other operators' strategies as input. Thus, no player can profit from deviating unilaterally.

We should state here that our model allows for different profit functions π_n^ι, and as a consequence, such games may, in general, not have pure strategy Nash equilibria (we provide an example in this section). From a theoretical perspective, we can study conditions under which pure strategy Nash equilibria are guaranteed to exist.

3.2.1. Unit Demand Stations. By this we mean that at most one customer is expected at station ι, and accordingly, each operator places at most one vehicle there. Under this assumption, our notation can be simplified because the notions of stations and locations become equivalent, and in particular, we have $q_n^\iota = a_n^i \in \{0, 1\}$, where i is the unique location associated with station ι. Also, under this assumption, we consider competitive locations rather than competitive stations; we let $\mathcal{Z} = \bigcup_n \mathcal{Z}_n$ denote the set of all locations and $\mathcal{Z}^C$ denote the set of competitive locations.

3.2.2. Indifferent Customer Choice. That is, customers do not have a preference among different operators and, thus, select vehicles at random and with equal probability. In particular, under this assumption, all customers are indifferent between all operators to which they are subscribed (although different customers may still have different subscriptions). This assumption is realistic if customers with multiple memberships choose the closest vehicle regardless of the operator. Note that, if we assume both unit demand stations as well as indifferent customer choice, the payoff function $\pi_n^i(\mathbf{q}^i)$ can be described rather succinctly: if $\bar{\pi}_n^i = \pi_n^i(1, 0_{-n})$ is the payoff that player n could extract by being the only operator at location i and $m = |\{n' \neq n : q_{n'}^i = 1\}|$ is the number of operators at location i (excluding n), then

$$\pi_n^i(\mathbf{s}^i) = \frac{\bar{\pi}_n^i}{m+1}.$$

In such a situation, we can efficiently specify the payoff functions with a single parameter $\bar{\pi}_n^i$ for each (location, operator) pair.

3.2.3. Player Homogeneous Payoffs. We assume that $\pi_n^\iota = \pi^\iota$ does not depend on n, for all stations ι, if n offers service at location ι. This assumption is justified by evidence that margins are driven by very similar revenues (see Balac et al. 2019). When combined with the first two assumptions, this further restricts the dimensionality of the problem: we only need to specify, for each location i, a value π^i for the payoff that any player could extract by being the sole operator at i.

Most of our theoretical results are proven under the combination of all three assumptions. From now on, we use the terminology *restricted C-PDOP model* whenever we refer to an instance satisfying all three assumptions (unit demand stations, indifferent customer choice, and player homogeneous payoffs). The restricted C-PDOP model has theoretical advantages: it allows for a congestion game formulation (see Lemma 1); thus, we can guarantee the existence of at least one pure strategy Nash equilibrium and that such equilibria can be reached via best-response dynamics (see Corollary 1). If these assumptions do not hold, cases without pure strategy Nash equilibria can exist, meaning that stability is not guaranteed. In Section 5, we indeed experiment with the generality of the model because the preceding assumptions can be very restrictive for some car-sharing systems.

3.3. Examples of Games Without Nash Equilibria

Dropping either of the three assumptions (even while keeping the other two), we can construct games without a pure strategy Nash equilibrium.

3.3.1. Multidemand Stations. In the example of Figure 1, we assume indifferent customer choice and player homogeneous payoffs but allow for multidemand stations. If a total of q vehicles are placed at a station ι, then a total revenue of $\pi^\iota(q)$ is extracted from this station. The revenue is split in proportion to the amount of vehicles by each player; thus, if players 1 and 2 place q_1 and q_2 vehicles, respectively (so that $q = q_1 + q_2$), then they receive payoffs of $\frac{q_1}{q}\pi^\iota(q)$ and $\frac{q_2}{q}\pi^\iota(q)$, respectively. Player 1 routes on the left half of the graph and has essentially two undominated strategies: put two vehicles in each of the stations A and B or put two vehicles in each of the stations C and D. Similarly, player 2 routes on the right half of the graph and has essentially two undominated strategies: put a vehicle in each of the stations A and C, or put a vehicle in each of the stations B and D. For simplicity, we assume that the routing costs of each of these strategies is normalized to zero, so we only need to worry about the way payoffs are split.

Finally, we choose the station payoffs as in Figure 1. For illustration purposes, let us compute the gross profits if player 1 takes the "AABB" tour and player 2 takes the "AC" tour. In this situation, player 1 extracts two thirds of the revenue from station A and full revenue from station B, and player 2 extracts one third of the revenue from station A and full revenue from station C. We, thus, have

$$\Pi_1 = \frac{2}{3}\pi^A(3) + \pi^B(2) = \frac{14}{3} + 7 = \frac{35}{3};$$

$$\Pi_2 = \frac{1}{3}\pi^B(3) + \pi^C(1) = \frac{7}{3} + 4 = \frac{19}{3}.$$

Note that player 2 would rather deviate to the "DB" tour and increase gross profit. By doing these calculations for all possible pairs of strategies, we see that no Nash equilibrium exists; in particular, any sequence of iterated best

Figure 1. Example of a C-PDOP Instance with No Pure Nash Equilibria

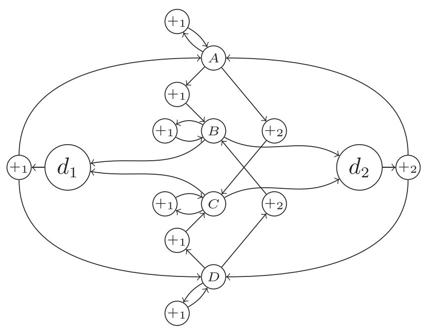

q	1	2	3
$\pi^A(q)$	4	6	7
$\pi^B(q)$	4	7	8
$\pi^C(q)$	4	7	8
$\pi^D(q)$	4	6	7

	0	AC	DB
0	$0, 0$	$0, \frac{24}{3}$	$0, \frac{24}{3}$
AABB	$\frac{39}{3}, 0$	$\frac{35}{3}, \frac{19}{3}$	$\frac{34}{3}, \frac{20}{3}$
DDCC	$\frac{39}{3}, 0$	$\frac{34}{3}, \frac{20}{3}$	$\frac{35}{3}, \frac{19}{3}$

Notes. There are four delivery districts, labeled A, B, C, D. Each player has a depot (d_1 or d_2) and vehicles at distinct pickup locations ($+_1$ and $+_2$, respectively) with zero payoff. Player 1 can put at most two vehicles in a district and has essentially three strategies (null, AABB tour, and DDCC tour). Player 2 can put at most two vehicles in a district and has essentially three strategies (null, AC tour, and DB tour). The concave payoffs are player homogeneous and follow indifferent customer choice. For a specific choice of location payoffs (center) and payoff matrix (right), no Nash equilibrium exists.

responses gets stuck in the loop (AABB, AC) → (AABB, DB) → (DDCC, DB) → (DDCC, AC) → (AABB, AC).

3.3.2. Differentiated Customer Choices. In the example of Figure 2, we assume unit demand stations and player homogeneous payoffs but allow for differentiated customer choice. Our example is built with two players. Each station ι has an associated unit revenue $\pi^\iota = 1$ that can be extracted. If both operators serve this station, the revenue is split unequally among the operators. In stations A and C, the revenue is split $75-25$ in favor of player 1, but in stations B and D the revenue is split $75-25$ in favor of player 2. Thus, if, for example, both players service station A, then they receive payoffs of 3/4 and 1/4, respectively. Player 1 routes on the left half of the graph and has essentially two undominated strategies: service stations A and B or service stations C and D. Similarly, player 2 routes on the right half of the graph and has essentially two undominated strategies: service stations A and C or service stations B and D. For simplicity, we assume that the routing costs of each of these strategies is normalized to zero, so we only need to worry about the way payoffs are split.

For illustration purposes, let us compute the gross profits if player 1 takes the "AB" tour and player 2 takes the "AC" tour. In this situation, player 1 extracts 75% of the revenue from station A and full revenue from station B, and player 2 extracts 25% of the revenue from station A and full revenue from station C. We, thus, have

$$\Pi_1 = \frac{3}{4} + 1 = \frac{7}{4};$$
$$\Pi_2 = \frac{1}{4} + 1 = \frac{5}{4}.$$

Note that player 2 would rather deviate to the "DB" tour and increase gross profit. By doing these calculations for all possible pairs of strategies, we see that no Nash equilibrium exists; in particular, any sequence of iterated best responses gets stuck in the loop (AB, AC) → (AB, DB) → (DC, DB) → (DC, AC) → (AB, AC).

3.3.3. Player Heterogenous Payoffs. In the example of Figure 3, we assume unit demand stations and indifferent customer choice but not player homogeneous payoffs. There are two players and five competitive delivery stations; each of the two players has five

Figure 2. Example of a C-PDOP Instance with No Pure Nash Equilibria

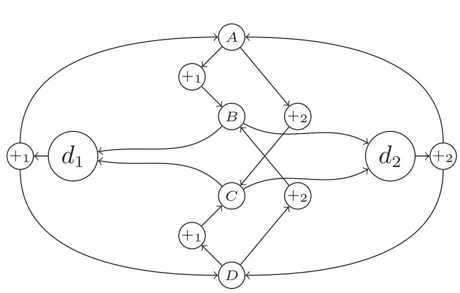

Split	1	2
A	3/4	1/4
B	1/4	3/4
C	1/4	3/4
D	3/4	1/4

	0	AC	DB
0	$0, 0$	$0, 2$	$0, 2$
AB	$2, 0$	$\frac{7}{4}, \frac{5}{4}$	$\frac{5}{4}, \frac{7}{4}$
DC	$2, 0$	$\frac{5}{4}, \frac{7}{4}$	$\frac{7}{4}, \frac{5}{4}$

Notes. There are four delivery unit capacity districts, labeled A, B, C, D. Each player has a depot (d_1 or d_2) and vehicles at distinct pickup locations ($+_1$ and $+_2$, respectively) with zero payoff. Player 1 has essentially three strategies (null, AB tour, and DC tour). Player 2 has essentially three strategies (null, AC tour, and DB tour). The payoffs are player homogeneous with differentiated customer choice (center). Looking at the payoff matrix (right), we can see that no Nash equilibrium exists.

vehicles at pickup stations. All arcs that appear in the network have equal travel cost of two. The arcs that do not appear in the network have a travel cost given by the induced directed graph metric (so, for example, the distance from d_1 to location 3 equals 12 because there is a six-arc path from d_1 to location 3 in the original graph). Thus, the problem instance even satisfies the triangular inequality. Note that player 1 routes on the left half of the graph and has essentially three strategies: do nothing, service locations (1,2,3), or service locations (1,4,5). Similarly, player 2 routes on the right half of the graph and has essentially three strategies: do nothing, service locations (5,4,3), or service locations (5,2,1). By our choice of arc costs, each of the nontrivial strategies has a travel cost of 14.

Finally, we choose the station payoffs as in Figure 3. Because we assume indifferent customer choice, this means that the payoff from a player at a station is halved if the other player also places a vehicle there. For illustration purposes, let us compute the gross profit if player 1 takes the "short" tour, servicing (1,2,3), and player 2 takes the "long" tour, servicing (5,2,1). In this situation, player 1 extracts full revenue from station 3,; player 2 extracts full revenue from station 5, and both players extract half revenue from stations 1 and 2 each. We, thus, have

$$\Pi_1 = \frac{13}{2} + \frac{2}{2} + 9 - 14 = \frac{1}{2};$$

$$\Pi_2 = \frac{9}{2} + \frac{5}{2} + 11 - 14 = 4.$$

Note that player 2 would rather deviate to the short tour and increase gross profit. By doing these calculations for all possible pairs of strategies, we see that no Nash equilibrium exists; in particular, any sequence of iterated best responses gets stuck in the loop $(S,S) \to (L,S) \to (L,L) \to (S,L) \to (S,S)$.

The counterexamples presented here are handcrafted to show that pure strategy Nash equilibria may not exist. For these examples, the structure of competitive locations and tours exhibits many symmetries, and revenues vary vastly between the operators and different locations. We could argue that this situation is in some sense "atypical" and would not arise in "realistic" networks. From a practical perspective, we observe in Sections 5.5–5.7 that pure strategy Nash equilibria can be found in most instances.

3.4. Properties of the Competitive Pickup and Delivery Orienteering Problem

We first observe the computational hardness of solving the PDOP (and, thus, calculating optimal routes of the C-PDOP).

Theorem 1 (NP-Hardness). *The following problems are all NP-hard even for the special case of two players.*

1. *Given a PDOP instance, compute a tour that maximizes the gross profit.*
2. *Given a C-PDOP instance, a (compact) joint strategy* $\mathbf{q}$*, and a player n, compute* $\Pi_n(\mathbf{q})$*.*
3. *Given a C-PDOP instance, a player n, and a joint strategy* $\mathbf{q}_{-n}$ *of the remaining players, compute an optimal response for player n, that is, a strategy* $\mathbf{q}_n$ *maximizing* $\Pi_n(\mathbf{q}_n, \mathbf{q}_{-n})$*.*
4. *Given a C-PDOP instance and a joint strategy* $\mathbf{q}$*, determine whether* $\mathbf{q}$ *is a Nash equilibrium.*

In the following, we focus on Nash equilibria; that is, we want to establish whether they exist and if they can be reached via "simple" dynamics. Because the restricted C-PDOP is a congestion game (Lemma 1), the existence of *pure strategy* Nash equilibria can be guaranteed. Congestion games are formally described by a set of players (in our model, the different operators) and a set of resources (in our restricted model, a resource for each location and a resource for each cost-minimizing tour). Each resource has an associated payoff function that depends only on *the number of* players accessing it (in our restricted model, the payoff achieved at a given location decreases with the number of players servicing it). Finally, each player has a set of feasible strategies or subsets of resources (in our model, this corresponds to the set of feasible tours), and the payoff incurred by a

Figure 3. Example of a C-PDOP Instance with No Pure Nash Equilibria

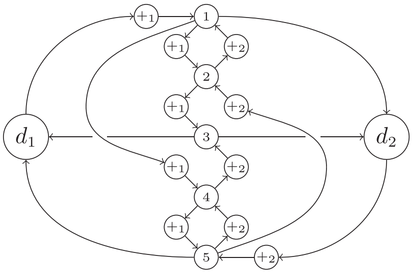

i	1	2	3	4	5
π_i^1	13	2	7	3	9
π_i^2	9	5	7	4	11

	0	S	L
0	0, 0	0, 8	0, 11
S	8, 0	4.5, 4.5	0.5, 4
L	11, 0	5, 0.5	0, 1

Notes. There are five delivery locations, labeled 1–5, which are all competitive. Each player has a depot (d_1 or d_2) and five vehicles at distinct pickup locations ($+_1$ and $+_2$, respectively) with zero payoff. Each player has essentially three strategies (null, short tour, and long tour). For a specific choice of location payoffs (center) and payoff matrix (right), no Nash equilibrium exists.

player at a joint strategy is just the sum of the payoffs of the strategies the player is using.

Lemma 1 (Congestion Game). *Any restricted C-PDOP instance can be transformed into a congestion game. This transformation induces a one-to-one correspondence between (compact) strategies in the C-PDOP instance and strategies in the congestion game that preserve improving deviations and, hence, the structure of Nash equilibria.*

Because of the generality of the definition of the congestion game and the proof (see Appendix A), this construction of the congestion game is valid for many different routing problems as long as the cost term $C_n(\mathbf{q}_n)$ depends only on the strategy of operator n. In particular, our construction can be adapted to integrate maintenance and refueling or recharging operations or to address other rebalancing modes, such as using a truck or minibus. From this congestion game formulation, the existence of pure strategy Nash equilibria follows directly (because of Rosenthal 1973, Monderer and Shapley 1996).

Corollary 1 (Existence of Pure Strategy Nash Equilibria). *Any restricted C-PDOP instance has at least one pure strategy Nash equilibrium. Moreover, for any starting strategy $\mathbf{q}^{(0)}$, any sequence $\mathbf{q}^{(0)} \to \mathbf{q}^{(1)} \to \ldots$ obtained by improving deviations (i.e., in which $\mathbf{q}^{(i+1)}$ is obtained from $\mathbf{q}^{(i)}$ by an improving deviation of any one player) must eventually terminate at a Nash equilibrium.*

We can prove not only the existence of pure strategy Nash equilibria, but also the convergence of best-response dynamics to Nash equilibria. We leverage this property in the iterated best response algorithm in Section 4. The Nash equilibrium is not necessarily unique; examples are discussed in Section 5.

3.5. Models for Comparison

To measure the impact of considering competition in the servicing and routing decision, we compare the optimal gross profit of the C-PDOP instance to problem variations in which (i) competitors collaborate to maximize the overall gross profit but still only serve their own customers (W-PDOP), (ii) operators merge (M-PDOP) or cooperate in their fleet relocation operations (e.g., by outsourcing to the same third-party provider; Coop-PDOP), or (iii) competitors ignore each other (quasi-monopolistic PDOP, optimistic: QMO, pessimistic: QMP).

3.5.1. Welfare-Maximizing PDOP. We can measure the losses in gross profit resulting from competition by comparing the results of the C-PDOP to the system optimum, that is, the welfare-maximizing strategy. Intuitively, this corresponds to a situation in which the operators relocate their fleets by themselves, but only the operator who benefits most from a competitive location serves it. We merge all constraints of Model (1a)–(1k) and add up the objective functions of each player:

$$\max \Pi(\mathbf{q}) = \sum_{n} \Pi_n(\mathbf{q}). \tag{4}$$

Regarding the gross profits, we observe the following:

Corollary 2 (Monotonicity of Profits (Welfare-Maximizing Solution)). *The optimal gross profit of the welfare-maximizing solution is never less than the joint gross profit of the competitive solution.*

The welfare-maximizing gross profit can, therefore—similar to the costs in the shared customer collaboration vehicle routing problem (Fernández, Roca-Riu, and Speranza 2018)—serve as an upper bound for the gross profit attainable in a Nash equilibrium.

For the restricted C-PDOP model, we can analytically quantify the loss of efficiency, interpreted in terms of the price of anarchy (Koutsoupias and Papadimitriou 1999, Roughgarden and Tardos 2002) and price of stability (Schulz and Stier Moses 2003, Anshelevich et al. 2008), which is defined as the worst-case bound on the ratio between the worst (respectively, best) joint gross profit of a pure strategy Nash equilibrium and the joint gross profit of a welfare-maximizing solution.

Lemma 2 (Price of Anarchy and Price of Stability). *The price of anarchy and price of stability of the C-PDOP can be arbitrarily large. For the restricted C-PDOP model with two players, we have the following results.*

1. *If $(\mathbf{q}_1, \mathbf{q}_2)$ is a pure strategy Nash equilibrium and $(\mathbf{q}'_1, \mathbf{q}'_2)$ is any strategy, then $\Pi(\mathbf{q}'_1, \mathbf{q}'_2) \le \Pi(\mathbf{q}_1, \mathbf{q}_2) + \frac{1}{2}\sum_{i\in\mathcal{Z}^C}\pi^i$.*
2. *The absolute difference in welfare between any two Nash equilibria is at most $\frac{1}{2}\sum_{i\in\mathcal{Z}^C}\pi^i$.*
3. *The absolute difference in welfare between any Nash equilibrium and any welfare-maximizing strategy is at most $\frac{1}{2}\sum_{i\in\mathcal{Z}^C}\pi^i$.*

Moreover, all bounds in points 1–3 are tight, and if $\pi^i = \pi \ \forall i \in \mathcal{Z}^C$, then all bounds in points 1–3 can be replaced by $\frac{1}{2}\pi\,|\mathcal{Z}^C|$.

The proof of points 1–3 in the lemma relies on a simple expression for the player gross profits in the restricted C-PDOP model; as such, it cannot trivially be extended to the more general case.

3.5.2. Monopoly PDOP and Coopetition PDOP. In the PDOP master model, we assume that only one operator is present, whereas in the competitive C-PDOP model, each operator relocates the operator's fleet separately while taking into account the strategy of the other operators. We now consider two alternative business models: the M-PDOP in which the competitors merge their fleets with the objective of reducing travel costs and gross profit losses because of competition and the Coop-PDOP in which the competitors combine their relocation efforts. In order to avoid the combinatorial explosion associated with all possible merge combinations between operators, in this section, we focus exclusively on the C-PDOP model with unit demand stations, indifferent customer choice (but not necessarily homogeneous payoffs), and two operators that consider merging or cooperating.

In the M-PDOP model, vehicles become indistinguishable. To construct an M-PDOP instance from a C-PDOP instance, we merge all locations for all operators in a station ι and devise a joint payoff function π^ι under the assumption that every customer returns the highest payoff over all operators ($\pi^\iota(\mathbf{q}) = \max_{\mathbf{q}'}\left(\sum_n \pi_n^\iota(\mathbf{q}') \mid \sum_n q_n^\iota = \sum_n q_n^{\iota'}\right)$. In the case of unit demand stations, servicing a location i contributes a payoff of $\pi_i = \max(\pi_i^1, \pi_i^2)$ in a monopoly. In general, the monopolist can keep both depots. In the numerical experiment, however, we assume that the depots are at the same geographical location and can, therefore, be merged.

In the Coop-PDOP, operators can collaboratively relocate their fleets (entering coopetition). This may be considered an alternative model when merging the fleets is not an option (because of strategic considerations or cartel law). One of the companies or a third party relocates the vehicles of both fleets, thereby maximizing the sum over both profit functions (with the assumption that a cost or profit-sharing mechanism is implemented at a later stage). The two key differences between this and the monopoly solution are that the payoff achieved in a competitive station depends on which operator(s) serve it, and noncompetitive delivery locations can only be reached from pickup locations of the same player (so, for example, one cannot generate revenue by moving a vehicle of player 1 to a location where only a customer of player 2 is expected). Similarly to the M-PDOP, we model the Coop-PDOP by devising a joint payoff function π^ι, which returns the sum of payoffs achievable in this station given the number of vehicles each operator places there ($\pi^\iota(\mathbf{q}) = \sum_n \pi_n^\iota(\mathbf{q})$). We also exclude arcs that connect a pickup location of one player with a delivery location of another player.

Note that it is possible (in both Coop-PDOP and M-PDOP models) that a vehicle of operator 1 is moved to some competitive station ι, and at the same time, a vehicle of operator 2 is removed from there. This can occur because of both inhomogeneous payoffs (the location is more attractive to player 1 than player 2) and a highly profitable station of player 2 elsewhere.

In practice, the number of workers in the monopoly, W_M, or in coopetition, W_O, often equals the number of workers under competition, but we do not restrict the model as such. As the M-PDOP and Coop-PDOP contain the PDOP as special cases, we immediately have the following results concerning their computational hardness.

Corollary 3 (NP-Hardness of M-PDOP and Coop-PDOP). *The following problems are NP-hard.*

1. *Given an M-PDOP instance, compute a tour that maximizes gross profit.*
2. *Given a Coop-PDOP instance, compute a tour that maximizes gross profit.*

Intuitively, one would assume that the gross profits of monopoly and coopetition solutions consistently exceed the gross profit of Nash solutions. Although this is not always true (for example, if the players serve distant operating areas with disjoint depots), we provide some assumptions that guarantee the validity of this intuition.

Lemma 3 (Monotonicity of Profits (Monopoly or Coopetition Solution)). *The following are true for the C-PDOP model and its monopoly and coopetition variants.*

1. *The optimal gross profit of a monopoly solution is not less than the optimal gross profit of the coopetition solution.*
2. *If the number of workers is at least the sum of people working for the first and second operator ($W_M, W_O \geq W_1 + W_2$), the optimal gross profit of the monopoly or coopetition solution is not less than the optimal gross profit attainable in any pair of strategies (which includes all pure strategy Nash equilibria as well as the welfare-maximizing solution).*

Thus, in realistic settings, competing is inferior to cooperation with respect to (short-term operational) gross profits.

3.5.3. Quasi-Monopolistic PDOP. We identify two slightly different strategies that model that either competition or the rationality of the competitor is ignored: First, we solve the PDOP assuming that the other operators have no vehicles available in any station and do not rebalance to these locations either (optimize against $\mathbf{q}_n = \langle 0, \ldots, 0 \rangle$ for each competitor n). We call this model the QMO. Second, we solve the PDOP assuming that the other operator has vehicles available at all locations after rebalancing (optimize against $\mathbf{q}_n = \langle \hat{q}_n^1, \ldots, \hat{q}_n^{|D_n|} \rangle$ for each competitor n, where $\hat{q}_n^\iota$ denotes the maximum number of vehicles that operator n can move to station $\iota \in D_n$). This model is called the QMP. Although both QMO and QMP seem to be good candidates for serving as lower bounds on the gross profit of either player as well as on the total gross profit, it is possible to generate instances in which these problems result in higher payoffs.

4. Algorithms for the Nash Equilibrium Calculation

The most basic approach for finding pure strategy Nash equilibria is full enumeration (FE): iterate through all strategy profiles, that is, combinations of strategies of all operators, and test if no player has incentive to unilaterally deviate. If this is true, the strategy profile constitutes a Nash equilibrium. Although this approach obviously finds the best pure strategy Nash equilibrium (with respect to social welfare or any other metric) whenever an equilibrium exists, it takes exponentially long in the number of competitive locations. We, therefore, consider alternative approaches: IBR and PFO. The two algorithms represent two approaches for finding a Nash equilibrium in congestion games: utilizing the improvement dynamics of alternately improving players to find a local optimum of the potential (IBR) or (centrally) finding the global optimum of the potential (PFO). Both algorithms find a Nash equilibrium in congestion games, but they do not necessarily find a welfare-maximizing Nash equilibrium. If the congestion game property is violated, the potential is undefined, but IBR might still be able to find pure strategy Nash equilibria if they exist.

4.1. Iterated Best Response Algorithm

Using the IBR, we locally search for a pure strategy Nash equilibrium. We first calculate the optimal strategy for one of the players (say, player 1) against a predefined strategy for the competitors, for example, assuming the competitors play the empty strategy $\mathbf{q}_n = \langle 0, \ldots, 0 \rangle$ (and, hence, do not have any vehicles at any competitive location). We then use the strategy of player 1 as input for calculating the optimal strategy of player 2, then player 3. We continue with our calculations of best responses until the strategies no longer change. Although the best response iterations cannot be implemented in the field, we assume that operators would calculate the Nash equilibrium theoretically and implement their equilibrium strategy.

Even if the IBR terminates, it may not necessarily return the best Nash equilibrium for one of the players or a welfare-maximizing Nash equilibrium. Yet, because of Lemma 2, we know that any two Nash equilibria do not differ by more than $\frac{1}{2}\sum_{i\in\mathcal{Z}^C}\pi^i$ for the restricted C-PDOP model with two players. Thus, implementing the IBR does not result in arbitrarily bad Nash solutions in such cases. Though operators do not necessarily find the "best" Nash equilibrium, we guarantee that, for two operators, the Nash equilibrium found using IBR is at least as good (for either player) than the optimistic quasi-monopolistic strategy.

Theorem 2 (Monotonicity of Profits (IBR vs. QMO)). *For the restricted C-PDOP model with two players, the gross profit of each player at the Nash equilibrium reached by IBR, starting from the $\langle 0, \ldots, 0 \rangle$ strategy, is at least as high as the gross profit with the optimistic quasi-monopolistic strategy.*

Thus, operators who currently calculate their routes using the optimistic quasi-monopolistic strategy can only benefit from calculating the Nash equilibrium. Conversely, however, we can generate instances in which the pessimistic quasi-monopolistic strategy beats the IBR solution.

Because the restricted C-PDOP is a congestion game, IBR terminates in a finite number of iterations. If one aims at implementing the IBR in practice, the number of iterations required to reach an equilibrium is critical. General congestion games with an arbitrary number of players belong to the class of polynomial local search–hard games. Thus, even if there exists a polynomial time algorithm for finding pure strategy Nash equilibria, the solution cannot, in general, be found in polynomial time by myopic players (Fabrikant, Papadimitriou, and Talwar 2004; Ackermann, Röglin, and Vöcking 2008). The C-PDOP, however, differs from general congestion games: the number of players is low. By presenting an upper bound on the number of iterations, we show that, assuming homogeneous payoffs, unit demand stations, and indifferent customer choice, the IBR for the C-PDOP does not require full enumeration of all strategies.

Theorem 3 (Termination of the Iterated Best Response Algorithm). *For the two-player restricted C-PDOP model, the IBR terminates after one player plays at most half of the player's strategies (thus, the maximum number of required recalculations is $|S| + 2$ instead of the $|S|^2$ we have in the FE algorithm).*

However, the upper bound on the runtime remains substantial as the number of strategies is exponential in the number of competitive stations, and obviously, no bound can be given in the general model because pure strategy Nash equilibria might not exist.

These results all require exact responses, that is, finding an optimal solution to the PDOP rather than a near-optimal feasible solution. However, some results also still hold when we relax the notion of optimality.

Theorem 4 (Iterated Best Response Algorithm for Approximate Nash Equilibria). *Let $s_n = (\mathbf{a}^n, \mathbf{x}^n)$ denote a full strategy of player n, that is, describing the servicing and routing decisions for all arcs and locations. For any*

$\epsilon > 0$, consider the following version of ϵ-approximate iterated best response (ϵ-IBR): given a player n and a full strategy $\mathbf{s}_{-n}$ of the other players, player n can compute a full strategy s_n with the property that

$$(1+\epsilon)\Pi_n(s_n, \mathbf{s}_{-n}) \geq \max_{s'_n} \Pi_n(s'_n, \mathbf{s}_{-n});$$

note that this amounts to finding an approximate solution of the PDOP problem in Equations (1a)–(1k). Player n can then choose to deviate to s_n if this improves gross profit.

For the restricted C-PDOP model, ϵ-IBR always terminates after a finite number of iterations, and the final joint strategy $\mathbf{s}$ is an ϵ-Nash equilibrium in the sense that, for any player n,

$$(1+\epsilon)\Pi_n(s_n, \mathbf{s}_{-n}) \geq \max_{s'_n} \Pi_n(s'_n, \mathbf{s}_{-n}).$$

Though no polynomial-time approximation of the PDOP is known, we can use this result to get a posteriori bounds on the quality of approximate equilibria. For example, if a commercial solver that implements branch-and-cut procedures obtains a solution with a provable optimality gap of $(1+\epsilon)$, this immediately implies an $(1+\epsilon)$ guarantee on the quality of the equilibrium.

4.2. Potential Function Optimizer

Another approach toward finding a Nash equilibrium in congestion games is optimizing the potential function that can be optimized using a standard solver. This is a global function that captures the local incentives for players to change their strategies and, therefore, a useful tool for analyzing equilibria. In particular, if a game admits a potential function, the Nash equilibria of the game coincide with the local optima of the potential function.

Because, in the general C-PDOP model, Nash equilibria are not guaranteed to exist, a potential function cannot be defined. However, for the restricted C-PDOP, it is possible to define the potential Φ as the sum

$$\Phi(\mathbf{q}) = \sum_{i \in \mathcal{Z}} H_{y^i}\pi^i - \sum_n C_n(\mathbf{q}_n), \tag{5}$$

where y^i is the number of operators servicing location i, π^i is the revenue that can be extracted from location i, and $H_k = 1 + 1/2 + \cdots + 1/k$ is the k th harmonic number. Notice that the potential is not the same as the social welfare (4): the potential function associates a payoff of $H_{y^i}\pi^i$ with competitive locations where y^i players are available, and the social welfare more realistically assumes that only one of them is able to service the customer. Intuitively, this means that the PFO tends to select equilibria in which multiple operators have a vehicle available at competitive locations.

Although a global optimum of the potential function is always a pure strategy Nash equilibrium, it is not necessarily a welfare-maximizing one. However, we can characterize instances in which both coincide.

Lemma 4 (Optimality of the Potential Function Optimizer). *For the restricted C-PDOP model (which has a well-defined potential function Φ),*

1. *Any potential function maximizer $\max_{\mathbf{s}} \Phi(\mathbf{q})$ is a Nash equilibrium.*
2. *If the PFO returns a Nash equilibrium in which at most one operator has a vehicle available at any competitive location ($\sum_n q^i_n \leq 1 \ \forall i \in \mathcal{Z}^C$), this Nash equilibrium is welfare-maximizing.*

Thus, the PFO is likely to return the welfare-maximizing Nash equilibrium if revenues are low, costs are high, and margins are tight as Nash equilibria mostly do not to contain locations where both operators are present.

5. Computational Study

In the following, we quantify the average-case gross profit gains and losses, not only for the restricted C-PDOP, but also for the generality of the model. Unless stated otherwise, we focus on player homogeneous payoffs, unit demand stations, indifferent customer choice, and two operators. Further, we conduct a sensitivity analysis if the number of operators increases, if payoffs are not player homogeneous, if stations are multidemand with decreasing marginal returns, and if customers are not strictly indifferent between operators. To quantify gains and losses, we present a case study featuring the competition between two major carsharing operators in Munich, Germany.

5.1. Experimental Design

We conduct our experiments on a Windows 10 computer restricted to a single 2.60 GHz core of an Intel Xeon E7-4860 CPU with 4 GB of RAM. We implement both algorithms in Java 10, using CPLEX 12.8 for solving the PDOPs. We start IBR against two different strategies: assuming that the competitors are absent (IBR-0) or starting against the welfare-maximizing strategy (IBR-WP).

To study the effects of competition and the various business models, we randomly generate 100 data sets for different combinations of the parameters mentioned in Table 1. When studying multiplayer settings, inhomogeneous payoffs, multidemand stations with marginal returns, and various different customer choice models, some of the parameters have to be defined slightly differently. These changes are introduced at a later stage. The parameter levels are motivated by the Munich car-sharing market. For each instance, we

Table 1. Parameters for the Experimental Design in the Base Case

Parameter		Level 1		Level 2
Substitution	F	Full Substitution	P	Partial Substitution (25%)
Margin	H	$\pi = 8$	L	$\pi = 4$
Density	H	$\lvert \mathcal{V}_n^+ \rvert = 12$	L	$\lvert \mathcal{V}_n^+ \rvert = 8$

randomly sample locations on a square with an edge length of 15 km and use Euclidean symmetric costs with a weight of 0.62 (25 km/h traffic speed, 12,5€/h worker wages (Wittenbrink 2014), 0.12€/km for fuel).

Substitution refers to the share of competitive locations, that is, if all delivery locations are shared (F) or if shared locations are randomly sampled (P). Albiński and Minner (2019) report that approximately 25% of all customers in Munich have multiple memberships; in the P level, we set 25% of all delivery locations as shared.

To quantify the impact of changes in the payoff/cost structure, we vary the payoff. In the high-margin scenario, we set the margin of all delivery and competitive pickup locations to $\pi = 8$ (1 otherwise), and in the low-margin scenario, we set the margin to $\pi = 4$ (0.5 otherwise). The payoffs that can be collected at delivery locations if no vehicle is available constitute the baseline profit toward which all relative improvements are measured. Pickup locations are associated with a small (but positive) payoff because of a low (but nonzero) probability that a customer rents a car from these locations.

The customer density refers to the number of locations that enter the model (possibly contingent to prior filtering). We generate instances with a high density ($\lvert\mathcal{V}_n^+\rvert=\lvert\mathcal{V}_n^-\rvert= 12$) and instances with a low density ($\lvert\mathcal{V}_n^+\rvert=\lvert\mathcal{V}_n^-\rvert= 8$). The customer density can vary between the operators. In all instances, we use one worker. Further, we assume that all locations are optional, that is, $S_n^1 = S_n^0 = \emptyset$.

5.2. Profit Increase Because of Considering the Presence of Competition

Table 2 lists relative gross profits (toward the baseline, that is, no rebalancing) for operators 1 and 2 for all parameter combinations if margins of both operators coincide, that is, are either high or low for both operators. More extensive results can be found in Appendix B in Table B.1. There, instances are addressed by four consecutive letters referring to substitution, margin (same for both operators), and the density of the first and second operator. For example, *F_H_L_L* refers to full substitution, high margins, and low densities for either operator. In particular, this allows us to study if the larger or smaller operator benefits more.

Averaging over all full competition instances, operators are even better off not to rebalance at all than to ignore the competitor (QMO). QMO even generates losses in five (operator 1)/five (operator 2) of eight instances with full competition, and all other approaches result in nonnegative gross profit gains (in all instances in which the other does not have lower margins). Thus, it makes sense to incorporate competition in the routing and servicing decision. However, Nash solutions outperform both quasi-monopolistic solutions. As to be expected, the improvement over QMO/QMP increases in the level of competition because more locations are shared. The maximum attainable gross profit gain is 251% under partial competition and if the second operator is larger than the first mover (versus no rebalancing operations). These very high relative values stem from the fact that the baseline and all absolute values are rather low and that (in particular, under full competition) quasi-monopolistic solutions often involve losses.

As is visible from Table 3, the player with more competitive locations can generate higher gross profit gains when operators have different numbers of vehicles to relocate. This is because the larger operator has more nonshared locations (and can, thus, build a more efficient route). This effect partially alleviates the disadvantage of the second mover. The smaller operator, however, has the larger benefit of considering competition. This is because the larger operator serves most competitive locations, and the small operator benefits from moving vehicles to the few remaining locations. It also becomes apparent that, even though gross profits compared with the baseline increase if the network becomes more dense, the relative benefit of considering

Table 2. Average, Minimum, and Maximum Percentage Profit Increase Toward Baseline (No Rebalancing) Under Various Models and Algorithms as Well as Different Experimental Settings for Either Operator (Operator 1 in Left Block, Operator 2 in Right Block)

	Operator 1					Operator 2				
Setting	IBR-0	IBR-WP	PFO	QMO	QMP	IBR-0	IBR-WP	PFO	QMO	QMP
Average (full competition)	83.6	72.4	71.2	−2.77	40.0	52.2	67.9	69.8	−2.11	39.7
Minimum (full competition)	5.51	4.88	4.88	−63.2	1.04	3.66	3.93	4.62	−61.7	1.15
Maximum (full competition)	225	204	206	85.9	180	191	207	215	103	199
Average (partial competition)	111	110	109	88.3	101	109	111	113	92.2	104
Minimum (partial competition)	6.74	6.74	6.46	3.45	4.89	8.28	8.28	8.92	4.49	5.54
Maximum (partial competition)	243	243	240	210	231	246	246	251	226	243

Table 3. Average Percentage Profit Increase Toward Baseline (No Rebalancing) Under Various Models and Algorithms if Operators Have Different Sizes (Operator 1 in Left Block, Operator 2 in Right Block)

	Operator 1					Operator 2				
Setting	IBR-0	IBR-WP	PFO	QMO	QMP	IBR-0	IBR-WP	PFO	QMO	QMP
H_H	113	105	103	47.6	81.1	89.5	103	103	48.4	79.9
H_L	130	125	125	84.6	109	49.9	55.3	59.8	13.9	40.5
L_H	65.6	61.3	60.7	15.3	44.3	121	126	130	91.8	117
L_L	80.8	73	71.7	23.6	47.7	60.9	73.9	73.6	25.9	51.1

competition decreases as QMO becomes more profitable. Averaging over all settings with two large operators, gross profit gains over the baseline increase from 79.9% (QMP) to 103% (IBR-WP/PFO) for operator 2 (28.9% increase), and for two small operators, the gross profit gain increases from 51.1% to 73.9% (44.6% increase). These high relative values are due to low absolute values, for example, the gross profit gain for operator 1 increases from 0.26 (QMP) to 0.43 (IBR-0) over the baseline in the setting with low density and low margins for both operators. In absolute numbers, however, the benefit of considering competition continues to increase.

We observe that IBR-0 privileges operator 1 over operator 2, and the other algorithms (IBR-WP, PFO, QMO, QMP) do not give a clear advantage to either player. This makes IBR-0 the best algorithm for player 1. For the second player, IBR-0 is outperformed by IBR-WP and PFO in almost all instances. IBR-WP tends to return higher gross profits than PFO for player 2 if both operators have the same size, and PFO tends to return higher gross profits if one player is larger than the other. A similar pattern can be observed with respect to welfare (sum over gross profits of both players). In most instances, QMO returns lower gross profits than QMP, but exceptions exist if the gross profit is low.

5.3. Profit Loss Because of the Presence of Competition

Table 4 lists the gross profits of the best found Nash equilibrium (profit-maximizing among IBR-0, IBR-WP, and PFO) compared with the welfare-maximizing solution, the coopetition solution, and the monopoly solution. Extended results can be found in Table B.2. There, the substitution level is addressed in the column header, and every row refers to margin (same for both operators), density of operator 1, and density of operator 2.

In general, obviously, all instances follow similar tendencies: gross profits increase in the number of vehicles that are rebalanced (either because of an increasing customer demand or increasing demand imbalance) and with increasing margins but decrease if competition increases. Although joint fleet management (monopoly or coopetition) results in a substantial gross profit increase, the benefit of welfare-maximization is little (consistently less than 2% under partial competition) and does not justify the additional coordination requirement. We observe a tendency that the percentage gap between the Nash solution and the other approaches closes with increasing instance sizes while absolute gaps continue to grow. This effect is less pronounced in the full competition case because pooling effects do not improve as much as in the partial competition case. This is mainly driven by better routing decisions because of larger pooling effects in the Nash solution. For Coop-PDOP and M-PDOP, we observe that full competition results in a lower improvement than partial competition. This might seem counterintuitive at first but can be explained as follows: In the partial competition case, the benefits of pooling increase as the total number of vehicles is higher. Thus, Coop-PDOP and M-PDOP find more efficient routes. In some cases, the routes of the M-PDOP/Coop-PDOP and the Nash equilibrium even coincide. High margins decrease the relative gross profit loss from competing because many

Table 4. Average, Minimum, and Maximum Percentage Profit Increase Toward Baseline (No Rebalancing) for Different Experimental Settings for Either Player (Full Substitution in Left Block, Partial Substitution in Right Block)

	Full competition				Partial competition			
Setting	NE	W-PDOP	Coop-PDOP	M-PDOP	NE	W-PDOP	Coop-PDOP	M-PDOP
Average	75.6	82.7	107	110	114	115	171	199
Minimum	6.89	6.89	11.2	11.2	10.1	10.1	37.7	59.6
Maximum	155	169	204	209	237	240	304	331

Table 5. Average Percentage Profit Increase over All Operators Toward Baseline (No Rebalancing) for an Increasing Number of Operators

		Two operators			Three operators			Four operators			Five operators		
Setting	Baseline	IBR-0	QMO	QMP	IBR-0	QMO	QMP	IBR-0	QMO	QMP	IBR-0	QMO	QMP
F_H_H	12	118	−2	61	77	−127	1	57	−171	3	48	−178	2
F_H_L	8	92	−34	34	67	−166	0	43	−212	0	44	−245	0
F_L_H	6	12	−2	2	19	−53	1	19	−81	0	11	−84	0
F_L_L	4	13	−3	5	5	−5	0	4	−1	0	5	−2	0
P1_H_H	12	216	164	203	237	191	221	237	211	231	232	196	221
P1_H_L	8	144	110	135	174	130	161	168	140	158	158	120	150
P1_L_H	6	17	11	12	23	9	19	32	28	30	20	16	13
P1_L_L	4	8	5	4	7	1	5	7	7	5	9	7	6
P2_H_H	12	225	173	213	217	155	205	218	151	208	201	137	195
P2_H_L	8	159	114	147	144	89	137	131	73	121	145	89	135
P2_L_H	6	28	13	19	18	2	12	20	−2	13	21	4	13
P2_L_L	4	10	10	7	5	5	4	6	0	3	4	0	2

customers are served in the competitive solution, and profit differences must, thus, be attributed to improved routing (and improved pooling does not contribute as much gross profit as serving additional customers).

5.4. Impact of an Increasing Number of Players

All previous experiments consider only two operators. Although this is sufficient to model competition in some cities, there are markets with more car-sharing operators.

Following the numerical design outlined in Table 1, we report average gross profit increases over all players for IBR-0, QMO, and QMP for an increasing number of operators (column title in Table 5), different levels of substitution, different densities, and different margins (all operators have the same density and margin to ensure comparability across different numbers of operators).

We observe that effects studied for the two-operator case get more pronounced as the number of operators increases, but tendencies remain the same. The gross profit increase over the baseline under all three models decreases if the number of operators increases and competition is either full or locations are shared among a subset of operators. This is to be expected because each operator services fewer customers on average. If the number of operators increases, it becomes even more critical to consider competition as ignoring competition results in significant losses (up to 245% for full competition and five operators) and, assuming that the competitors service all locations, results in refraining from any rebalancing if the number of operators increases. When considering the gain relative to the case in which all competitive locations are serviced by the competitors, the improvement of considering competition slightly decreases if the number of operators increases but remains substantial in all instances.

5.5. Impact of Inhomogeneous Payoffs

If payoffs for players are inhomogeneous, that is, differ between players and locations, pure strategy Nash equilibria do not provably exist. However, in many instances, equilibria appear nonetheless and can be found using IBR. For two players, we consider full and partial substitution, and either player can have high or low location density, following the numerical design outlined in Table 1. We alter the definition of margins because the case $\pi_i^1 = k \cdot \pi_i^2$ is a special case of homogeneous payoffs in which pure strategy Nash equilibria provably exist. Instead, margins for either player are randomly drawn from a high ($\pi \in [6, 10]$) or low ($\pi \in [2, 6]$) interval. With 100 repetitions of 32 instances, Nash equilibria existed in all cases, which is partially because of the full graph with Euclidean costs.

The most central results for inhomogeneous payoffs are depicted in Table 6, and all results for 32 different

Table 6. Average Percentage Profit Increase Toward Baseline (No Rebalancing) Considering Inhomogeneous Payoffs

Setting	IBR-0 1	IBR-0 2	QMO 1	QMO 2	QMP 1	QMP 2
Average (full competition)	167	111	−41	−44	88	83
Minimum (full competition)	−26	−47	−237	−239	−77	−66
Maximum (full competition)	404	371	236	226	358	365
Average (partial competition)	311	289	311	289	311	289
Minimum (partial competition)	157	122	157	122	157	122
Maximum (partial competition)	434	433	434	433	434	433

instances are reported in Table B.3. Compared with the case of homogeneous payoffs, effects become more pronounced: gross profit increases over the quasi-monopolistic solutions become larger for full competition (on average, the improvement over the baseline is twice as high for the Nash equilibrium than for QMP) although, under partial competition, considering competition never improves the solution (compared with very small improvements in the homogeneous payoffs case). When ignoring the presence of competition, the operator with lower contribution margins often faces gross profit losses. This operator can circumvent or at least partially counteract these by considering competition.

5.6. Impact of Stations with Diminishing Marginal Payoffs

To study the impact of diminishing payoffs in larger stations, we fix the density at a high value (12) and assign the vehicles to stations with varying size $\omega \in [1,2,3]$, where 1 is the case without diminishing payoffs. We consider full and partial substitution and assume that every station is either competitive or noncompetitive, but there are no stations with some competitive and some noncompetitive locations. In absence of any competition, margins are given by

$$\pi_\iota(0, q_n^\iota) = \sum_{i=1}^{q_n^\iota} \pi_\iota^\star \cdot \lambda^{(i-1)},$$

and the average margin over all locations in a station is eight (high, H) or four (low, L), which results in a maximum margin of $\pi_\iota^\star = \frac{\pi}{\omega}\sum_{i=1}^{\omega}\lambda^{(i-1)}$, where π is the average margin. If multiple operators have vehicles at station ι, gross profits are split fairly depending on the number of vehicles either operator has at this station. $\lambda \in \{0.5, 0.7, 0.9\}$ is the deterioration rate. Instances are then addressed by substitution (full or partial), margin (high or low), deterioration rate, and station size. For example *F_H_0.7_3* refers to the instance with full substitution, high margins, medium deterioration rate (0.7), and three locations per station. In Table 7, the number of locations per station is moved to the column head.

In total, no Nash equilibria was found in two cases (out of 100 repetitions of 36 different instances). Both affected instances have three locations per station and high margins but different levels of substitution and deterioration rates. With an increasing number of locations per station, the benefit of considering competition decreases slightly as the improvement over the baseline decreases for the Nash equilibrium but increases for the quasi-monopolistic solutions. The former can be explained by moving not too many vehicles to the same station to achieve high payoffs per location, and the latter occurs because operators gain some payoff from stations even if the other operator is also having vehicles there. This effect is not very strong except that QMO does not result in a negative gross profit increase over the baseline if the number of locations exceeds one. Interestingly, the gap between the two operators' gross profits closes with an increasing number of locations per station (under full competition, it decreases from a factor of two difference to approximately 10% difference) under partial competition and at three locations per station, operator 2 can even achieve a higher gross profit than operator 1. This is mostly because of the second operator no longer omitting high payoff stations and correlates with a higher number of iterations of the IBR.

5.7. Impact of Other Customer Choice Behaviors

To establish how much gross profit can be gained if customers do not strictly choose vehicles at random and with equal probability, we generate instances with varying competition, margins, and densities for

Table 7. Average Percentage Profit Increase Toward Baseline (No Rebalancing) with Diminishing Returns for an Increasing Number of Locations per Station

		One location per district						Two locations per district						Three locations per district					
		Operator 1			Operator 2			Operator 1			Operator 2			Operator 1			Operator 2		
Setting	Baseline	IBR-0	QMO	QMP	IBR-0	QMO	QMP	IBR-0	QMO	QMP	IBR-0	QMO	QMP	IBR-0	QMO	QMP	IBR-0	QMO	QMP
F_H_0.9	12	161	2	74	102	2	88	134	14	84	115	22	92	113	14	63	103	24	85
F_H_0.7	12	161	5	81	101	3	87	135	31	85	113	31	94	128	48	96	118	56	82
F_H_0.5	12	158	8	86	102	5	78	147	56	110	99	40	85	143	78	105	124	72	98
F_L_0.9	6	36	−11	9	17	−14	9	31	10	10	22	8	12	23	11	10	28	19	13
F_L_0.7	6	39	−9	9	16	−16	9	25	12	10	23	14	10	37	29	12	31	24	12
F_L_0.5	6	33	−19	8	22	−15	11	33	3	11	30	8	13	59	27	18	40	14	15
P_H_0.9	12	249	202	232	237	203	232	192	172	185	195	179	191	141	117	132	168	148	160
P_H_0.7	12	244	201	227	237	204	234	208	187	202	200	184	196	181	157	174	174	155	168
P_H_0.5	12	243	197	227	234	202	230	223	206	215	214	201	214	220	203	215	215	200	209
P_L_0.9	6	36	27	27	29	25	24	33	29	29	29	27	27	25	24	22	33	32	28
P_L_0.7	6	35	22	25	36	28	30	33	30	31	32	30	28	38	35	33	38	36	33
P_L_0.5	6	35	26	26	31	24	24	42	39	37	39	37	36	54	44	46	61	52	55

either operator and address them in an analogy to Section 5.2. If both operators are available at a competitive location i, a customer chooses operator 1 with probability $\alpha \in \{0.5, 0.75, 1\}$, where $\alpha = 0.5$ is strict availability-based substitution. Thirty-six different instances are repeated 100 times, and pure strategy Nash equilibria are found in all cases even though they do not provably exist. In Table B.4, instances are addressed by substitution (in the column header), margin (same for both), density operator 1, density operator 2, and α. For example, $H\_L\_H\_0.75$ in the column F refers to the instance with full substitution, high margins, small operator 1, large operator 2, and customers preferring operator 1 over 2 (selecting it in three out of four cases).

Table 8 gives high-level insights into the trends for other customer choice behaviors. Unsurprisingly, the higher the preference for operator 1, the more the gross profits of the two operators diverge. If customers have a strict preference for operator 1 ($\alpha = 1$), IBR-0 and QMO coincide for operator 1, and operator 2 always collects at least as high of payoffs in the equilibrium as in any of the quasi-monopolistic solutions (operator 2 might be able to collect additional revenue at locations that "do not fit" into operator 1's tour). For operator 1 (the "preferred" operator), the benefit of considering competition, thus, decreases if α increases, and for operator 2, the benefit of considering competition increases.

5.8. Case Study for Munich Car-Sharing

To quantify profit gains and losses, we consider a Munich, Germany, case study. We use publicly available data from two Munich car-sharing providers containing start and end locations and times for car-sharing trips collected in August 2019. Because the data set does not contain any data about the customers, we assume that all customers have both memberships and have no preference for one operator over the other. Having both memberships is realistic for frequent users and, thus, most trips. The car-sharing operators have large fleets of ≈ 500 and ≈ 700 vehicles, respectively. We aggregate trips by assigning them to a start and end district. Districts are hexagons with a radius of $\approx$ 500 m, which is commonly assumed to be a reasonable walking distance and provides sufficient flexibility to operators (Ströhle, Flath, and Gärttner 2019). Districts are approximated by stations at the center of the district. We focus on 21 stations with the highest demand during the observation period (16 days, Monday–Thursday, during August 2019). The average demand of operator 1 is 185 trips, and the average demand of operator 2 is 153 trips. First, we count the number of trips starting and ending in every station. The differences between arrivals and departures ("demand imbalance") can be described by a cumulative arrival probability $P(\hat{\imath} \geq k)$, that is, the probability that the kth vehicle moved to station ι is used. The probability $P(\hat{\imath} \geq k)$ is derived from the available data. During rebalancing, external influences on demand and supply are sufficiently little, and all remaining differences in demand can be attributed to randomness. Probability $P(\hat{\imath} \geq k)$ is independent of n because all customers are shared and have no preference for one operator over the other. We define a (joint) payoff function

$$\pi_\iota^n(q_1^\iota, q_2^\iota) = \frac{q_n^\iota}{q_1^\iota + q_2^\iota} \sum_{k=1}^{q_1^\iota + q_2^\iota} P(\hat{\imath} \geq k) \cdot \pi,$$

where π is the contribution margin associated with serving additional customers as a result of rebalancing a vehicle. We set $\pi = 15$ to account not only for direct revenues of the first customer, but also all future users of that car until it must be rebalanced again as well as the benefit of preventing customer dissatisfaction (because of a low level of service). We chose this value because the data suggests that vehicles are rebalanced after approximately 10 trips, customers pay approximately 0.35€ per minute (minus direct costs), and trips often take at least 15 minutes.

The rebalancing costs are calculated by the travel time between the stations (given a velocity of 20 km/h) and five minutes for additional tasks (e.g., loading/unloading the foldable bike, searching for a parking spot) at an hourly wage of 10€/h and vehicle cost of 0.3€/km. Thus, the minimum rebalancing cost is $\approx$ 3.4€ (moving back and forth between two stations). We use this minimum rebalancing cost as a bound on the maximum number of vehicles that can be profitably moved to a location. Then, there are 28 delivery locations in nine stations (with one to five locations per station). Of the 26 pickup locations distributed across 10 stations, 13 belong to operator 1 and 13 belong to operator 2. The remaining two stations inherently have balanced demand.

Table 9 reports the gross profits in euros of either operator when using IBR, QMO, QMP, and M-PDOP to find routes.

Table 8. Average Percentage Profit Increase Toward Baseline (No Rebalancing) Considering Different Customer Choice

	Operator 1			Operator 2		
Setting	IBR-0	QMO	QMP	IBR-0	QMO	QMP
Average ($\alpha = 0.5$)	102	58	33	85	57	36
Minimum ($\alpha = 0.5$)	6	−52	0	5	−53	0
Maximum ($\alpha = 0.5$)	254	254	121	260	260	136
Average ($\alpha = 0.75$)	117	89	105	71	27	1
Minimum ($\alpha = 0.75$)	5	2	0	3	−174	0
Maximum ($\alpha = 0.75$)	262	262	252	255	255	3
Average ($\alpha = 1$)	119	119	119	70	−3	0
Minimum ($\alpha = 1$)	6	6	6	1	−302	0
Maximum ($\alpha = 1$)	258	258	258	257	257	0

Table 9. Results for the Munich Case Study (Absolute Profits in the Left Block, Relative Gaps in the Right Block)

	IBR	QMO	QMP	M-PDOP	$\frac{IBR}{QMO}$	$\frac{IBR}{QMP}$	$\frac{IBR}{M\text{-}PDOP}$
Operator 1	35.61	25.39	31.43		1.40	1.13	
Operator 2	35.28	26.96	31.59		1.31	1.12	
Operators 1 and 2	70.89	52.35	63.02	79.18	1.35	1.12	0.90

Because M-PDOP and Coop-PDOP coincide if all delivery locations are competitive, we do not report this additionally. The gross profit gains resulting from considering competition range between 31% and 40% (35%, on average, over both operators) compared with assuming that the competitor does not move any vehicles to locations with demand for additional vehicles. The high gains stem from the fact that both operators move their vehicles to the same locations, leaving a demand imbalance of eight vehicles even though almost all demand could be served (in equilibrium, no vehicle is moved to three potential customers). However, operators can also reach high profits by assuming their competitor serves all locations of all stations. On average, the operators improve their profits by 12% by considering competition compared with assuming that their competitor is omnipresent. Because of the large fleet and the large stations/districts, the benefit of merging or outsourcing the rebalancing operations is small (≈10%). In conclusion, the price of ignoring competition is very high at 35%, and the price of ignoring the rationality of the competitor is lower at 12% (most likely, operators currently assume some strategy between these two extrema, and the gain from considering competition is most likely closer to 12% than to 35%). The price of competition is not too high with 10%.

5.9. Algorithmic Performance Results

With respect to performance of the algorithms, we now focus on the following aspects: How large are the instances for car-sharing relocation that we can solve, in reasonable time, under the various business models; in particular, can IBR and PFO solve the C-PDOP on real life-sized instances? How large are the differences in gross profits between Nash equilibria found by the different algorithms?

5.9.1. Size of Solvable Instances.

All rebalancing problems under consideration are NP-hard problems (Theorem 1). However, we can solve medium-sized instances with up to 50 vehicles that shall be rebalanced. Weikl and Bogenberger (2015) record 36 relocations in Munich during one night for one operator. Munich's one-way, free-floating car-sharing fleet is among the largest in the world; thus, the size of solvable instances is most likely sufficient in other cities as well. Further, operators are now reducing the number of necessary relocations by offering incentives for user-based relocation (e.g., Ströhle, Flath, and Gärttner 2019) and by increasing the fleet size (e.g., George and Xia 2011).

In Figure 4, we show the average runtimes of 10 instances of increasing size on a logarithmic scale for full and partial substitution (the latter with different substitution rates). The instances are solved with the IBR-0 and the PFO.

As the relocation problem in car-sharing is an operational problem repeated every night, it should not run for more than 30 minutes in the average case (or 10 minutes per iteration in the IBR because IBR solves most instances in three to four iterations). This is possible for instances with up to 50 locations with both algorithms if there is only a little substitution. For the IBR, it is still computationally feasible to solve instances with 50 pickup and delivery locations per operator under realistic substitution (25%), but in these instances, we already observe the computational advantage of the IBR-0 over PFO. Under full substitution, both algorithms perform substantially worse, but the IBR-0 can solve instances that are roughly twice as large. Thus, with respect to runtime, the iterated best response is the method of choice with both algorithms

Figure 4. (Color online) Runtime (in Seconds) on a Logarithmic Scale for Full Substitution, Partial Substitution (25%), and Partial Substitution (10%)

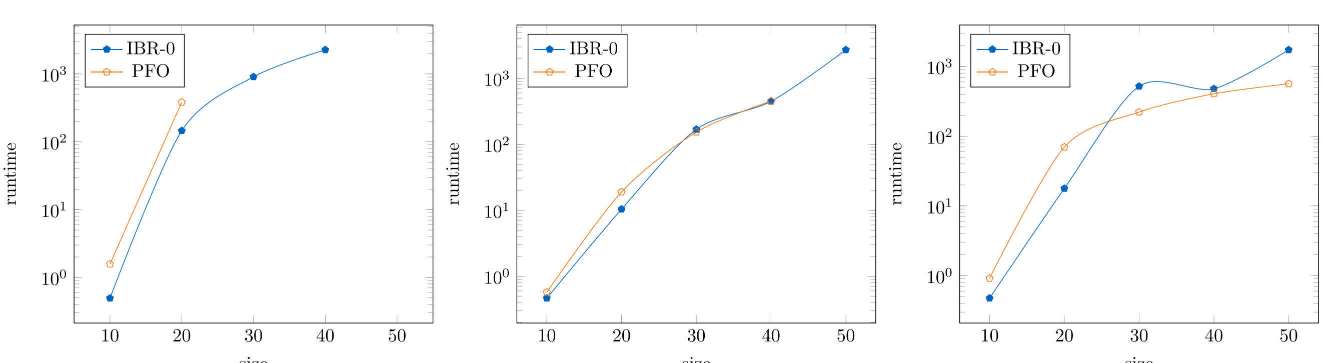

Figure 5. (Color online) Fraction of Instances Solved to Provable/1%/5% Optimality for Full Substitution and Partial Substitution (25%) Using IBR-0 and PFO

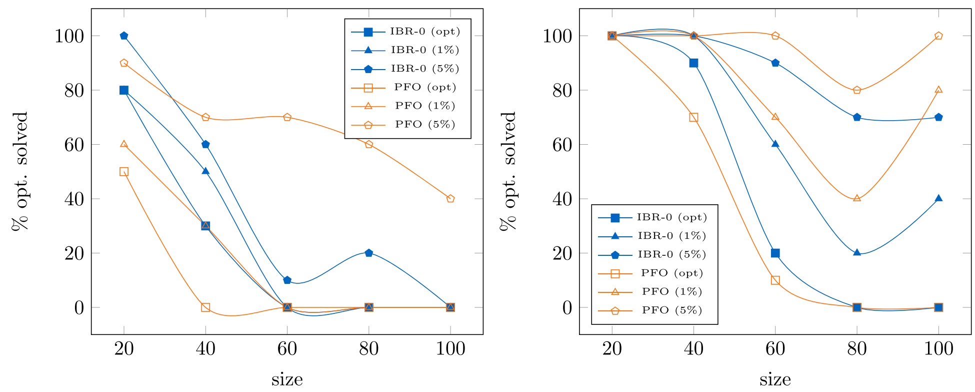

performing roughly the same for low substitution (10%).

Both algorithms only find exact equilibria reliably on comparably small instances. Leveraging Theorem 4, we know that provable optimality of the PDOP is not necessary to reach a "sufficiently stable" solution. Thus, Figure 5 reports the fraction of instances of given size that solve (i) to optimality, (ii) to 1% optimality, or (iii) to 5% optimality using IBR-0 and PFO within a time limit of 10 minutes per iteration (IBR-0) or 30 minutes in total (PFO).

Similar to the results from Figure 4, we observe that IBR-0 and PFO solve a similar percentage of small and medium-sized instances. Surprisingly, PFO solves 80% of the instances with 100 customers and vehicles and partial competition to 1% optimality although IBR-0 fails to solve almost any instances with 100 customers and vehicles to 1% optimality (even though the average runtime to optimality is higher for PFO than for IBR). A similar effect can be observed for full competition. The reason for this is that, with IBR, terminating with a high optimality gap in any iteration results in a weak approximation guarantee. This is also observable from Figure 6, which reports the average gap for PFO and the average over the worst gap of all iterations using IBR-0, which gets as high as 9.7% for full substitution and 60 locations (only if at least 50% of all instances provide a feasible solution). Also, a deeper look into the branch-and-bound behavior for partial competition reveals that already the first found integer solution often provides a reasonably good bound. Thus, if any feasible solution is found, it frequently already provides a reasonably tight approximation guarantee.

5.9.2. Trade-off Because of Equilibrium Selection. Both of the algorithms we presented for the C-PDOP come at a price: Neither of the algorithms provably returns the best Nash equilibrium. In Figure 7, we denote the actual gaps between different Nash equilibria.

If the number of Nash equilibria increases (derived using full enumeration), we empirically observe the following ordering for the welfare (sum over all gross profits): the average Nash equilibrium (derived using FE) results in lower gross profits than IBR-0, which, in turn, has lower gross profits than PFO. The profit of PFO is exceeded by IBR-WP, which has a lower gross profit than the best Nash equilibrium (derived using

Figure 6. (Color online) Average Optimality Gap for Full and Partial Substitution (25%) Using IBR-0 and PFO (if Incumbent Is Found for at Least 50% of Instances)

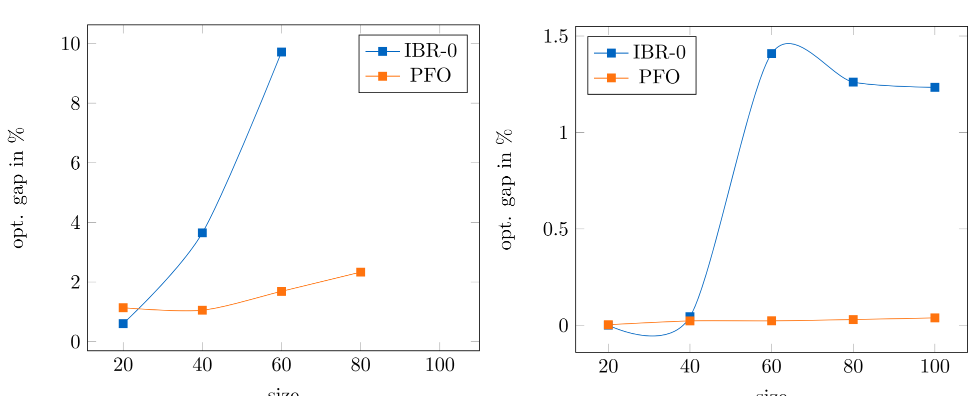

Figure 7. (Color online) Quality of Nash Equilibria (for Players 1 and 2 and with Respect to Welfare, Respectively)

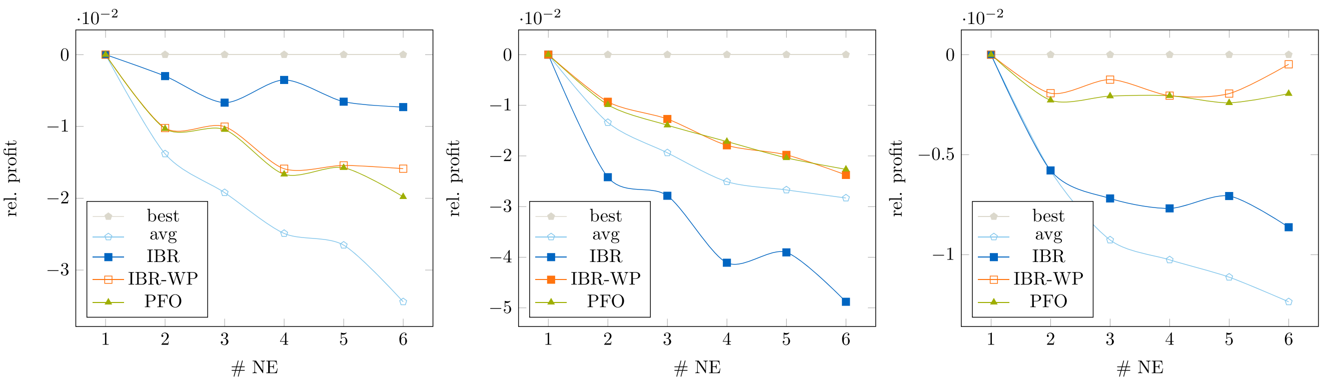

FE). This ordering is inverse to the ordering by runtime. Further, we can see that player 1 (the operator who moves first) profits more from IBR-0 than player 2 (which makes it the second-best and worst algorithm, respectively). Thus, there is a trade-off between solution quality and runtime. If the number of expected Nash equilibria is low, if computation time is a scarce resource, or if any stable solution rather than the best solution is sufficient, IBR-0 is the preferred method, and PFO and IBR-WP (or even FE) are preferred if margins are low and solution quality is paramount.

6. Conclusion

In this paper, we study the profitability of relocation operations under competition in a station-based car-sharing system. We present a mathematical model for the relocation problem that arises in car-sharing, which we PDOP. We further present variations of the PDOP that capture different business models under competition: The C-PDOP models direct competition, and the M-PDOP and Coop-PDOP model merger/monopoly and coopetition/outsourcing relocation operations, respectively. In the C-PDOP, we introduce a competitor who also optimizes the fleet and then solve the problem for Nash equilibria, that is, stable states of the system. The C-PDOP assumes that each operator plans the tour before executing any relocations. In a future line of research, one could investigate a dynamic setting or multistage game in which operators can change their decision during relocation as they observe the competitors' moves. We present two algorithms to find pure strategy Nash equilibria, namely IBR and PFO, both of which are considerably faster than out-of-the-box and brute-force algorithms. Pure strategy Nash equilibria provably exist if both operators receive the same revenue from servicing a customer (player homogeneous payoffs), stations hold at most one vehicle or customer and demand processes are independent (unit demand stations), and if customers choose a vehicle at random if multiple operators have a vehicle available (indifferent customer choice). Even if these assumptions do not hold, we find that equilibria do exist and are reached by IBR in most settings.

Though there are examples in which considering competition is worse than ignoring the presence of competition for some operators, we show numerically that profitability for all operators increases in realistic settings. If margins are low, this improvement (or, vice versa, the *cost of ignorance*) can be up to several orders of magnitude. The main drivers for a high cost of ignorance or benefit of considering competition are fierce competition, a high number of operators, inhomogeneous payoffs, not too large stations, and customer preferences for one operator. In a case study, the gross profit improvement resulting from considering competition is 35% over assuming that no competition exists and 12% over assuming that competition is omnipresent. The more candidate locations there are, the more important relocation becomes as routes become more efficient. The more of these locations are shared, the more important it becomes to consider competition. We observe that operators might be worse off by ignoring the presence of competition in their routing decision than not relocating any vehicles and might even lose money, in particular, if competition is fierce. For each of the three assumptions (player homogeneous payoffs, unit demand stations, and indifferent customer choice), we show that lifting the assumption results in similar tendencies if equilibria exist. Equilibria exist in many realistic instances. Though the studied algorithms do not necessarily find the best equilibrium, we show numerically that they still yield higher gross profits than solutions that do not consider competition. Hence, operators have an incentive to adopt and implement game-theoretic strategies in their relocation decisions.

Acknowledgments

The authors are grateful to two anonymous reviewers for their comments that led to a much improved version of this paper. The authors thank Yiannis Giannakopoulos for his numerous helpful comments.

Appendix A. Proofs

A.1. Proof of Theorem 1 (NP-Hardness)

We prove hardness even for the restricted C-PDOP model, which implies hardness for the general model as well. To prove point 1, we reduce from the TSP, which is well known to be NP-hard (Karp 1972). The NP-hardness of the TSP on bipartite instances follows from Krishnamoorthy (1975), whereas the NP-hardness of the prize-collecting TSP (on general graphs) follows from Feillet, Dejax, and Gendreau (2005). Because our setting combines both modifications, we provide for completeness an NP-hardness proof, adapting the reduction techniques in those papers as well as in Volgenant and Jonker (1987).

Given an instance of the TSP comprising a complete graph K_n with vertices labeled $1,\ldots,n$ and arc routing costs c_e for $e \in E_n$, construct a PDOP instance G with $4n+2$ locations, one worker, and $S^0 = S^1 = \emptyset$ (all locations are optional) as follows. For each node $i \in K_n$ we include four copies i^0, i^1, i^2, i^3 in G such that $i^0, i^2 \in \mathcal{Z}^-$ and $i^1, i^3 \in \mathcal{Z}^+$ (if the costs are known to satisfy the triangular inequality, for example, in the Euclidean TSP, then the TSP remains NP-hard, and a somewhat simpler reduction can be used in which only two copies i, i' of each node are required). We further include additional depot nodes d and p. Fix a node $1 \in K_n$; then, arc $\langle dj^3\rangle$ has cost $c_{\langle 1j\rangle}$ for each $j \neq 1$; further, arc $\langle 1^0 p\rangle$ has cost zero. For each arc $\langle ij\rangle \in A$, the corresponding arc $\langle i^0 j^3\rangle$ has cost $c_{\langle ij\rangle}$. For each node $i \in K_n$, the arcs $\langle i^0 i^1\rangle$, $\langle i^2 i^1\rangle$, and $\langle i^2 i^3\rangle$ have cost zero. Any other arc has cost Cn, where $C = \max_{e\in E_n} c_e$. As far as profits go, for each node $i \in K_n$, the corresponding node i^2 has profit Cn. Any other node has profit zero. Finally, set all travel times to be zero and an arbitrary positive time window T so that we can drop the restrictions on the travel time of the worker.

To finish the proof of the reduction, simply observe that a tour $(i_1 i_2 \ldots i_n i_1)$ in K_n having cost L with $i_1 = 1$ can be lifted into a d–p path $(d[i_2]\ldots[i_n][i_1]p)$ in G, where $[i]$ denotes the sequence $i^3 i^2 i^1 i^0$. This path has profit $Cn^2 - L \in [Cn(n-1), Cn^2]$. Moreover, any d–p path not of this form has profit at most $Cn(n-1)$. Thus, K_n admits a tour of cost at most L if and only if G admits a d–p path of profit at least $Cn^2 - L$.

To prove point 2, we observe that a PDOP instance (for one player) is a special case of a C-PDOP instance (with two players) in which the other player routes on a trivial graph and $\mathcal{Z}^C = \emptyset$. To prove points 3 and 4, we observe that a PDOP instance is a special case of a C-PDOP instance in which all locations are competitive but the distances from depot nodes d_2, p_2 to the rest of the graph are prohibitively large (so that player 2's best strategy is to play the empty strategy $\langle 0,\ldots,0\rangle$ and the Nash equilibria correspond to player 1's best responses to this strategy). □

A.2. Proof of Lemma 1 (Congestion Game)

Given a restricted C-PDOP instance with c locations (and n players), we construct a congestion game with at most $c + n2^c$ resources (and also with n players). The congestion game is described by the following components: a set of common resources R; a payoff function p_r for each resource $r \in R$; and a set of valid strategies ξ^n for each player, where each strategy in ξ^n is a subset of R.

Recall that, from the discussion in Section 3.2, we can specify the payoff function at each location i by a single value π^i that represents the payoff that any player could extract by being the sole operator at i. For each location $i \in \mathcal{Z}$, we include a resource, also denoted $i \in R$, with payoff function

$$p_i(y^i) = \frac{\pi^i}{y^i},$$

where y^i is the number of players having a vehicle available at location i. Moreover, for each player n and each (compact) strategy $\mathbf{q}$ of that player in the original C-PDOP instance, our congestion game includes a resource, which is denoted $(\mathbf{q}, n) \in R$, with a constant, negative payoff function

$$p_{(\mathbf{q},n)} = -C_n(\mathbf{q});$$

these can be, in theory, obtained by solving the PDOP instance described in Equations (1a)–(1k) while setting fixed the variables corresponding to locations.

A valid strategy of player n consists of exactly one resource of the second type, $(\mathbf{q}, n)$, and all associated locations i such that $q^i = 1$; that is, the set ξ_n of valid strategies is given by

$$\xi_n = \Big\{\{(\mathbf{q},n)\} \cup \{i : q^i = 1\} : \mathbf{q} \text{ is a compact strategy} \times \text{of player } n\Big\}.$$

In this way, we obtain a valid formulation of a congestion game as laid out by Rosenthal (1973). The profit of a player n playing strategy $x_n = \{(\mathbf{q}_n, n)\} \cup \{i : q_n^i = 1\}$ is defined as $\sum_{r\in\{(\mathbf{q}_n,n)\}\cup\{i:q_n^i=1\}} p_r(y^r)$, where y^r is the number of players accessing resource r. There are, however, two key differences from the usual formulation: players are maximizing payoffs instead of minimizing costs, and payoffs may assume positive or negative values. These differences are without loss of generality because the standard potential argument can still be applied as we show in the proof of Corollary 1.

The (compact) strategies $\mathbf{q}$ in the standard C-PDOP instance are in one-to-one correspondence to the valid strategies $\mathbf{x}$ in the congestion game, in which $x_n = \{(\mathbf{q}_n, n)\} \cup \{i : q_n^i = 1\}$. This correspondence preserves profits; if P_n is the profit function of player n in the congestion game, then

$$\begin{aligned} P_n(\mathbf{x}) &= \sum_{r\in\{(\mathbf{q}_n,n)\}\cup\{i:q_n^i=1\}} p_r(y^r) = \sum_{i:q_n^i=1} \frac{\pi^i}{y^i} + \pi^{(\mathbf{q}_n,n)} \\ &= \sum_{i:q_n^i=1} \pi^i(\mathbf{q}^i) - C_n(\mathbf{q}_n) = R_n(\mathbf{q}) - C_n(\mathbf{q}) = \Pi_n(\mathbf{q}). \end{aligned}$$

Thus, a deviation is improving in the C-PDOP instance if and only if it is improving in the congestion game. It follows that $\mathbf{q}$ is a Nash equilibrium for the C-PDOP instance if and only if the corresponding strategy $\mathbf{x}$ is a Nash equilibrium for the congestion game. □

A.3. Proof of Corollary 1 (Existence of Pure Strategy Nash Equilibria)

We apply Rosenthal's (1973) potential argument. For a given strategy profile $\mathbf{q}$, let $y^i = y^i(\mathbf{q})$ denote the number of operators placing a vehicle at location i. Define the potential function

$$\Phi(\mathbf{s}) = \sum_{i\in\mathcal{Z}}\sum_{q=1}^{y^i} p_i(q) - \sum_n C_n(\mathbf{q}_n) = \sum_{i\in\mathcal{Z}} H_{y^i}\pi^i - \sum_n C_n(\mathbf{q}_n),$$

where $H_{y^i} = 1 + \frac{1}{2} + \cdots + \frac{1}{y^i}$ denotes the harmonic number of order y^i.

Next observe that the potential function keeps track of the changes in profit when a player deviates. For example, suppose player n deviates from strategy $\mathbf{q}_n$ to $\tilde{\mathbf{q}}_n$ while the other players keep to strategy $\mathbf{q}_{-n}$, and let $\tilde{y}^i$ denote the number of vehicles at location i in the strategy profile $(\tilde{\mathbf{q}}_n, \mathbf{q}_{-n})$. We prove that the change in potential equals the change in the profit of player n:

$$\begin{aligned}
&\Phi(\tilde{\mathbf{q}}_n, \mathbf{q}_{-n}) - \Phi(\mathbf{q}_n, \mathbf{q}_{-n})\\
&= \sum_{i\in\mathcal{Z}}\sum_{q=1}^{\tilde{y}^i} p_i(q) - C_n(\tilde{\mathbf{q}}_n) - \sum_{n'\neq n} C_{n'}(\mathbf{q}_{n'}) - \left(\sum_{i\in\mathcal{Z}}\sum_{q=1}^{y^i} p_i(q) - C_n(\mathbf{q}_n) - \sum_{n'\neq n} C_{n'}(\mathbf{q}_{n'})\right)\\
&= \sum_{i:q_n^i=0,\tilde{q}_n^i=0}\left(\sum_{q=1}^{\tilde{y}^i} p_i(q) - \sum_{q=1}^{y^i} p_i(q)\right) + \sum_{i:q_n^i=1,\tilde{q}_n^i=1}\left(\sum_{q=1}^{\tilde{y}^i} p_i(q) - \sum_{q=1}^{y^i} p_i(q)\right)\\
&\quad + \sum_{i:q_n^i=0,\tilde{q}_n^i=1}\left(\sum_{q=1}^{\tilde{y}^i} p_i(q) - \sum_{q=1}^{y^i} p_i(q)\right) + \sum_{i:q_n^i=1,\tilde{q}_n^i=0}\left(\sum_{q=1}^{\tilde{y}^i} p_i(q) - \sum_{q=1}^{y^i} p_i(q)\right)\\
&\quad - C_n(\tilde{\mathbf{q}}_n) + C_n(\mathbf{q}_n)\\
&= \sum_{i:q_n^i=1,\tilde{q}_n^i=1} p_i(\tilde{y}^i) - \sum_{i:q_n^i=1,\tilde{q}_n^i=1} p_i(y^i) + \sum_{i:q_n^i=0,\tilde{q}_n^i=1} p_i(\tilde{y}^i)\\
&\quad - \sum_{i:q_n^i=1,\tilde{q}_n^i=0} p_i(y^i) - C_n(\tilde{\mathbf{q}}_n) + C_n(\mathbf{q}_n)\\
&= \sum_{i:\tilde{q}_n^i=1} p_i(\tilde{y}^i) - C_n(\tilde{\mathbf{q}}_n) - \left(\sum_{i:q_n^i=1} p_i(y^i) - C_n(\mathbf{q}_n)\right)\\
&= \Pi_n(\tilde{\mathbf{q}}_n, \mathbf{q}_{-n}) - \Pi_n(\mathbf{q}_n, \mathbf{q}_{-n}).
\end{aligned}$$

In the second equality, we split the sums over $i \in \mathcal{Z}$ into four cases, depending on whether each of q_n^i, $\tilde{q}_n^i$ is zero or one. In the third equality, we observe

- If $q_n^i = 0, \tilde{q}_n^i = 0$, then $\tilde{y}^i = y^i$ and $\sum_{q=1}^{\tilde{y}^i} p_i(q) - \sum_{q=1}^{y^i} p_i(q) = 0$.
- If $q_n^i = 1, \tilde{q}_n^i = 1$, then $\tilde{y}^i = y^i$ and $\sum_{q=1}^{\tilde{y}^i} p_i(q) - \sum_{q=1}^{y^i} p_i(q) = 0 = p_i(\tilde{y}^i) - p_i(y^i)$.
- If $q_n^i = 0, \tilde{q}_n^i = 1$, then $\tilde{y}^i = y^i + 1$ and $\sum_{q=1}^{\tilde{y}^i} p_i(q) - \sum_{q=1}^{y^i} p_i(q) = p_i(\tilde{y}^i)$.
- If $q_n^i = 1, \tilde{q}_n^i = 0$, then $\tilde{y}^i = y^i - 1$ and $\sum_{q=1}^{\tilde{y}^i} p_i(q) - \sum_{q=1}^{y^i} p_i(q) = -p_i(y^i)$.

We conclude that a deviation by any player is improving if and only if the potential increases. Because there are finitely many strategies, Φ possesses a global maximum. The corresponding strategy must be a Nash equilibrium because no player's deviation would increase the value of the potential function and, thus, not increase the player's profit. Moreover, if $\mathbf{q}^{(0)} \to \mathbf{q}^{(1)} \to \ldots$ is a sequence in which $\mathbf{q}^{(n+1)}$ is obtained from $\mathbf{q}^{(n)}$ by an improving deviation from one of the players, then the potential must strictly increase through the sequence. Such a sequence must then terminate at a local maximum, which is a Nash equilibrium. □

A.4. Proof of Corollary 2 (Monotonicity of Profits (vs. Welfare-Maximizing Solution))

By definition, the welfare-maximizing solution is the solution to the PDOP with two competing operators in which the joint profit is maximal. Thus, no other solution, including any Nash equilibrium solution, can be better. □

A.5. Proof of Lemma 2 (Price of Anarchy and Price of Stability)

Consider an instance of the restricted C-PDOP model with two players. We first derive a useful relation between the profits of a player when the other player changes strategy. Let $\mathbf{q}_1$ be any strategy for player 1 and $\mathbf{q}_2, \mathbf{q}_2'$ be any two strategies for player 2. By definition of the profit function, we have

$$\Pi_1(\mathbf{q}_1, \mathbf{q}_2) = \sum_{i:q_1^i=1,q_2^i=0} \pi^i + \sum_{i:q_1^i=1,q_2^i=1} \frac{\pi^i}{2} - C_1(\mathbf{q}_1);$$
$$\Pi_1(\mathbf{q}_1, \mathbf{q}_2') = \sum_{i:q_1^i=1,q_2'^i=0} \pi^i + \sum_{i:q_1^i=1,q_2'^i=1} \frac{\pi^i}{2} - C_1(\mathbf{q}_1);$$

putting these two equations together, we see that

$$\Pi_1(\mathbf{q}_1, \mathbf{q}_2) - \Pi_1(\mathbf{q}_1, \mathbf{q}_2') = \frac{1}{2}\sum_{i:q_1^i=1,q_2'^i=1,q_2^i=0} \pi_i - \frac{1}{2}\sum_{i:q_1^i=1,q_2'^i=0,q_2^i=1} \pi_i. \tag{A.1}$$

To prove point 1, let $(\mathbf{q}_1, \mathbf{q}_2)$ be any Nash equilibrium and $(\mathbf{q}_1', \mathbf{q}_2')$ be any strategy. Applying Equation (A.1) and the definition of Nash equilibrium,

$$\begin{aligned}
\Pi_1(\mathbf{q}_1', \mathbf{q}_2') &= \Pi_1(\mathbf{q}_1', \mathbf{q}_2) + \frac{1}{2}\sum_{i:q_1'^i=1,q_2^i=1,q_2'^i=0} \pi_i - \frac{1}{2}\sum_{i:q_1'^i=1,q_2^i=0,q_2'^i=1} \pi_i\\
&\le \Pi_1(\mathbf{q}_1, \mathbf{q}_2) + \frac{1}{2}\sum_{i:q_1'^i=1,q_2^i=1,q_2'^i=0} \pi_i - \frac{1}{2}\sum_{i:q_1'^i=1,q_2^i=0,q_2'^i=1} \pi_i
\end{aligned}.$$

Similarly, for player 2, we have

$$\Pi_2(\mathbf{q}_1', \mathbf{q}_2') \le \Pi_2(\mathbf{q}_1, \mathbf{q}_2) + \frac{1}{2}\sum_{i:q_2'^i=1,q_1^i=1,q_1'^i=0} \pi_i - \frac{1}{2}\sum_{i:q_2'^i=1,q_1^i=0,q_1'^i=1} \pi_i;$$

adding both equations and observing that the two sets $\{i : q_1'^i = 1, q_2^i = 1, q_2'^i = 0\}$ and $\{i : q_2'^i = 1, q_1^i = 1, q_1'^i = 0\}$ are disjoint subsets of $\mathcal{Z}^C$, as well as that $\pi_i \ge 0$, we obtain

$$\begin{aligned}
\Pi(\mathbf{q}_1', \mathbf{q}_2') &= \Pi_1(\mathbf{q}_1', \mathbf{q}_2') + \Pi_2(\mathbf{q}_1', \mathbf{q}_2') \le \Pi_1(\mathbf{q}_1, \mathbf{q}_2)\\
&\quad + \Pi_2(\mathbf{q}_1, \mathbf{q}_2) + \frac{1}{2}\sum_{i\in\mathcal{Z}^C} \pi_i = \Pi(\mathbf{q}_1, \mathbf{q}_2) + \frac{1}{2}\sum_{i\in\mathcal{Z}^C} \pi_i.
\end{aligned}$$

Because $(\mathbf{q}_1', \mathbf{q}_2')$ was taken to be any strategy, the preceding equation implies that the difference in welfare between any two Nash equilibria is at most $\frac{1}{2}\sum_{i\in\mathcal{Z}^C} \pi_i$, which proves point 2. Now, let $(\mathbf{q}_1, \mathbf{q}_2)$ be any Nash equilibrium and $(\mathbf{q}_1^*, \mathbf{q}_2^*)$ be a welfare-maximizing strategy. We get that

$$\Pi(\mathbf{q}_1, \mathbf{q}_2) \le \Pi(\mathbf{q}_1^*, \mathbf{q}_2^*) \le \Pi(\mathbf{q}_1, \mathbf{q}_2) + \frac{1}{2}\sum_{i\in\mathcal{Z}^C} \pi_i,$$

and thus, the difference between any Nash equilibrium and any welfare-maximizing strategy is at most $\frac{1}{2}\sum_{i\in\mathcal{Z}} \pi_i$, proving point 3.

Figure A.1. Example of a C-PDOP Instance (Left)

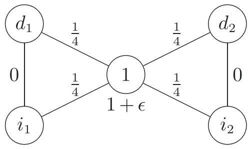

	0	1
0	$0,0$	$0,\frac{1}{2}+\epsilon$
1	$\frac{1}{2}+\epsilon,0$	$\frac{\epsilon}{2},\frac{\epsilon}{2}$

	0	1
0	$0,0$	$0,\frac{1}{2}$
1	$\frac{1}{2},0$	$0,0$

Notes. There is only one delivery location, labeled 1, which is competitive. Each player has a depot (d_1 or d_2) and one vehicle at a separate location (i_1 or i_2) with null payoff. The corresponding payoff matrix (center) and the special case $\epsilon = 0$ (right).

To prove that the price of anarchy and the price of stability can be arbitrarily high, consider the game depicted in Figure A.1. There is only one competitive location with payoff $\pi = 1+\epsilon$ for both players, for some small positive ϵ. Either player incurs a traveling cost of $c = 1/2$ to relocate a vehicle to the competitive location and return to the depot. Thus, each player has only two strategies (one for "move" or zero for "don't move"). Looking at the payoff matrix, we see that the only Nash equilibrium occurs when both players decide to service the location for a welfare of ϵ, whereas the maximum possible welfare is $\frac{1}{2}+\epsilon$, occurring when only one player services the location. Thus, both the price of anarchy and the price of stability equal $\frac{1/2+\epsilon}{\epsilon} = 1 + \frac{1}{2\epsilon}$, which is arbitrarily large as ϵ can be arbitrarily small. Because the C-PDOP generalizes the restricted C-PDOP, the result also carries over to the C-PDOP.

To show that the bounds obtained in points 1–3 are tight, consider the same game as before, but take $\epsilon = 0$. Now, both strategies $(q_1, q_2) = (1, 0)$ or $(0, 1)$ are welfare maximizing as well as Nash equilibria, achieving welfare $\frac{1}{2} = \frac{1}{2}\sum_{i\in\mathcal{Z}^C}\pi_i$. Moreover, $(q_1, q_2) = (1, 1)$ is still a Nash equilibrium with welfare exactly equal to zero. Thus, the absolute difference between the welfare-maximizing strategy and the worst Nash equilibrium as well as between best and worst Nash equilibria is exactly $\frac{1}{2}\sum_{i\in\mathcal{Z}^C}\pi_i$. □

A.6. Proof of Corollary 3 (NP-Hardness of M-PDOP and Coop-PDOP)

A PDOP instance is a special case of a Coop-PDOP (M-PDOP) instance (derived from two players) in which the second player routes on a trivial graph and $\mathcal{Z}^C = \emptyset$. Thus, both problems remain NP-hard. □

A.7. Proof of Lemma 3 (Monotonicity of Profits (vs. Monopoly or Coopetition Solution))

To prove points 1 and 2, simply observe that the space of feasible solutions increases as we move from C-PDOP to Coop-PDOP to M-PDOP. In other words, a feasible strategy (x^1, a^1, x^2, a^2) with W_1, W_2 workers in the C-PDOP model can be merged into a feasible strategy $(\bar{x}, \bar{a})$ with $W_1 + W_2$ workers in the Coop-PDOP model, and a feasible strategy (x, a) with W_O workers in the Coop-PDOP model is also a feasible strategy (with the same number of workers) in the M-PDOP model. Moreover, the payoffs associated to competitive locations only increase as we move from C-PDOP to Coop-PDOP to M-PDOP. To see this, let ι be a station and $\mathbf{q}^\iota$ represent the vehicles of each operator at station ι. The joint payoff at location ι is the same for the C-PDOP and Coop-PDOP models and equals $\sum_n \pi_n^\iota(\mathbf{q}^\iota)$. This, in turn, is less than or equal to the joint payoff for the M-PDOP model, which is defined as $\max_{\mathbf{q}'}(\sum_n \pi_n^\iota(\mathbf{q}') \mid \sum_n q_n^\iota = \sum_n q_n^{\iota'})$. Therefore, the optimal profit does not decrease as we move from C-PDOP to Coop-PDOP to M-PDOP (as long as the number of workers is consistent, i.e., $W_M \ge W_O \ge W_1 + W_2$). □

A.8. Proof of Theorem 2 (Monotonicity of Profits (IBR vs. QMO))

We begin by deriving the following relation between the potential and payoff functions, valid for the restricted C-PDOP with two players:

$$\begin{aligned}
\Phi(\mathbf{q}_1, \mathbf{q}_2) &= \sum_{i\in\mathcal{Z}} H_{y^i}\pi^i - C_1(\mathbf{q}_1) - C_2(\mathbf{q}_2)\\
&= \sum_{i:q_1^i=1, q_2^i=0}\pi^i + \sum_{i:q_1^i=0, q_2^i=1}\pi^i + \sum_{i:q_1^i=1, q_2^i=1}\frac{3}{2}\pi^i - C_1(\mathbf{q}_1) - C_2(\mathbf{q}_2)\\
&= \left(\sum_{i:q_1^i=1, q_2^i=0}\pi^i + \sum_{i:q_1^i=1, q_2^i=1}\frac{\pi^i}{2} - C_1(\mathbf{q}_1)\right) + \left(\sum_{i:q_2^i=1}\pi^i - C_2(\mathbf{q}_2)\right)\\
&= \Pi_1(\mathbf{q}_1, \mathbf{q}_2) + \Pi_2(\langle 0, \ldots, 0\rangle, \mathbf{q}_2);
\end{aligned} \tag{A.2}$$

similarly, one has $\Phi(\mathbf{q}_1, \mathbf{q}_2) = \Pi_1(\mathbf{q}_1, \langle 0, \ldots, 0\rangle) + \Pi_2(\mathbf{q}_1, \mathbf{q}_2)$.

Consider a sequence of iterated best responses starting from the $\langle 0, \ldots, 0\rangle$ strategy and ending at a Nash equilibrium with player 1 moving first into an optimistic strategy:

$$(\langle 0,\ldots,0\rangle, \langle 0,\ldots,0\rangle) \xrightarrow{1} \left(\mathbf{q}_1^O, \langle 0,\ldots,0\rangle\right) \xrightarrow{2} \left(\mathbf{q}_1^O, \mathbf{q}_2^{(1)}\right) \to \cdots \to \left(\mathbf{q}_1^N, \mathbf{q}_2^N\right).$$

Because $\mathbf{q}_2^{(1)}$ is a best response to player 1 playing $\mathbf{q}_1^O$, the following is also a sequence of iterated best responses:

$$\left(\mathbf{q}_1^O, \mathbf{q}_2^O\right) \xrightarrow{2} \left(\mathbf{q}_1^O, \mathbf{q}_2^{(1)}\right) \to \cdots \to \left(\mathbf{q}_1^N, \mathbf{q}_2^N\right).$$

In particular, this implies that the potential value at the Nash equilibrium retrieved through IBR is at least the potential value at the optimistic quasi-monopolistic strategy, that is,

$$\Phi\left(\mathbf{q}_1^N, \mathbf{q}_2^N\right) \ge \Phi\left(\mathbf{q}_1^O, \mathbf{q}_2^O\right).$$

Note that the optimistic strategy maximizes a player's profit with respect to the other player playing the empty strategy; in particular, we have $\Pi_1(\mathbf{q}_1^O, \langle 0,\ldots,0\rangle) \ge \Pi_1(\mathbf{q}_1^N, \langle 0,\ldots,0\rangle)$ and $\Pi_2(\langle 0,\ldots,0\rangle, \mathbf{q}_2^O) \ge \Pi_2(\langle 0,\ldots,0\rangle, \mathbf{q}_2^N)$. Using (A.2), we get

$$\begin{aligned}
\Pi_1\left(\mathbf{q}_1^N, \mathbf{q}_2^N\right) &= \Phi\left(\mathbf{q}_1^N, \mathbf{q}_2^N\right) - \Pi_2\left(\langle 0,\ldots,0\rangle, \mathbf{q}_2^N\right)\\
&\ge \Phi\left(\mathbf{q}_1^O, \mathbf{q}_2^O\right) - \Pi_2\left(\langle 0,\ldots,0\rangle, \mathbf{q}_2^O\right)\\
&= \Pi_1\left(\mathbf{q}_1^O, \mathbf{q}_2^O\right).
\end{aligned}$$

Using a similar reasoning for player 2, we get that $\Pi_2(\mathbf{q}_1^N, \mathbf{q}_2^N) \ge \Pi_2(\mathbf{q}_1^O, \mathbf{q}_2^O)$. Thus, both players are better off if they agree on a Nash equilibrium obtained by iterated best responses from the empty strategy. □

A.9. Proof of Theorem 3 (Termination of the Iterated Best Response Algorithm)

We first define the inverse $\neg\mathbf{q}$ of a strategy $\mathbf{q}$ as the strategy in which a vehicle is available at precisely the competitive locations in which $\mathbf{q}$ does not have a vehicle available (formally, for every $i \in \mathcal{Z}^C$, we have that $\neg q^i = 1 - q^i$). To prove this theorem, we need the following key auxiliary result: for any strategies $\mathbf{q}_1$ and $\mathbf{q}_2$, if $\mathbf{q}_1$ is a best response to $\mathbf{q}_2$, then $\mathbf{q}_1$ is also a best response to its inverse $\neg\mathbf{q}_1$. To see this, let $\mathbf{q}_1'$ be an arbitrary strategy for player 1; applying Equation (A.1) twice,

$$\begin{aligned}
\Pi_1(\mathbf{q}_1', \neg\mathbf{q}_1) &= \Pi_1(\mathbf{q}_1', \mathbf{q}_2) + \frac{1}{2}\sum_{i:q_1'^i=1, q_2^i=1, \neg q_1^i=0} \pi_i - \frac{1}{2}\sum_{i:q_1'^i=1, \neg q_1^i=1, q_2^i=0} \pi_i \\
&\le \Pi_1(\mathbf{q}_1, \mathbf{q}_2) + \frac{1}{2}\sum_{i:q_1'^i=1, q_2^i=1, q_1^i=1} \pi_i - \frac{1}{2}\sum_{i:q_1'^i=1, q_1^i=0, q_2^i=0} \pi_i \\
&= \Pi_1(\mathbf{q}_1, \neg\mathbf{q}_1) + \frac{1}{2}\sum_{i:q_1^i=1, q_1^i=0, q_2^i=0} \pi_i - \frac{1}{2}\sum_{i:q_1^i=1, q_2^i=1, q_1^i=1} \pi_i \\
&\quad + \frac{1}{2}\sum_{i:q_1'^i=1, q_2^i=1, q_1^i=1} \pi_i - \frac{1}{2}\sum_{i:q_1'^i=1, q_1^i=0, q_2^i=0} \pi_i \\
&= \Pi_1(\mathbf{q}_1, \neg\mathbf{q}_1) - \frac{1}{2}\sum_{i:q_1^i=1, q_2^i=1, q_1'^i=0} \pi_i - \frac{1}{2}\sum_{i:q_1'^i=1, q_1^i=0, q_2^i=0} \pi_i \le \Pi_1(\mathbf{q}_1, \neg\mathbf{q}_1).
\end{aligned}$$

Next, let us consider a sequence of iterated best responses,

$$\left(\mathbf{q}_1^0, \mathbf{q}_2^0\right) \xrightarrow{1} \left(\mathbf{q}_1^1, \mathbf{q}_2^0\right) \xrightarrow{2} \left(\mathbf{q}_1^1, \mathbf{q}_2^1\right) \xrightarrow{1} \dots \xrightarrow{2} \left(\mathbf{q}_1^N, \mathbf{q}_2^N\right) \xrightarrow{1} \left(\mathbf{q}_1^{N+1}, \mathbf{q}_2^N\right),$$

starting with player 1 and ending at a Nash equilibrium after $2N + 1$ iterations. For ease of exposition, we only consider the case in which the sequence ends with a movement of player 1. For a fixed $0 < i \le N$, there are two possibilities:

- If $\mathbf{q}_1^{i+1} = \neg\mathbf{q}_2^i$, then we have reached a Nash equilibrium: because $\mathbf{q}_2^i$ is a best response to $\mathbf{q}_1^i$, it must be a best response to its inverse $\neg\mathbf{q}_2^i = \mathbf{q}_1^{i+1}$; note that this would imply $i = N$.
- If $\mathbf{q}_1^{i+1} \ne \neg\mathbf{q}_2^i$, then player 1 never plays $\neg\mathbf{q}_2^i$ on subsequent iterations: assume otherwise that $\mathbf{q}_1^{j+1} = \neg\mathbf{q}_2^i$ for some $j > i$. Then, $\neg\mathbf{q}_2^i$ would be an (equally) best response to $\mathbf{q}_2^i$ for player 1 so that $(\neg\mathbf{q}_2^i, \mathbf{q}_2^i)$ would be a Nash equilibrium as before. In particular, $\mathbf{q}_2^i$ would be a best response to $\mathbf{q}_1^{j+1} = \neg\mathbf{q}_2^i$ for player 2. Putting all these together, we get a contradiction as

$$\Phi\left(\mathbf{q}_1^{i+1}, \mathbf{q}_2^i\right) = \Phi\left(\neg\mathbf{q}_2^i, \mathbf{q}_2^i\right) \ge \Phi\left(\neg\mathbf{q}_2^i, \mathbf{q}_2^j\right) = \Phi\left(\mathbf{q}_1^{j+1}, \mathbf{q}_l^j\right) > \Phi\left(\mathbf{q}_1^{i+1}, \mathbf{q}_2^i\right).$$

By similar principles, we have that, if $0 \le i < j \le N$, then $\mathbf{q}_2^i \ne \mathbf{q}_2^j$ (i.e., player 2 will not repeat strategies); if $0 < i < j \le N$, then $\mathbf{q}_1^i \ne \mathbf{q}_1^j$ (i.e., player 1 will not repeat strategies except possibly for $\mathbf{q}_1^0$); if $0 < i \le j \le N$, then $\mathbf{q}_2^j \ne \neg\mathbf{q}_1^i$; and if $0 < i < j \le N$, then $\mathbf{q}_1^j \ne \neg\mathbf{q}_2^i$.

In other words, we conclude that the strategies $\mathbf{q}_1^1, \neg\mathbf{q}_2^1, \mathbf{q}_1^2, \neg\mathbf{q}_2^2, \dots, \mathbf{q}_1^N, \neg\mathbf{q}_2^N$ must be all different. As there are only $|S|$ possible strategies, it follows that $2N \le |S|$; in other words, player 2 can play at most $|S|/2$ different strategies, player 1 can play at most $|S|/2 + 1$ different strategies, and the total number of iterations is at most $|S| + 1$ (for a maximum number of $|S| + 2$ recalculations). □

A.10. Proof of Theorem 4 (Iterated Best Response Algorithm for Approximate Nash Equilibria)

We start by presenting a different conversion from C-PDOP to a congestion game that can handle the description of full strategies. We construct a congestion game having a resource for each location and for each arc in the network. To each arc $e \in A$, we associate a constant negative profit $p_e = -c_e$. Similarly to the proof of Lemma 1, for each location $i \in \mathcal{Z}$, we include a resource with profit function

$$p_i(y^i) = \frac{\pi^i}{y^i},$$

where y^i is the number of players having a vehicle available at location i.

The valid strategies for each player correspond to feasible tours that only visit locations/arcs associated with that player, that is, satisfying Constraints (1b)–(1k) from the PDOP model. In other words, for each valid strategy $s_n = (\mathbf{a}^n, \mathbf{x}^n)$ of player n, we associate a corresponding strategy in the congestion game consisting of those arcs $e \in A$ and locations $i \in \mathcal{Z}$ for which $a_e^n, x_i^n = 1$. As in the proof of Lemma 1, this defines a valid congestion game, albeit in a profit-maximizing instead of cost-minimizing formulation, and in which resources may assume positive or negative values. The standard potential argument as in the proof of Corollary 1 can then be applied to conclude that any sequence of improving deviations must eventually reach a Nash equilibrium.

Next, we consider the ϵ-IBR as described in the statement of the theorem. By definition, a player only deviates if the ϵ-approximately optimal routing found is a strict improvement to that player's profit. Therefore, the ϵ-IBR dynamics still yield a sequence of improving deviations and must terminate after a finite number of iterations. All is left is to prove the quality guarantee of the final strategy, that is, that $\mathbf{s}$ is an ϵ-Nash equilibrium. Let s_n' be the ϵ-approximate best response to $\mathbf{s}_{-n}$ found by player n. Because player n opts to not deviate from s_n, it follows that

$$(1+\epsilon)\Pi_n(s_n, \mathbf{s}_{-n}) \ge (1+\epsilon)\Pi_n(s_n', \mathbf{s}_{-n}) \ge \max_{s_n'} \Pi_n(s_n', \mathbf{s}_{-n});$$

as this relation holds for every player n, the final strategy is an ϵ-Nash equilibrium. □

A.11. Proof of Lemma 4 (Optimality of the Potential Function Optimizer)

Point 1 follows directly from Lemma 1, as the maxima of the potential function correspond directly to Nash equilibria of the congestion game. To prove point 2, let $\mathbf{q}$ be a maximum potential Nash equilibrium in which each location is visited by at most one player ($\sum_n q_n^i \le 1 \ \forall i \in \mathcal{Z}$). Note that this implies, for each location i, that

$$H_{y^i}\pi^i = y^i \pi^i = \sum_n \pi^i(\mathbf{s}) q_n^i,$$

and as such $\Phi(\mathbf{q}) = \Pi(\mathbf{q})$. Now, if $\tilde{\mathbf{q}}$ is any other solution, we have

$$\begin{aligned}
\Pi(\tilde{\mathbf{q}}) &= \sum_{i \in \mathcal{Z}} \sum_n \pi^i(\tilde{\mathbf{q}}) \tilde{\mathbf{q}}_n^i - \sum_n C_n(\tilde{\mathbf{q}}_n) \\
&\le \sum_{i \in \mathcal{Z}} H_{\tilde{y}^i} \pi^i - \sum_n C_n(\tilde{\mathbf{q}}_n) \\
&= \Phi(\tilde{\mathbf{q}}) \le \Phi(\mathbf{q}) = \Pi(\mathbf{q}).
\end{aligned}$$

We conclude that $\mathbf{q}$ is also a welfare-maximizing strategy, and thus, there cannot be a better Nash equilibrium. □

Appendix B. Tables

In the following, we give detailed results for Sections 5.2–5.7. Table B.1 shows the percentage increase toward a

Table B.1. Percentage Profit Increase Toward Baseline (Operator 1 in Left Block, Operator 2 in Right Block) Under Various Models and Algorithms as Well as Different Experimental Settings with Substitution Rates, Margins, and Densities for Either Operator

	Operator 1						Operator 2					
Setting	Baseline	IBR-0	IBR-WP	PFO	QMO	QMP	Baseline	IBR-0	IBR-WP	PFO	QMO	QMP
F_H_H_H	12	160	139	130	−5.46	77.8	12	86.2	125	123	−13.1	66.8
F_H_H_L	12	225	204	206	85.9	180	8	38.5	59.6	67	−61.7	10.1
F_H_L_H	8	79.5	63.6	64.9	−63.2	19.6	12	191	207	215	103	199
F_H_L_L	8	136	110	106	−42.9	30.5	8	59	98.9	97.6	−44.9	32.4
F_L_H_H	6	24.9	21.4	21.4	−12.5	1.9	6	11.1	17.2	17.2	−14.9	1.26
F_L_H_L	6	28.5	28.6	28.7	21.1	8.05	4	3.66	3.93	4.62	−4.4	1.15
F_L_L_H	4	9.59	8.03	7.82	−6.67	1.04	6	21.1	22.9	24.7	13.8	4.25
F_L_L_L	4	5.51	4.88	4.88	1.57	1.07	4	7.04	8.9	8.9	5.31	2.49
P_H_H_H	12	238	234	234	190	223	12	231	240	239	199	228
P_H_H_L	12	243	243	240	210	231	8	149	149	159	117	143
P_H_L_H	8	167	167	164	128	152	12	246	246	251	226	243
P_H_L_L	8	171	166	165	129	152	8	168	178	178	137	164
P_L_H_H	6	28.4	27.4	27.4	18.3	21.2	6	30	32.3	32.3	23.1	23.4
P_L_H_L	6	24.5	24.5	24.2	21.3	17.7	4	8.28	8.28	8.92	4.49	7.25
P_L_L_H	4	6.74	6.74	6.46	3.45	4.89	6	27	27	27.5	24.6	20.9
P_L_L_L	4	10.8	10.6	10.6	6.44	6.66	4	9.29	9.61	9.61	6.04	5.54

Table B.2. Percentage Profit Increase Toward Baseline (Full Substitution in Left Block, Partial Substitution in Right Block) for Different Experimental Settings with Substitution Rates, Margins, and Densities for Either Player

	Full substitution (*F_*)					Partial substitution (*P_*)				
Setting	Baseline	NE	W-PDOP	Coop-PDOP	M-PDOP	Baseline	NE	W-PDOP	Coop-PDOP	M-PDOP
H_H_H_H	24	132	156	184	184	24	237	240	304	331
H_H_H_L	20	151	160	198	204	20	207	209	276	305
H_H_L_H	20	155	169	204	209	20	216	218	287	316
H_H_L_L	16	104	113	143	143	16	172	175	249	277
L_L_H_H	12	19.3	19.4	42.4	42.5	12	29.9	29.9	91.5	121
L_L_H_L	10	19.1	19.3	38.1	44.6	10	18.1	18.2	57.8	88.7
L_L_L_H	10	17.9	17.9	36.7	41.8	10	19.1	19.1	65.9	93.8
L_L_L_L	8	6.89	6.89	11.2	11.2	8	10.1	10.1	37.7	59.6

Table B.3. Percentage Profit Increase Toward Baseline with Inhomogeneous Payoffs (Full Substitution in Left Block, Partial Substitution in Right Block) for Different Experimental Settings with Substitution Rates, Margins, and Densities for Either Player

	Full substitution								Partial substitution							
	Operator 1				Operator 2				Operator 1				Operator 2			
Setting	Baseline	IBR-0	QMO	QMP	Baseline	IBR-0	QMO	QMP	Baseline	IBR-0	QMO	QMP	Baseline	IBR-0	QMO	QMP
H_H_H_H	12	142	−3	100	12	118	−1	98	12	434	434	434	12	430	430	430
H_H_H_L	12	279	143	251	8	82	−62	46	12	424	424	424	8	355	355	355
H_H_L_H	8	98	−85	35	12	260	141	258	8	358	358	358	12	433	433	433
H_H_L_L	8	152	−59	84	8	71	−72	52	8	354	354	354	8	349	349	349
H_L_H_H	12	321	55	274	6	1	−163	−37	12	423	423	423	6	257	257	257
H_L_H_L	12	404	236	358	4	−46	−215	−61	12	425	425	425	4	147	147	147
H_L_L_H	8	212	−21	108	6	113	−23	96	8	362	362	362	6	244	244	244
H_L_L_L	8	316	39	194	4	−47	−239	−66	8	358	358	358	4	122	122	122
L_H_H_H	6	20	−170	−38	12	273	70	279	6	295	295	295	12	416	416	416
L_H_H_L	6	145	−18	92	8	140	−15	96	6	292	292	292	8	336	336	336
L_H_L_H	4	−26	−237	−77	12	371	226	365	4	169	169	169	12	427	427	427
L_H_L_L	4	0	−222	−67	8	211	39	154	4	170	170	170	8	336	336	336
L_L_H_H	6	162	−110	35	6	57	−117	8	6	294	294	294	6	254	254	254
L_L_H_L	6	257	63	169	4	−17	−187	−65	6	277	277	277	4	133	133	133
L_L_L_H	4	62	−166	−54	6	174	48	155	4	182	182	182	6	245	245	245
L_L_L_L	4	130	−105	−49	4	8	−135	−56	4	157	157	157	4	141	141	141

Table B.4. Percentage Profit Increase Toward Baseline with Other Customer Choice Behaviors (Full Substitution in Left Block, Partial Substitution in Right Block) for Different Experimental Settings with Substitution Rates, Margins (Equal), and Densities for Either Player and Variable Customer Preferences

	Full substitution								Partial substitution							
	Operator 1				Operator 2				Operator 1				Operator 2			
Setting	Baseline	IBR-0	QMO	QMP	Baseline	IBR-0	QMO	QMP	Baseline	IBR-0	QMO	QMP	Baseline	IBR-0	QMO	QMP
H_H_H_0.5	12	145	−2	65	12	106	3	79	12	254	254	121	12	251	251	125
H_H_H_0.75	12	248	125	242	12	15	−139	0	12	254	254	242	12	253	253	2
H_H_H_1	12	254	254	254	12	14	−266	0	12	258	258	258	12	257	257	0
H_H_L_0.5	12	229	96	113	8	48	−53	19	12	250	250	116	8	179	179	42
H_H_L_0.75	12	251	176	242	8	13	−174	0	12	262	262	252	8	185	185	3
H_H_L_1	12	251	251	251	8	7	−302	0	12	253	253	253	8	194	194	0
H_L_H_0.5	8	93	−52	18	12	173	88	113	8	182	182	33	12	260	260	136
H_L_H_0.75	8	165	67	157	12	111	18	1	8	191	191	171	12	255	255	3
H_L_H_1	8	188	188	188	12	103	−63	0	8	189	189	189	12	249	249	0
H_L_L_0.5	8	149	−43	28	8	48	−49	26	8	186	186	21	8	182	182	33
H_L_L_0.75	8	187	76	169	8	14	−159	0	8	178	178	158	8	182	182	2
H_L_L_1	8	186	186	186	8	7	−271	0	8	182	182	182	8	180	180	0
L_H_H_0.5	6	27	−25	3	6	10	−28	0	6	24	24	1	6	26	26	0
L_H_H_0.75	6	26	2	7	6	6	−51	0	6	24	24	9	6	28	28	0
L_H_H_1	6	24	24	24	6	10	−68	0	6	29	29	29	6	23	23	0
L_H_L_0.5	6	26	10	1	4	5	−15	1	6	33	33	3	4	9	9	1
L_H_L_0.75	6	25	21	10	4	3	−13	0	6	27	27	8	4	4	4	0
L_H_L_1	6	33	33	33	4	1	−28	0	6	34	34	34	4	7	7	0
L_L_H_0.5	4	6	−7	1	6	19	13	2	4	10	10	1	6	30	30	2
L_L_H_0.75	4	10	3	5	6	21	13	0	4	7	7	3	6	26	26	1
L_L_H_1	4	6	6	6	6	19	3	0	4	8	8	8	6	25	25	0
L_L_L_0.5	4	6	0	1	4	6	1	1	4	6	6	0	4	9	9	1
L_L_L_0.75	4	5	2	0	4	5	−1	0	4	10	10	6	4	11	11	0
L_L_L_1	4	6	6	6	4	7	−4	0	4	8	8	8	4	10	10	0

baseline under different algorithms for either operator for all combinations of parameters.

Analogously, Table B.2 lists detailed results for the total profit (sum over both operators) toward the baseline and shows that the profit decreases in all instances when competition is present.

Table B.3 presents detailed information about attainable profit increases if the assumption of homogeneous payoffs does not hold.

Table B.4 outlines profit increases if customers do not choose vehicles completely at random if both operators have a vehicle available at this location.

References

Ackermann H, Röglin H, Vöcking B (2008) On the impact of combinatorial structure on congestion games. *J. ACM* 55(6):25:1–25:22.

Albiński S, Minner S (2019) Competitive rebalancing in one-way car-sharing. Working paper, Technical University of Munich, Munich, Germany.

Anshelevich E, Dasgupta A, Kleinberg J, Tardos E, Wexler T, Roughgarden T (2008) The price of stability for network design with fair cost allocation. *SIAM J. Comput.* 38(4):1602–1623.

Archetti C, Speranza MG, Vigo D (2014) Vehicle routing problems with profits. Toth P, Vigo D, eds. *Vehicle Routing: Problems, Methods, and Applications*, 2nd ed. (Society for Industrial and Applied Mathematics, Philadelphia), 273–297.

Balac M, Becker H, Ciari F, Axhausen KW (2019) Modeling competing free-floating carsharing operators—A case study for Zurich, Switzerland. *Transportation Res. Part C Emerging Tech.* 98:101–117.

Bellos I, Ferguson M, Toktay LB (2017) The car sharing economy: Interaction of business model choice and product line design. *Manufacturing Service Oper. Management* 19(2):185–201.

Berbeglia G, Cordeau J-F, Gribkovskaia I, Laporte G (2007) Static pickup and delivery problems: A classification scheme and survey. TOP 15(1):1–31.

BMW Group (2018) BMW Group und Daimler AG vereinbaren B"undelung ihrer Mobilit"atsdienste. Accessed January 23, 2019, https://www.press.bmwgroup.com/deutschland/article/detail/T0279654DE/bmw-groupund-daimler-ag-vereinbaren-buendelung-ihrer-mobilitaetsdienste.

Brook D (2004) Carsharing—Start up issues and new operational models. *Transportation Res. Board 83rd Annual Meeting*.

Bruglieri M, Colorni A, Luè A (2014) The relocation problem for the one-way electric vehicle sharing. *Networks* 64(4):292–305.

Bruglieri M, Pezzella F, Pisacane O (2017) Heuristic algorithms for the operator-based relocation problem in one-way electric carsharing systems. *Discrete Optim.* 23:56–80.

Bruglieri M, Pezzella F, Pisacane O (2018) A two-phase optimization method for a multiobjective vehicle relocation problem in electric carsharing systems. *J. Combin. Optim.* 36:162–193.

Cepolina EM, Farina A (2014) A methodology for planning a new urban car sharing system with fully automated personal vehicles. *Eur. Transport Res. Rev.* 6(2):191–204.

Fabrikant A, Papadimitriou C, Talwar K (2004) The complexity of pure Nash equilibria. Babai L, ed. *Proc. 36th Annual ACM Sympos. Theory Comput.* (Association for Computing Machinery, New York), 604–612.

Fagnant DJ, Kockelman KM, Bansal P (2015) Operations of shared autonomous vehicle fleet for Austin, Texas, market. *Transportation Res. Record J. Transportation Res. Board* 2536(1):98–106.

Feillet D, Dejax P, Gendreau M (2005) Traveling salesman problems with profits. *Transportation Sci.* 39(2):188–205.

Fernández E, Roca-Riu M, Speranza MG (2018) The shared customer collaboration vehicle routing problem. *Eur. J. Oper. Res.* 265(3):1078–1093.

George DK, Xia CH (2011) Fleet-sizing and service availability for a vehicle rental system via closed queueing networks. *Eur. J. Oper. Res.* 211(1):198–207.

Ghosh S, Varakantham P, Adulyasak Y, Jaillet P (2017) Dynamic repositioning to reduce lost demand in bike sharing systems. *J. Artificial Intelligence Res.* 58:387–430.

He L, Hu Z, Zhang M (2020) Robust repositioning for vehicle sharing. *Manufacturing Service Oper. Management* 22(2):241–256.

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

Karp RM (1972) Reducibility among combinatorial problems. Miller RE, Thatcher JW, Bohlinger JD, eds. *Complexity of Computer Computations*, The IBM Research Symposia Series, (Springer, Boston), 85–103.

Kortum K, Schönduwe R, Stolte B, Bock B (2016) Free-floating carsharing: City-specific growth rates and success factors. *Transportation Res. Procedia* 19:328–340.

Koutsoupias E, Papadimitriou C (1999) Worst-case equilibria. Meinel C, Tison S, eds. *Annual Sympos. Theoretical Aspects Comput. Sci., Lecture Notes Comput. Sci.*, vol. 1563 (Springer, Berlin, Heidelberg), 404–413.

Krishnamoorthy MS (1975) An NP-hard problem in bipartite graphs. *ACM SIGACT News* 7(1):26.

Krumke SO, Quilliot A, Wagler AK, Wegener J-T (2013) Models and algorithms for carsharing systems and related problems. *Electron. Notes Discrete Math.* 44:201–206.

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

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

Lien RW, Iravani SMR, Smilowitz K, Tzur M (2011) An efficient and robust design for transshipment networks. *Production Oper. Management* 20(5):699–713.

Monderer D, Shapley LS (1996) Potential games. *Games Econom. Behav.* 14(1):124–143.

Nair R, Miller-Hooks E (2011) Fleet management for vehicle sharing operations. *Transportation Sci.* 45(4):524–540.

Nash J (1951) Non-cooperative games. *Ann. Math.* 54(2):286–295.

Pandey V, Monteil J, Gambella C, Simonetto A (2019) On the needs for MaaS platforms to handle competition in ridesharing mobility. *Transportation Res. Part C Emerging Tech.* 108:269–288.

Parragh SN, Doerner KF, Hartl RF (2008) A survey on pickup and delivery models. Part II: Transportation between pickup and delivery locations. *J. Betriebswirtschaft* 58(2):81–117.

Paterson C, Kiesmüller G, Teunter R, Glazebrook K (2011) Inventory models with lateral transshipments: A review. *Eur. J. Oper. Res.* 210(2):125–136.

Pavone M (2015) Autonomous mobility-on-demand systems for future urban mobility. Maurer M, Gerdes JC, Lenz B, Winner H, eds. Autonomes Fahren (Springer Vieweg, Berlin), 399–416.

Perboli G, Ferrero F, Musso S, Vesco A (2018) Business models and tariff simulation in car-sharing services. *Transportation Res. Part A Policy Practice* 115:32–48.

Rosenthal RW (1973) A class of games possessing pure-strategy Nash equilibria. *Internat. J. Game Theory* 2(1):65–67.

Roughgarden T, Tardos É (2002) How bad is selfish routing? *J. ACM* 49(2):236–259.

Schulz AS, Stier Moses N (2003) On the performance of user equilibria in traffic networks. *Proc. Sympos. Discrete Algorithms* (Association for Computing Machinery and Society for Industrial and Applied Mathematics, New York), 86–87.

Sixt (2019) Sixt Share: Maximale Flexibilität und Vielfalt – Sixt verbindet Autovermietung und Carsharing. Accessed April 30, 2019, https://about.sixt.com/websites/sixt_cc/German/2999/news-details.html?newsID=1757303.

Ströhle P, Flath CM, Gärttner J (2019) Leveraging customer flexibility for car-sharing fleet optimization. *Transportation Sci.* 53(1):42–61.

Trentini A, Losacco F (2017) Sampling car-sharing data to evaluate urban traffic behaviour. Isaías P, Weghorn H, eds. *Internat. Conf. Appl. Comput.* (IADIS Press, Lisbon, Portugal), 295–299.

Vansteenwegen P, Souffriau W, Van Oudheusden D (2011) The orienteering problem: A survey. *Eur. J. Oper. Res.* 209(1):1–10.

Vasconcelos AS, Martinez LM, Correia GHA, Guimarães DC, Farias TL (2017) Environmental and financial impacts of adopting alternative vehicle technologies and relocation strategies in station-based one-way carsharing: An application in the city of Lisbon, Portugal. *Transportation Res. Part D Transportation Environ.* 57:350–362.

Volgenant T, Jonker R (1987) On some generalizations of the travelling-salesman problem. *J. Oper. Res. Soc.* 38(11): 1073–1079.

Vosooghi R, Puchinger J, Jankovic M, Sirin G (2017) A critical analysis of travel demand estimation for new one-way carsharing systems. *2017 IEEE 20th Internat. Conf. Intelligent Transportation Systems* (IEEE), 199–205.

Weikl S, Bogenberger K (2015) A practice-ready relocation model for free-floating carsharing systems with electric vehicles—Mesoscopic approach and field trial results. *Transportation Res. Part C Emerging Tech.* 57:206–223.

Wittenbrink P (2014) *Transportmanagement: Kostenoptimierung, Green Logistics und Herausforderungen an der Schnittstelle Rampe* (Springer-Verlag, Wiesbaden, Germany).