Competitive Analysis of Vehicle-Sharing Systems with Cournot Queueing Games

Matthew Sheldon, *Student Member, IEEE,* Dario Paccagnan *Member, IEEE*, Giuliano Casale, *Member, IEEE*

***Abstract*—Product-form closed queueing networks are a useful formalism for analyzing the availability of vehicles within a vehicle sharing system with stochastic user behavior. However, existing models assume that the provider has full control over the system. In reality, competition between providers within a single geographic area is quite common. In this paper, we introduce a non-cooperative game that extends a closed queueing model into a competitive environment, in which players decide the number of jobs that they wish to submit. This can be used to model vehicle sharing systems with multiple providers, in which they receive fares from trips but are responsible for the costs of their own fleet. The core technical results of this paper include conditions that guarantee the existence of a pure Nash equilibrium, and an efficient equilibrium-finding algorithm. We then present a case study, using our model, of a vehicle sharing system in Oslo with multiple providers. In this case study, we find that adding an additional competitor can increase the number of trips by up to** 18.9%**, and a highly competitive market can increase this by up to** 30%**.**

***Index Terms*—Vehicle Sharing, Closed Queueing Networks, Rational Queueing**

I. Introduction

The rise of vehicle sharing systems has led to situations in which multiple providers must compete for ridership in a given area. Many cities feature multiple ride-hailing providers such as Uber, Bolt, or Lyft. Furthermore, at present the city of Milan has three distinct bikesharing providers, the city of Bangkok has four, and the city of Berlin has seven [1]. In order to accurately model the effect of competition, policy makers and service providers may wish to analyze these transportation systems from a game-theoretic perspective.

A well-established modeling formalism for vehicle sharing systems is closed queueing networks [2]–[4]. This enables the analysis of vehicle sharing networks under stochasticity in customer arrivals and routing behavior. Furthermore, it is useful for finding performance measures of the system, such as the mean number of vehicles at a location and the probability of a user finding a vehicle. A review by Zardini, *et al.*, identifies steady-state queueing models as one of three primary methods of analysis of mobility-on-demand systems [4].

However, in order to analyze the impact of competition in vehicle sharing systems, we wish to analyze a situation where there are multiple providers which receive benefits from fares, with operational costs that vary with the number of vehicles. This necessitates a model where a closed queueing network is *shared* between multiple players, which receive a reward based on the throughput, while they incur a cost that varies with the number of jobs. We introduce the term *Cournot queueing games* for these models. This reflects the similarity to Cournot competition [5], [6]. In the case of a vehicle sharing system, the throughput is used to model the number of trips taken within the network. The pricing functions correspond to the operational cost of maintaining the fleet. Cournot games are typically defined over a continuous strategy space [6], but there are some existence guarantees in subsets of Cournot games with a discrete strategy space [7], [8].

a) Related Work: There are several papers that analyze the intersection of queueing theory and game theory. Beginning with a seminal paper by Naor [9], one of the core research areas is rational queueing. Rational queueing analyzes systems in which customers decide which queues to join, based on the performance metrics of the queueing system [9], [10]. Another important stream of research is networking games, in which players choose paths through an open queueing network for their [11]–[17]. There is significantly less work on games featuring closed queueing networks, where players separately choose network parameters. In [18], the authors analyze a closed queueing model in which each player controls the service rate of a queue.

There has been a significant amount of work on finding equilibria among transit behavior of users or drivers within shared mobility systems. [19] proposes a bilevel program for the planning of high-occupancy traffic lanes for ride-hailing systems, under strategic behavior of drivers. [20] includes the planning of origin-destination pricing strategies with shared rides. In [21], the equilibrium user behavior with a multimodal system provider is considered.

In contrast, there is a lesser degree of prior work that analyzes the impact of competition among providers within the context of vehicle sharing. In [22], competition is modeled as a Stackelberg game, with a municipality deciding prices and individual providers deciding prices and fleet sizes. However, this model is a deterministic graph-theoretic model, rather than a stochastic queueing-theoretic model. Therefore, it is less able to consider vehicle availability based on stochastic user arrivals. Furthermore, there is an assumption that the flow into and out of a node is balanced, which may not be the case in all circumstances. Finally, no problem-specific methods are proposed and therefore one must fall back on general methods for Stackelberg games, which are NP-hard in the general case [23]. In [24], the authors consider a

Matthew Sheldon, Dario Paccagnan, and Giuliano Casale are affiliated with Imperial College London

For correspondence: matthew.sheldon20@imperial.ac.uk

Manuscript received July 1, 2024; revised February 17, 2025.

The work of Dario Paccagnan was partly supported by the EPSRC grant EP/Y001001/1, funded by the International Science Partnerships Fund (ISPF) and UKRI.

profit-maximization problem among a duopoly, but it does not consider a queueing-theoretic model or generalize to more providers. [25] considers competing mobility providers who decide whether or not to closely monitor consumer honesty, and analyze evolutionary stable strategies. In [26], the authors introduce a method for dynamic pricing in a shared mobility system with multiple providers using reinforcement learning, but their method is focused on pricing rather than fleet sizing, and it is approximate rather than exact. In contrast with these prior works, our mobility application focuses on exact solutions to finding equilibrium fleet sizes within the context of a queueing model, which accounts for stochasticity, and enables detailed analysis of vehicle availability.

b) Contribution: In this work, we define and analyze a class of games in which providers decide the number of jobs that they submit, at cost, into a closed queueing network. We provide equilibrium proofs in multiple general cases, and then present an efficient, polynomial-time, algorithm for finding equilibria within these games. Finally, we present an experimental study that uses our formalism to analyze the impact of competition on the number of vehicles in a vehicle sharing system with competition between providers.

This paper extends our extended abstract in [27]. Firstly, we generalize our initial result in the symmetric case by allowing for non-linear prices, as shown in Theorem 3. Secondly, we provide applications using a variety of commonly-used cost functions in the literature. Thirdly, we present a two-priority system that allows for equilibrium existence with fewer restrictions, presented in Section V-C. Next, we present a more comprehensive description of the equilibrium-finding algorithm in Section VI. Finally, we present a more complete analysis of our empirical results, including the impact on the quantity of vehicles in addition to the throughput, as well as a more detailed view of the impact of inhomogeneous prices on the quantity of individual players in Section VII.

The rest of the paper is organized as follows. Section II introduces definitions and theoretical preliminaries for the underlying concepts in game theory and queueing theory. Section III introduces the game formulation. Section IV relates the queueing-theoretic model and the corresponding Cournot game to a vehicle sharing model. Section V presents an equilibrium proof in selected cases of Cournot queueing games. Section V presents an efficient algorithm for equilibrium finding. Section VI presents an application of this model to a new vehicle sharing system in Oslo with multiple providers. All proofs can be found in the online supplement [28].

II. Preliminaries

A. Queueing-Theoretic Preliminaries

We begin with a few queueing-theoretic preliminaries. Closed queueing networks (CQNs) are used to model a system with a finite and constant number of entities, commonly referred to as jobs, that transition throughout the network to receive service at different stations, which are referred to as queues. At each queue, jobs may wait for other jobs to complete service before entering service. Once service is completed, they transition to another queue. The service distribution is often described by a random variable, representing variability in the service time. Jobs within closed queueing networks can be assigned to different classes. In this paper, we use classes to differentiate jobs owned by different players, and we will refer to class-r jobs as jobs owned by player r. We consider networks in which all parameters, excluding the number of jobs, are identical for all classes. For more detail on closed queueing networks with heterogeneous classes, we refer to the reader to [23]. A CQN with identical parameters for each class excluding the number of jobs, can be described by a set of N_q queues, a routing matrix P and a number of jobs M [23]. The routing matrix $P = [\boldsymbol{P}_{ij}]$ is an $N_q \times N_q$ matrix, where P_{ij} is the probability that a job transitions to queue j after finishing service at queue i. The relative number of visits to a queue at i, when the system is in steady state, is called the visit ratio, e_i. The vector of visit ratios, $\boldsymbol{e}$, is equal to the left-eigenvector of P that corresponds to eigenvalue 1.

Each queue in an closed queueing network, with infinite-capacity queues, can be represented by Kendall's notation, $-/S/c/D$, where S represents the service process, c represents the number of servers, and D represents the service discipline. When S is an exponential distribution, it is represented by M (Markovian). An arbitrary real-valued distribution with positive support is represented as G, meaning general. It is common to assume that G queues are phase-type (PH) distributed, since as all positive-valued distributions with a rational Laplace transform can be approximated arbitrarily well by a PH distribution [29]. The number of servers is represented as $c \geq 1$. When c is infinite, no queueing effect exists and such a queue may be referred to as a delay. Following standard terminology, we refer to all stations within the queueing network as "queues", although there is no queueing effect at delays. The service discipline D represents the order in which jobs in the queue enter service. Common service disciplines include FCFS (first-come first-serve), PS (processor sharing) and SIRO (service in random order).

Several performance measures are used to analyze queueing networks. The throughput at a given queue i for jobs of class r, $T_{i,r}$, is the mean number of times class r jobs at queue i receive service per unit time. Where i is an arbitrary queue in the network, the system throughput of class r jobs is equal to

$$T_r = \frac{1}{e_i} T_{i,r} \tag{1}$$

This quantity is equal for any selected station [23]. Typically this is defined with regards to a reference station with $e_i = 1$. The overall system throughput, T, is equal to

$$T = \sum_r T_r \tag{2}$$

The response time, W_i, is the average time a given job takes to receive service. L_i represents the mean queue length at queue i. A less common performance measure, but one that is critical for our analysis, is the throughput per job, or response rate, $Z = \frac{T}{M} = W_i^{-1}$.

Product-form queueing networks are among the most analytically tractable queueing networks, and include single-class, Markovian closed queueing networks. The steady-state dynamics of a closed, product-form queueing network can be solved with Mean-Value Analysis (MVA) [30], [31]. MVA depends on the arrival theorem [32], which states that an arrival to a queue in steady state with M jobs, sees the same performance metrics as a system in steady state with $M-1$ jobs. This avoids the computation of normalizing constants, which may be expensive [23]. We point to [23], [33] for more background on queueing theory.

A typical closed queueing network model for a vehicle sharing system uses jobs to represent vehicles within the network. The queues represent locations at which vehicles wait for customers. This may be a physical point in the network, such as a bike rack or a taxi queue, or it may represent an abstraction of a larger area in which customers may choose between different vehicles. The service process of each queue represents customers arriving to access a vehicle, which is often modeled with an exponential distribution [4], [34], [35]. Delays are used to represent vehicles that are in transit between queues, which may be either physical roads or as an abstraction of the travel time. The queueing network can then be used to model different performance metrics that are of interest to VSS providers.

B. Game-Theoretic Preliminaries

In this section, we introduce relevant game-theoreticpreliminaries. We first introduce normal-form games, in which players simultaneously submit their strategies, and then receive utility U_r [36].

Definition 1. *A normal-form game consists of the following,*

- *A set of R players*
- *For each player $r \in 1 \dots R$, there is a set of pure strategies $\mathbf{S}_r$*
- *A set of all possible strategy profiles, $\mathbf{S} = \mathbf{S}_1 \times \mathbf{S}_2 \times \dots \times \mathbf{S}_R$*
- *For each player $r \in 1 \dots R$, there is a function $U_r : \mathbf{S} \to \mathbb{R}$, called a utility function, that represents the reward to player r.*

The notion of Nash equilibrium [37] is a widely-used tool for describing the behavior of non-cooperative, normal form games. A Nash equilibrium is a strategy profile where each individual player r has no incentive to change their strategy.

Definition 2. *A strategy profile $\mathbf{x} \in \mathbf{S}$ is a Nash equilibrium among pure strategies (equivalently, pure Nash equilibrium) if, for all players r and strategies $y_r \in \mathbf{S}_r$ (using the convention where x_r is the strategy for player r and $\mathbf{x}_{-r}$ is the strategy for all other players),*

$$U_r(y_r, \mathbf{x}_{-r}) \leq U_r(x_r, \mathbf{x}_{-r}) \tag{3}$$

Another important concept is the best response for player r, $b_r(\mathbf{S}_{-r})$. This defines a set of strategies-that maximize the utility, when the strategies for other players is $\mathbf{S}_{-r}$. In this paper, we focus on the least best response, $\underline{b}_r(\mathbf{S})$, as the set of strategies is ordered. Therefore, we can characterize the best response with a scalar rather than a set.

Not all finite games have pure equilibria, but there are some classes of games that do so. A special class of these is the class of potential games and their generalizations. These allow the game to be described by a single global function, known as a potential function. The generalization we use in this case is the ordinal potential game. In a finite ordinal potential game, there is guaranteed to be at least one Nash equilibrium among pure strategies, at a maximum of the potential function [38], [39].

TABLE I
TABLE OF SELECTED NOTATION

Notation	*Description*
R	Number of players
r, s	Player indices
Φ	A potential function
$\boldsymbol{S}$	Set of strategies
$x_r^{\min}$	Minimum number of jobs for player r
$x_r^{\max}$	Maximum number of jobs for player r
x_r	Number of jobs submitted by player r
$\boldsymbol{x}$	Vector of submitted jobs from each player
M	Total number of jobs in the system
x_{-r}	Total number of jobs submitted by all players except r
m	Background quantity of jobs
ψ_r	Cost function for player r
T	Total system throughput
Z	The response rate, $\frac{T(M)}{M}$
T_r	System throughput for player r
U_r	Utility function for player r

Definition 3. *A game G is an ordinal potential game if there is a function $\Phi : \boldsymbol{S} \to \mathbb{R}$, called an ordinal potential function, where*

$$\begin{aligned} &U_i(x_r, x_{-r}) - U_i(x'_r, x_{-r}) > 0 \\ &\Leftrightarrow \Phi(x_r, x_{-r}) - \Phi(x'_r, x_{-r}) > 0 \end{aligned} \tag{4}$$

III. GAME FORMULATION

Definition 4. *We define a Cournot queueing game as follows*

- *A set of R players*
- *A set of strategies $\mathbf{S} = [x_1^{\min}, x_1^{\max}] \times [x_2^{\min}, x_2^{\max}] \times \dots \times [x_R^{\min}, x_R^{\max}]$, where $[x_r^{\min}, x_r^{\max}]$ is an interval of $\mathbb{N}$*
- *A closed queueing network*
- *A fixed background quantity $m \in \mathbb{N}_{\geq 0}$, which represents a constant quantity of jobs in the network that is independent of the strategy of any player*
- *For each player, a function $\psi_r(x_r) : \mathbb{N} \to \mathbb{R}$, which represents the cost of submitting x_r jobs into the system*

The strategy of each player is represented by a scalar value $x_r \in \mathbb{N} \cap [x_r^{\min}, x_r^{\max}]$, which represents the number of jobs they submit into the system, with maximum value $x_r^{\max}$. Where T_r represents the throughput for player r and $x_{-r} = \sum_{s \neq r} x_s$, the reward for player r equals

$$U_r(x_r, \mathbf{x}_{-r}) = \begin{cases} T_r(x_r, \mathbf{x}_{-r}, m) - \psi_r(x_r) & x_r \neq 0 \\ 0 & x_r = 0 \end{cases} \tag{5}$$

Our first goal is to prove the existence of pure equilibria. Emphasizing pure, integral strategies allows for a seamless integration of game-theoretic analysis and queueing analysis. Many techniques for solving closed queueing networks, such as mean-value analysis [30], [31] or techniques based on the normalizing constant [23], depend on a number of recursive relationships of the performance over the number of jobs. Accurate analysis of these performance measures, such as queue lengths or waiting time, may require full knowledge of the distribution of jobs at a given mixed equilibria, whereas they can be calculated easily if an integral solution is found.

This game is most tractable in the case where the closed queueing network is product-form [23]. In addition, we add another restriction that the service rates and routing probabilities are identical for the jobs of all players.

Assumption 1. *The underlying closed queueing network contains contains $-/M/c$ queuing queues and $-/G/\infty$ delays, and all players have identical service processes and routing probabilities. Furthermore, the service discipline at each queuing queue is one of the following,*

- *First-come, first-served (FCFS)*
- *Last-come, first served with preemptive resume (LCFS-PR)*
- *Processor sharing (PS)*
- *Service-in-random-order (SIRO)*

Proposition 1. *Let G be a Cournot queueing game that fulfills Assumption 1. Then, where T is the system throughput,*

$$T_r(x_r, \mathbf{x}_{-r}) = M(x_r, \sum_{s \neq r} x_s + m) = \frac{x_r}{\sum_s x_s + m} T(\sum_s x_s + m) \tag{6}$$

$$U_r(x_r, \mathbf{x}_{-r}) = \begin{cases} \frac{x_r}{\sum_s x_s + m} T(\sum_s x_s + m) - \psi_r(x_r) & x_r \neq 0 \\ 0 & x_r = 0 \end{cases} \tag{7}$$

With this assumption, which is quite general in practice, we can now treat the game as a quantity game. Using this, we will let $x_{-r} = \sum_{s \neq r} x_s$ represent the sum of quantities of all players, with the exclusion of r.

IV. Vehicle Sharing System Model

Next, we describe how the Cournot queueing games model can be used to describe a Vehicle Sharing System with multiple providers. Each mobility provider corresponds to a player within the game. The jobs within the queueing network represent vehicles. There are N_q queues within the network, of which there are N_s $-/M/1$ queues and N_d $-/G/\infty$ delays. We let the single-server queues be indexed first. The $-/M/1$ queues represent locations at which empty vehicles are parked, and may be accessed by customers. The $-/G/\infty$ delays represent the travel process of vehicles with customers. We also use τ_i to represent the average travel time in delay i.

We consider profit functions of the form, where $\alpha_{k,r}$ are arbitrary constants, and $f^c_{l,r}$ are convex functions of x_r.

$$P_r(\boldsymbol{x}) = \sum_k F_k(\boldsymbol{x}) \tag{8}$$

$$F_k(\boldsymbol{x}) = \alpha_{k,r} T_r(\boldsymbol{x}) - f^c_{k,r}(x_r) \tag{9}$$

Under the case that $\sum_k \alpha_{k,r} > 0$, we can set

$$\psi_r(x_r) = \frac{1}{\sum_k \alpha_{k,r}} \sum_k f^c_{k,r}(x_r) \tag{10}$$

This function is convex, and the following utility function is both the form used in the Cournot queueing game formalism. It is also ordinally equivalent to P_r, which implies that the best-response correspondences and therefore the equilibria are identical when it is used instead [40].

$$U_r(\boldsymbol{x}) = T_r(\boldsymbol{x}) - \psi_r(x_r) \tag{11}$$

Several other costs that may be relevant to VSS providers can be expressed as functions that are affine over the throughput and convex over the number of jobs, and can be used in the same format as (9). Therefore, they can be represented in a Cournot queueing game. We present a few options for terms F_k that may be of use in modeling real-world systems.

A. Heterogeneous Fares

There may be substantial differences in the price paid by customers when renting vehicles from different locations. Furthermore, trips between different pairs of locations may be priced differently. We let d_i be the mean fare at $i \leq N_s$, while the number of trips departing at i is equal to $T_{i,r} = e_i T_r$. Thus, we can account for these in the profit function, using the following subterm

$$F^{fare}(\boldsymbol{x}) = \sum_{i=1}^{N_s} d_i e_i T_r(\boldsymbol{x}) \tag{12}$$

B. Transit Costs

While the customer may be responsible for the majority of the transit costs, there may be additional costs to the provider. This can include mileage costs, depreciation and wear-and-tear. To model this, we can include arbitrary fixed costs that vary with the throughput along each delay, as well as variable costs that reflect the time spent along each delay.

Using Little's law, we can find the mean number of player r vehicles in a particular delay i.

$$L_{i,r} = \tau_i e_i T_r \tag{13}$$

Let c^t_i be the cost per unit of time a vehicle is in transit, and c^f_i be a fixed cost for transiting. Then, we can represent the mean cost of transiting vehicles along delay i as

$$(c^f_i e_i + c^t_i \tau_i e_i) T_r \tag{14}$$

Then, we can represent the transit cost as $F^{transit}(\boldsymbol{x})$.

$$F^{transit}(\boldsymbol{x}) = -\sum_{i=N_s+1}^{N_q} (c^f_i e_i + c^t_i \tau_i e_i) T_r \tag{15}$$

C. Vehicle Availability Penalties

A lack of vehicle availability may have more significant effects for providers than just lost revenue. This may include reputational costs or future lost trips. A common approach in fleet sizing is to include additional penalties based on the utilization of each queue [41]–[43] We now show that these penalties can be included without loss of generality.

The utilization of vehicles owned by player r at queue i is

$$B_{i,r}(\boldsymbol{x}_r) = \frac{e_i}{\mu_i} T_r(\boldsymbol{x}_r) \tag{16}$$

Let c_i^a be a penalty term for the lack of availability at a station i. Then, the total dis-availability cost can be represented by a function $F^{avail}(\boldsymbol{x})$.

$$F^{avail}(\boldsymbol{x}) = \sum_{i=1}^{N_s} c_i^a(1 - B_{i,r}) \tag{17}$$

Rearranging, we get a formulation with the form in (9).

$$F^{avail}(\boldsymbol{x}) = -\sum_{i=1}^{N_s} c_i^a \frac{e_i}{\mu_i} T_r(\boldsymbol{x}) + N_s \tag{18}$$

D. Parking Costs

An important operational consideration in vehicle sharing systems is the cost of parking [44], [45]. This may concern parking fees, payment for chargers, fixed costs for wear-and-tear when vehicles are not in motion, or wages for drivers when they are not taking a customer.

The number of jobs that are in transit can be calculated as follows, using Little's law.

$$x_r^{transit} = \sum_{i=N_s+1}^{N_q} e_i \tau_i T_r(\boldsymbol{x}) \tag{19}$$

Therefore, the number of waiting jobs at each queue is

$$x_r^{waiting} = x_r - (\sum_{i=N_s+1}^{N_q} e_i \tau_i) T_r(\boldsymbol{x}) \tag{20}$$

Where c^p is the cost of parking, we can derive $F^{park}(\boldsymbol{x})$, a term that represents the parking cost within the system.

$$F^{park}(\boldsymbol{x}) = c^p \sum_{i=N_s+1}^{N_q} e_i \tau_i T_r(\boldsymbol{x}) - c^p x_r \tag{21}$$

E. Rebalancing Costs

A common method to model rebalancing, as proposed in [2], [4], [35], is to use the service rate to model both manual rebalancing trips from a given station, as well as customers accessing vehicles. Rebalancing trips, either using drivers [2], or autonomous rebalancing [35], can be modeled as "virtual customers". We can decompose the service rate μ_i into μ_i^c and μ_i^r. In this case, μ_i^c represents the rate at which customers arrive to access vehicles, and μ_i^{rb} is the rate at which vehicles exit a queue to be moved elsewhere through rebalancing.

$$\mu_i = \mu_i^c + \mu_i^{rb} \tag{22}$$

There is a condition that the service demands must be equal for all classes, which implies that there must be homogeneity in the rebalancing behavior. The throughput for trips with vehicles is equal to

$$T_{i,r}^c(\boldsymbol{x}) = \frac{\mu_i^c}{\mu_i^c + \mu_i^{rb}} T_{i,r}(\boldsymbol{x}) \tag{23}$$

Likewise, the throughput for rebalancing trips is equal to

$$T_{i,r}^{rb}(\boldsymbol{x}) = \frac{\mu_i^{rb}}{\mu_i^c + \mu_i^{rb}} T_{i,r}(\boldsymbol{x}) \tag{24}$$

As the throughput is no longer equal to the number of trips taken from a particular queue, we must revise F^{fare}. We label the revision as F^{fr}.

$$F^{fr}(\boldsymbol{x}) = \sum_{i=1}^{N_s} d_i e_i \frac{\mu_i^c}{\mu_i^c + \mu_i^{rb}} T_r(\boldsymbol{x}) \tag{25}$$

We can also include the fixed costs of rebalancing as F^{reb}, where c_i^{rb} is the fixed cost for rebalancing from station i.

$$F^{reb}(\boldsymbol{x}) = -\sum_{i=1}^{N_s} c_i^{rb} e_i \frac{\mu_i^{rb}}{\mu_i^c + \mu_i^{rb}} T_r(\boldsymbol{x}) \tag{26}$$

Variable costs for rebalancing can be constructed by using different delays for rebalancing vehicles and vehicles with customers. Where $1^{rb}(i)$ is the indicator function that indicates if a delay is used for rebalancing, c_i^t is a parameter that represents the cost per unit time that a vehicle is being rebalanced along delay i, we represent this term as $F^{rt}(\boldsymbol{x})$.

$$F^{rt}(\boldsymbol{x}) - \sum_{i=N_s+1}^{N_q} 1^{rb}(i)(c_i^t \tau_i e_i) T_r \tag{27}$$

V. Equilibrium Existence

A. Case 1: Identical Pricing Functions

We first present a more tractable case where all pricing functions are identical across all players. The following hold in general for discrete symmetric Cournot games with a concave demand function, and are relatively straightforward yet novel results. However, they are of particular interest to our Cournot queueing game formalism, and we present the results in terms of a Cournot queueing game.

Theorem 1. *Let G be a Cournot queueing game, under Assumption 1, where all players have an identical, linear pricing function $\psi_r(x_r) = px_r$, where p is a constant price. Furthermore, let $x_r^{\min} > 0$ for all players r. Then G is an ordinal potential game with potential function,*

$$\Phi(\boldsymbol{x}) = \frac{\prod_s x_s}{m + \sum_s x_s} T(m + \sum_s x_s) - p \prod_s x_s \tag{28}$$

Furthermore, G has at least one Nash equilibrium among pure strategies at the maximum of $\Phi(\boldsymbol{x})$.

When pricing functions are convex, there is no longer a potential function. However, we can still guarantee the existence of an equilibrium. Throughout the paper, we will also assume that prices are convex.

Assumption 2. *ψ_r is convex over x_r.*

The following theorem establishes a few properties of the shape of the throughput function. These are particularly useful for establishing our existence and uniqueness results, as presented later in the paper.

Theorem 2. *[46], [47] In a closed queueing network that meets Assumption 1, throughput is concave, non-decreasing, and anti-starshaped as a function of the number of jobs.*

Theorem 2 describes the well-known fact that steady-state throughputs for single-class Markovian queueing networks are concave over the number of jobs. The following result shows that this remains true among the number of jobs that an individual player submits into the system. It must be noted that this is not an immediate property of the throughput of T, as concavity is not necessarily preserved by multiplication, so this is a slightly stronger result.

Proposition 2. *Under Assumption 1, the throughput for player r, $T_r(x_r, x_{-r}+m) = x_r Z(x_r + x_{-r} + m)$, is concave over x_r.*

The next result is a generalization and simplification of Theorem 2 in [8]. In particular, the original theorem is a proof by contradiction that rests on the inverse demand function being concave and non-increasing. In this case, the inverse demand function is the response rate Z, which is non-increasing but not necessarily concave.

Theorem 3. *Let G be a Cournot queueing game where all players have an identical, convex pricing function, ψ. Then, G has at least one equilibrium among pure strategies.*

B. Case 2: Heterogeneous Cost Functions

The game is no longer symmetric when the cost functions are allowed to differ between players, and therefore the analysis is more complicated. Broadly speaking, the proofs in this paper result from the fact that the best responses are decreasing with respect to the background quantity, a property known as strategic substitutes. This does not imply the existence of an equilibrium in general [48], but we establish it in a few specific cases. In this section, we qualify when pure equilibria are guaranteed, under moderate conditions on the strategy set. We require the strategy set to fulfill at least one of two criteria under every pure strategy profile. The first, which we term the non-monopoly criterion, allows providers to submit one more job than the total of all other jobs in the system, but no more than that. The second, which we term the light-load criterion, imposes a restriction on the second finite difference of the response rate, that is consistent with a queueing network in light load. Strategic substitutes may not exist when neither of these conditions apply.

Assumption 3. *The strategy set $\mathbf{S}$ is restricted such that at least one of these two conditions hold for all possible pure strategy profiles $\mathbf{x}$,*

- *(non-monopoly) $x_r \le x_{-r} + m + 1$, for all players r, or*
- *(light load) $(\sum x_s + 1)\Delta^2 Z(m + \sum x_s) + \Delta Z(m + \sum x_s) \le 0$*

When a particular provider has a large market share, the best response to a particular quantity may not be decreasing. Qualitatively, this may be interpreted as monopolistic providers having an incentive to maintain both vehicle availability as well as their market share, and they may occasionally increase their fleet size to maintain this. However, when there are few vehicles in the system, this market share effect is lessened. Based on the decreasing responses property and the fact that this is a quantity game, we can apply Tarski's fixed-point theorem [49] to get the following existence result.

Theorem 4. *Let G be a Cournot queueing game that meets Assumptions 1-3. Then there is at least one equilibrium among pure strategies.*

One way in which the non-monopoly criterion may be enforced by adequately restricting the strategy set. For example, this is true where $x^{\max} \le (R-1)x^{\min} + 1 + q$, assuming all players share the same $x^{\min}$ and $x^{\max}$. This may be limiting in cases where there are a low number of players and a low background quantity m. The next result implies that the equilibrium quantity, under the least best response correspondence, is also unique under heterogeneity. This corresponds to an intuitive assumption that providers will not submit extra vehicles if they have the same estimated profit. In practice, since the throughput function is rarely linear, multiple best responses will be rare.

Theorem 5. *Let G be a Cournot queueing game, that meets Assumptions 1 and 2, and has a common pricing function ψ. Then, the equilibrium quantity based on the least best response correspondence, $M^{eq}(\underline{b}_r)$, is unique.*

C. Ensuring an Equilibrium with a Two-Priority System

Assumption 3 may be restrictive, particularly in cases where there are a small number of players. We can generalize this result in the case in which the queueing system has two priority classes. If at least one job with priority 1 is in a queue, it will be served first. In the case where only priority 1 jobs are restricted, then this game can be guaranteed to have an equilibrium in all cases with heterogeneous prices. However, Assumption 1 is not met, and the algorithm described later in Section VI does not work in this case.

When a monopolist is present, their best response may be increasing over the quantity of other jobs. This may result in instability of fleet sizes as well as vehicle oversupply. By implementing a priority-based system, this guarantees that additional jobs from a monopolist only improve the availability of vehicles, rather than being used to maintain a larger proportion of trips.

Theorem 6. *Let G be a Cournot queueing game, that meets Assumption 2, in which jobs are given two priorities. If any player has more than half the total number of jobs, any further jobs are allocated into the second priority class. Then, G has a unique equilibrium among pure strategies.*

VI. Equilibrium Computation

In this section, we present an efficient, unified algorithm for finding equilibria and solving the closed queueing network, if the pricing function is convex. In [50] and [51], Novshek introduces an elegant method for analyzing and solving aggregative games. These works introduce the concept of backwards reply correspondence, or finding optimal strategies that can lead to a given total quantity M. In other words, finding strategies x_r, from M, such that $x_r \in b_r(M - x_r)$. If there is a strategy profile that sums to M, an equilibrium is found. When the best responses are decreasing, Kukushkin showed that an equilibrium can be found in polynomial time [52], with

Algorithm 1 Equilibrium Finding Algorithm

$L_i, R_i, T, \Delta T \leftarrow 0$
$minPivots \leftarrow ARRAY(R, init = \infty)$
$\mathbf{x}^{lo}, \mathbf{x}^{hi} \leftarrow ARRAY(R, init = 0)$
for $M = \sum_s x_s^{\min} \dots \sum_s x_s^{\max}$ **do**
  $R_i \leftarrow ResponseTime(L_i, e_i, \mu_i)$
  $\Delta T \leftarrow \frac{M}{\sum e_i R_i} - T$
  $T \leftarrow T + \Delta T$
  $isValid \leftarrow TRUE$
  for $r = 1 \dots R$ **do**
    $x_r^{pivot} \leftarrow GetPivot(T, \Delta T, \psi_r, \mathbf{x}^{\min}, \mathbf{x}^{\max}, M)$
    if $pivot = NIL$ **then**
      $isValid \leftarrow FALSE$
      continue
    end if
    $x_r^{lo}[r] \leftarrow x_r^{pivot}$
    $x_r^{hi}[r] \leftarrow \min(x_r^{\max}, M - \sum_{s \neq r} x_s^{\min},$
        $minPivots[r] + M)$
    if $x_r^{lo}[r] > x_r^{hi}[r]$ **then**
      $isValid \leftarrow FALSE$
      break
    else
      $minPivots[r] = \min(minPivots[r], x_r^{pivot} - M)$
    end if
  end for
  if $M \in [\sum_r x_r^{lo}, \sum_r x_r^{hi}]$ **and** $isValid$ **then**
    return $[M, T, \mathbf{x}^{lo}, \mathbf{x}^{hi}]$
  end if
end for

$O(R^3|\mathbf{S}|^2)$ time complexity. We present a novel and efficient algorithm in this section, with a lower time complexity, under Assumptions 1 and 2. This can be solved efficiently in tandem with the closed queueing network, using the MVA algorithm. A pseudocode of this is provided in Algorithm 1. We first present a few theoretical results that guarantee the correctness of this algorithm, and then show how to determine the pivot values. We find equlibria based off the least-best response correspondence, $\underline{b}_r$, to ensure uniqueness of best replies.

We first express the utility as a function of the strategy of player r, x_r, and the total number of jobs M. Note that this is a slightly different definition than $U_r(x_r, x_{-r})$.

$$U_r(x_r, M) = x_r Z(M) - \psi_r(x_r) \tag{29}$$

The forward difference is equal to, where M increases along with the strategy x_r,

$$\Delta_{x_r} U_r(x_r, M) = x_r \Delta Z(M) + Z(M+1) - \Delta\psi_r(x_r) \tag{30}$$

Theorem 7. *Let Assumptions 1 and 2 be true. Then, the following, together, are necessary and sufficient conditions for x_r to be the least best response of $x_{-r} + q$, where $M = x_r + x_{-r} + q$,*

1) $\Delta_{x_r} U_r \leq 0$, *if* $x_r < x_r^{\max}$
2) $x_r \in [x_r^{\min}, x_r^{\max}]$
3) $x_r + \sum_{s \neq r} x_s^{\max} + q \geq M$
4) $x_r + \sum_{s \neq r} x_s^{\min} + q \leq M$
5) *There is not another pair (M', x_r') such that $x_r' < x_r, M' < M, x_r - x_r' = M - M'$, and all previous conditions apply*

This implies that the set of best responses corresponding to a total quantity is an interval for each player, if it exists. Then, we can construct an algorithm that finds an equilibrium in polynomial time. The following result gives the form of this interval, in terms of a quantity known as $x_r^{pivot}(M)$. This is a lower bound for the backward reply correspondence to M, and can be used to find upper bounds for larger quantities. In particular, we use these results to construct 3 intervals, $I_r^1(M)$, $I_r^2(M)$ and $I_r^3(M)$, that meet increasing subsets of each of these conditions for a given player r and total quantity M. $I_r^3(M)$ is the most relevant for the algorithm, but the other two are useful for estimating the pivot value. A pseudocode of this algorithm is given in Algorithm 1.

Lemma 1. *The set of x_r that meets conditions 2-4 for a given total quantity M, is either the empty set $\emptyset$, or an interval of the form,*

$$I_r^1(M) = [\max(x_r^{\min}, M - m - \sum_{s \neq r} x_s^{\max}), \\ \min(x_r^{\max}, M - m - \sum_{s \neq r} x_s^{\min})] \tag{31}$$

Lemma 2. *Under Assumptions 1 and 2, the set of x_r that meets conditions 1-4 for a given total quantity M, is either the empty set $\emptyset$, or an interval of the form,*

$$I_r^2(M) = [x_r^{pivot}(M), \min(x_r^{\max}, M - q - \sum_{s \neq r} x_s^{\min})] \tag{32}$$

Where $x_r^{pivot} \in I_r^1(M)$. Furthermore, either $x_r^{pivot} = x_r^{\max}$ or x_r^{pivot} is the least strategy x_r such that $\Delta_{x_r} U_r(x_r, M) \leq 0$

Theorem 8. *Let Assumptions 1 and 2 be true. Let M be the total job quantity under a given strategy profile, $M = \sum_r x_r$. Then, the set of x_r such that $x_r = \underline{b}_r(M - x_r)$ is either the empty set $\emptyset$, or an interval of the form*

$$I_r^3(M) = \left[x_r^{pivot}(M), I_r^{3\max}(M)\right] \tag{33}$$

$$I_r^{3\max}(M) = \min(\min_{M' < M}\{x_r^{pivot}(M') - M'\} \\ + M - 1, \\ x_r^{\max}, M - m - \textstyle\sum_{s \neq r} x_s^{\min}) \tag{34}$$

In addition to establishing an efficient algorithm for finding pure equilibria, another direct result is a guarantee of the existence of an equilibrium for symmetric Cournot queueing games, that is almost symmetric. It is worth noting that this relies on an equilibrium existence result which is shown later in Theorem 3.

Corollary 1. *If G is a symmetric Cournot queueing game what follows Assumption and 1 and 2, then there is an equilibrium where the minimum strategy among all players and the maximum strategy differ by no more than 1 job.*

The throughput can be efficiently calculated for each background quantity M, using the MVA algorithm. We have shown that the set of all best responses $x_r = \underline{b}_r(M - x_r)$, if it exists, is an interval, which we temporarily denote as $[x_r^{lo}, x_r^{hi}]$. Then, it is clear that a Nash equilibrium can be found at an index where $M \in [\sum_s x_s^{lo}, \sum_s x_s^{hi}]$, where all intervals are non-empty for each player, which can be established in linear time. With only a marginal amount of additional computation, the equilibrium vector $\mathbf{x}$ can be found via a greedy algorithm. A detailed pseudocode of this algorithm can be found in Algorithm 1. If $\Delta Z(M) < 0$, then x_r^{pivot} can be found in logarithmic (in the general case, using binary search) or constant time (when prices are linear or affine, with an analytic solution) by considering the floor and ceiling by first considering the following equation, as an estimate of the value of x_r^{pivot}, by equating (30) with 0 and inverting,

$$\tilde{x}_r^{pivot} = \frac{\Delta\psi_r(x_r) - Z(M+1)}{\Delta Z(M)} \tag{35}$$

Given that x_r^{pivot} must be discrete, it can be found by the following equation, when $\Delta Z(M) < 0$,

$$x_r^{pivot} = \begin{cases} \lfloor \tilde{x}_r^{pivot} \rfloor & \Delta_{x_r} U_r(\lfloor \tilde{x}_r^{pivot} \rfloor, M) \leq 0, \\ & \quad \lfloor \tilde{x}_r^{pivot} \rfloor \in I_r^1(M) \\ \lceil \tilde{x}_r^{pivot} \rceil & \lceil \tilde{x}_r^{pivot} \rceil \in I_r^1(M), \\ & \quad \text{and the prior case does not hold} \\ x_r^{\max} & \lfloor \tilde{x}_r^{pivot} \rfloor > I_r^1(M) \\ \min I_r^1(M) & \lceil \tilde{x}_r^{pivot} \rceil \leq I_r^1(M) \end{cases} \tag{36}$$

This equation follows immediately from the fact that x_r^{pivot} is the least value within $I_r^1(M)$ where $\Delta U \leq 0$, if it exists. If a value does not exist, then it is equal to $x_r^{\max}$. Otherwise, if $\Delta Z(M) = 0$, then $x_r^{pivot} = \min I_r^1(M)$ if $\Delta\psi_r(x) \geq Z(M+1)$, and it does not exist otherwise. If none of these cases apply, then x_r^{pivot} does not exist and $I_r^2(M) = \emptyset$.

Proposition 3. *Let $F^{\max}$ represent the maximum number of strategies per player, $\max_s 1 + x_s^{\max} - x_s^{\min}$, and $M^{\max}$ represent the maximum feasible quantity, $\sum_s x_s^{\max}$. Algorithm 1 has a time complexity of*

$$O(RF^{\max}(N_q + R\log(F^{\max}))) \tag{37}$$

If this the pricing function is affine, the time complexity is

$$O(RF^{\max}(N_q + R)) \tag{38}$$

In practice, this algorithm is very efficient. In a random experiment, we analyzed equilibria for $10,000$ closed queueing networks. With 20 players, linear prices, precomputed throughputs, and strategies within $[10, 190]$, we find equilibria for all $10,000$ games within a total of 578 seconds (or 0.0578 seconds per individual game). In comparison, the total number of strategy profiles is 1.27×10^{45}.

VII. Experimental Results

In this section, we present an empirical case study, involving the design of a mobility system with multiple distinct providers in Oslo. We use a dataset[1] of 1,116,344 trips in 2021, in order to infer the service rate and routing probability at each queue. In order to parametrize the model, we represent each bikesharing rack within the network as a queue. Although the travel time between queues may be heterogeneous, the mean queue length and throughput at each of the queueing queues in steady-state only varies with the mean time at each delay. Therefore, we only need to use a single delay with exponentially-distributed service rates to model the transit process of vehicles in use between queues. We use the average hourly number of trips between each queue pair between 10:00 and 15:00, on weekdays, as the service rate at each queue. As the mean travel time is equal to 12.74 minutes, the service

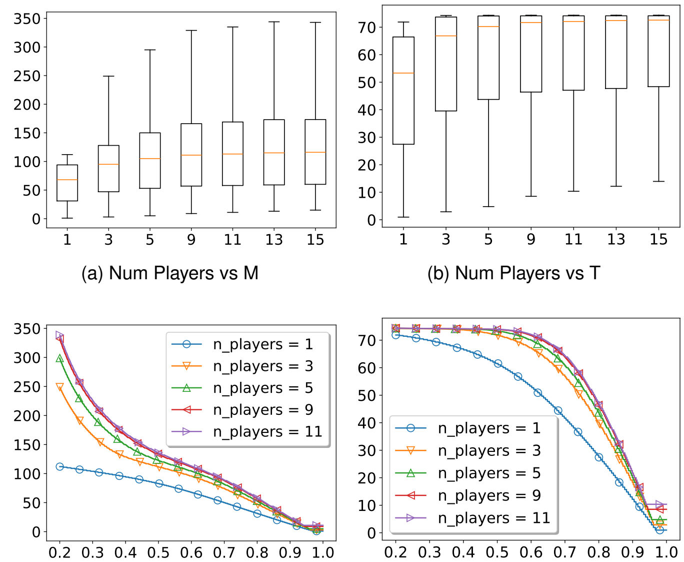

(a) Num Players vs M (b) Num Players vs T

(c) Price vs M (d) Price vs T

Fig. 1. Experiment with Single Prices (M: Total Quantity, T: Throughput)

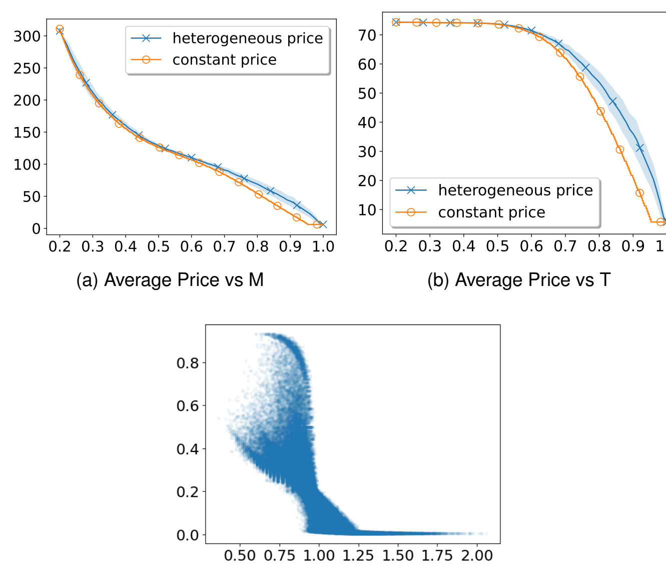

(a) Average Price vs M (b) Average Price vs T

(c) $p_r / \frac{1}{n}\sum p_r$ vs x_r/M

Fig. 2. Experiment with Heterogeneous Prices (6 players, M: Quantity, T: Throughput)

[1]https://oslobysykkel.no/en/open-data/historical

rate at this delay is set to $\frac{1}{12.74} = 0.078$ vehicles/min.

We present two experiments, one with constant prices and a varying number of customers, and another with 6 providers with heterogeneous prices. In the first case, we consider a varying number of players within the range $[1, 15]$, with prices sampled uniformly within the range $[0.2, 1.0]$. Each player can choose, in principle, between 1 and 5000 jobs, although no equilibrium features a strategy even close to this upper bound. For each player count, $20,000$ games were analyzed. Results for this experiment are presented in Figure 1. Figure 1a shows the quantity against the number of players. The results suggest that the relationship is concave. Figure 1b shows the number of players versus the throughput. It should be noted that the throughput appears to reach large-market behavior faster than the quantity. Figures 1c and 1b represent both the quantity and the throughput as a function of price, both of which imply that the market dynamics quickly approaches the large market behavior when the number of players is greater than five.

On average, adding a single competitor increases the number of trips taken by 14.1%, for a given network. Adding two competitors increases this amount by 18.9%. On the other hand, increasing the number of providers to 20 results in a 30.0% increase in the number of trips. On the other hand, increased competition may result in vehicle oversupply, as a single competitor results, on average, in a 29.5% increase in the number of vehicles. This indicates that there is a tendency towards vehicle oversupply under competition. Qualitatively, this may be explained by the fact that a provider r adding a new vehicle to the network increases the number of trips in two ways. The first is that the new vehicle may increase availability, allowing more trips across the network to be filled. The second is that the new vehicle increases the likelihood that an arriving customer chooses provider r over the other providers. Only the first effect increases the aggregate number of trips in the system, but providers still have an incentive to add more vehicles to increase their market share.

In Figure 2, we present the results of a heterogeneous experiment with 6 players. We analyze $20,000$ games. We select a common price p_{common} uniformly within the range $[0.2, 1.0]$. Then, the price for each player, p_r, is distributed according to $\min(\max(p_{common} + \mathcal{N}(0, 0.1), 0.2), 1.0)$. The strategy range for each player is set to $[1, 66]$ to ensure compatibility with Assumption 3.

Figure 2a shows the average price versus the quantity, in comparison with a homogeneous game with similar prices. Figure 2b shows the throughput in a similar manner. Figure 2 shows the ratio of the price of a player to the average price in the game, versus the ratio of the equilibrium quantity of the player to the total equilibrium quantity. We note that the quantity and throughput is generally higher than the game with a common price of $\frac{1}{r}\sum(p_r)$, especially at higher prices. In Figure 2c, we present a graph comparing the ratio of the quantity of a given player to the total quantity, versus the ratio of their price to the average price. This graph shows a steep S-shaped boundary, which suggests that small changes in price can lead to a significant difference in the equilibrium quantity for a given player. Furthermore, it implies that even a slightly higher price results in a dramatically reduced quantity at equilibrium. Together, these imply that while heterogeneity can result in more vehicle access, it can also cause monopolistic behavior in the market. These results may be relevant to policymakers that have the choice to admit providers with different cost-structures, such as taxis as opposed to ride-hailing systems. In both cases, we observe a decreasing relationship between the M and the price. Furthermore, the relationship between T and the price appears to be concave, when all players submit below the maximum quantities.

VIII. Conclusion

In conclusion, we have presented the notion of Cournot queueing games, a game in which a closed queuing network is shared between players, who must decide their job population. We have then shown that this class of games has an equilibrium among pure strategies under moderate conditions, and we have presented and demonstrated an efficient algorithm for finding equilibria. Then, we have applied this formalism to a vehicle-sharing system model and presented a case study parametrized by Bysykkel bike-sharing system in Oslo. In the future, it may be useful to extend these results to networks including finite-capacity queues, in order to more accurately model situations in which there are limitations on the queue length, such as when there are a finite number of parking spaces or chargers.

References

[1] "List of bicycle-sharing systems," publication Title: Wikipedia. [Online]. Available: https://en.wikipedia.org/wiki/List_of_bicycle-sharing_systems

[2] R. Zhang and M. Pavone, "A Queueing Network Approach to the Analysis and Control of Mobility-On-Demand Systems," in *2015 American Control Conference (ACC)*. IEEE, 2015, pp. 4702–4709.

[3] A. Waserhole and V. Jost, "Pricing in vehicle sharing systems: optimization in queuing networks with product forms," *Euro journal on transportation and logistics*, vol. 5, pp. 293–320, 2014.

[4] G. Zardini, N. Lanzetti, M. Pavone, and E. Frazzoli, "Analysis and control of autonomous mobility-on-demand systems," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, pp. 633–658, 2022.

[5] A. A. Cournot, *Recherches sur les principes mathématiques de la théorie des richesses*. L. Hachette, 1838.

[6] C. Morrison, "Cournot, bertrand, and modern game theory," *Atlantic Economic Journal*, vol. 26, no. 2, 1998.

[7] P. Davis, "Estimation of Quantity Games in the Presence of Indivisibilities and Heterogeneous Firms," *Journal of Econometrics*, vol. 134, pp. 187–214, 2006.

[8] T. Iimura and T. Watanabe, "Pure strategy equilibria in finite symmetric concave games and an application to symmetric discrete Cournot games," Tokyo Metropolitan University, Tech. Rep., 2015.

[9] P. Naor, "The regulation of queue size by levying tolls," *Econometrica*, vol. 37, no. 1, 1969.

[10] R. Hassin, *Rational Queueing*. CRC press, 2016.

[11] E. Altman, T. Boulogne, R. El-Azouzi, T. Jiménez, and L. Wynter, "A survey on networking games in telecommunications," *Computers & Operations Research*, vol. 33, no. 2, pp. 286–311, 2006.

[12] E. Altman, "A Markov Game Approach for Optimal Routing into a Queueing Network," *Annals of Dynamic Games*, vol. 5, 1999.

[13] E. Altman and H. Kameda, "Equilibria for Multiclass Routing in Multi-Agent Networks," 2001.

[14] T. Boulogne, E. Altman, H. Kameda, and O. Pourtallier, "Mixed equilibrium in multiclass routing games," *IEEE Transactions on Automatic Control*, vol. 47, no. 6, 1997.

[15] B. Calvert, W. Solomon, and I. Ziedins, "Braess's Paradox in a Queueing Network with State-Dependent Routing," *Journal of Applied Probability*, vol. 34, 1997.

[16] A. Economides and J. Silvester, "A game theory approach to cooperative and non-cooperative routing problems," 1990.
[17] D. Paccagnan and M. Gairing, "In congestion games, taxes achieve optimal approximation," in *Proceedings of the 22nd ACM Conference on Economics and Computation*, 2021, pp. 743–744.
[18] L. Xia and Q.-S. Jia, "Game theoretical control of service rates in closed Jackson networks," 2013.
[19] X. Di, R. Ma, H. Liu, and X. Ban, "A link-node reformulation of ridesharing user equilibrium with network design," *Transportation Research Part B*, vol. 112, pp. 230–255, 2018.
[20] J. Ma, M. Xu, Q. Meng, and L. Cheng, "Ridesharing user equilibrium problem under OD-based surge pricing strategy," *Transportation Research Part B*, vol. 134, pp. 1–24, 2020.
[21] R. Nair and E. Miller-Hooks, "Equilibrium network design of shared-vehicle systems," *European Journal of Operational Research*, vol. 235, pp. 47–61, 2014.
[22] G. Zardini, N. Lanzetti, E. Frazzoli, and F. Dorfler, "Game theory to study interactions between mobility stakeholders," in *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, 2021.
[23] Bolch, *Queueing Networks and Markov Chains : Modeling and Performance Evaluation with Computer Science Applications*. Wiley-Interscience, 2006.
[24] B. Turan and M. Alizadeh, "Competition in electric autonomous mobility-on-demand systems," *IEEE Transactions on Control of Network Systems*, vol. 9, no. 1, pp. 295–307, 2021.
[25] L. Lei and S. Gao, "Evolutionary game analysis of ridesourcing industry between transportation network companies and passengers under new policies of ridesourcing," *IEEE Access*, vol. 6, pp. 71 918–71 931, 2018.
[26] G. Sun, G. O. Boateng, D. Ayepah-Mensah, and G. Liu, "Dynamic Pricing for Vehicle Dispatching in Mobility-as-a-Service Market via Multi-Agent Deep Reinforcement Learning," *IEEE Transactions on Vehicular Technology*, 2024.
[27] M. Sheldon, D. Paccagnan, and G. Casale, "Cournot Queueing Games With Applications to Mobility Systems," in *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, 2024, pp. 2462–2464.
[28] ——, "Addendum: Detailed proofs of theorem 6 and 9 (competitive analysis of vehicle-sharing systems with cournot queueing games)," 2024.
[29] P. Buchholz, J. Kriege, and I. Felko, *Input Modeling with Phase-Type Distributions and Markov Models: Theory and Applications*, ser. Springer Briefs In Mathematics, 2014.
[30] P. Schweitzer, "Approximate analysis of multiclass closed networks of queues," 1979.
[31] Y. Bard, "Some Extensions to Multiclass Queueing Network Analysis," 1979.
[32] M. Reiser and S. Lavenberg, "Mean-Value Analysis of Closed Multichain Queuing Networks," *Journal of the ACM*, vol. 27, no. 2, 1980.
[33] C. H. Sauer and K. M. Chandy, *Computer Systems Performance Modeling*. Prentice-Hall, Inc., 1981.
[34] C. Fricker and N. Gast, "Incentives and redistribution in homogeneous bike-sharing systems with stations of finite capacity," *Euro journal on transportation and logistics*, vol. 5, no. 3, pp. 261–291, 2016.
[35] R. Iglesias, F. Rossi, R. Zhang, and M. Pavone, "A BCMP network approach to modeling and controlling autonomous mobility-on-demand systems," *The International Journal of Robotics Research*, vol. 38, no. 2-3, pp. 357–374, 2019.
[36] G. Owen, *Game Theory*. Emerald Group Publishing, 2013.
[37] J. Nash, "Non-cooperative games," *Annals of mathematics*, pp. 286–295, 1951.
[38] D. Monderer and L. Shapley, "Potential Games," *Games and Economic Behavior*, vol. 14, 1996.
[39] M. Voorneveld and H. Norde, "A Characterization of Ordinal Potential Games," *Games and Economic Behavior*, vol. 19, no. 2, 1997.
[40] P. Hammond, "Rationality in Economics," *Rivista internazionale di scienze sociali*, vol. 105, no. 3, pp. 247–288, 1997.
[41] G. Beaujon and M. Turnquist, "A Model for Fleet Sizing and Vehicle Allocation," *Transportation Science*, vol. 25, no. 1, pp. 19–45, 1991.
[42] D. George and C. Xia, "Fleet-sizing and service availability for a vehicle rental system via closed queueing networks," *European Journal of Operational Research*, vol. 211, pp. 198–207, 2011.
[43] R. Pascual, A. Martinez, and R. Giesen, "Joint optimization of fleet size and maintenance capacity in a fork-join cyclical transportation system," *Journal of the Operational Research Society*, vol. 64, no. 7, pp. 982–994, 2013.
[44] H. Zhang, C. Sheppard, T. Lipman, and S. Moura, "Joint Fleet Sizing and Charging System Planning for Autonomous Electric Vehicles," *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 11, pp. 4725–4738, 2020.
[45] J. Barrios and J. Godier, "FLEET SIZING FOR FLEXIBLE CARSHARING SYSTEMS: A SIMULATION-BASED APPROACH," *Transportation Research Record*, vol. 2416, no. 1, pp. 1–9, 2014.
[46] L. Dowdy, D. Eager, K. Gordon, and L. Saxton, "Throughput Concavity and Response Time Convexity," *Information Processing Letters*, vol. 19, pp. 209–212, 1984.
[47] J. Shanthikumar and D. Yao, "Second-Order Properties of the Throughput of a Closed Queueing Network," *Mathematics of Operations Research*, vol. 13, no. 3, pp. 524–534, 1988.
[48] S. Roy and T. Sabarwal, "On the (Non-)Lattice Structure of the Equilibrium Set in Games with Strategic Substitutes," *Economic Theory*, vol. 37, no. 1, pp. 161–169, 2008.
[49] A. Tarski, "A lattice-theoretical fixpoint theorem and its applications," *Pacific Journal of Mathematics*, vol. 5, no. 2, pp. 285–309, 1955.
[50] W. Novshek, "Finding All n-Firm Cournot Equilibria," *International Economic Review*, vol. 25, no. 1, pp. 61–70, 1984.
[51] ——, "On the Existence of Cournot Equilibrium," *The Review of Economic Studies*, vol. 52, no. 1, pp. 85–98, 1985.
[52] N. S. Kukushkin, "Novshek's trick as a polynomial algorithm," *Discrete Applied Mathematics*, vol. 162, pp. 28–31, 2014.

Matthew Sheldon received the BS degree in Industrial and Systems Engineering in 2019 and the MSc degree in Computer Science in 2021. He is currently a PhD student in the department of Computing at Imperial College London. His research interests focus on mobility systems, rational queueing, and reinforcement learning.

Dario Paccagnan is a Senior Lecturer in the Department of Computing at Imperial College London. In 2018, Dario obtained a Ph.D. degree from the Information Technology and Electrical Engineering Department, ETH Zürich, Switzerland. He received his B.Sc. and M.Sc. in Aerospace Engineering in 2011 and 2014 from the University of Padova, Italy. In 2014, he also received his M.Sc. in Mathematical Modelling from the Technical University of Denmark; all with Honours. Dario was a visiting scholar at the University of California, Santa Barbara in 2017, and at Imperial College London in 2014. He is recipient of the SNSF fellowship for his work in Distributed Optimization and Game Design.

Giuliano Casale joined the Department of Computing at Imperial College London in 2010, where he is currently a Reader. He does research in performance engineering and cloud computing, topics on which he has published more than 150 refereed papers. He has served on the technical program committee of several conferences in the area of performance and dependability. His research work has received multiple recognitions, including best paper awards at ACM SIGMETRICS, IEEE/IFIP DSN, and IEEE INFOCOM. During 2019-2023 he was the chair of ACM SIGMETRICS. He serves on the editorial boards of ACM TOMPECS and ACM TAAS and as Editor-in-Chief of Elsevier Performance Evaluation.