Price Optimization for Round Trip Car Sharing

Christine S.M. Currie[a] Rym M'Hallah[b] Beatriz B. Oliveira[c]

[a]University of Southampton, UK
[b]King's College London, UK
[c]INESC TEC, Faculty of Engineering, University of Porto, Portugal

Abstract

Car sharing, car clubs and short-term rentals could support the transition toward net zero but their success depends on them being financially sustainable for service providers and attractive to end users. Dynamic pricing could support this by incentivizing users while balancing supply and demand. We describe the usage of a round trip car sharing fleet by a continuous time Markov chain model, which reduces to a multi-server queuing model where hire duration is assumed independent of the hourly rental price. We present analytical and simulation optimization models that allow the development of dynamic pricing strategies for round trip car sharing systems; in particular identifying the optimal hourly rental price. The analytical tractability of the queuing model enables fast optimization to maximize expected hourly revenue for either a single fare system or a system where the fare depends on the number of cars on hire, while accounting for stochasticity in customer arrival times and durations of hire. Simulation optimization is used to optimize prices where the fare depends on the time of day or hire duration depends on price. We present optimal prices for a given customer population and show how the expected revenue and car availability depend on the customer arrival rate, willingness-to-pay distribution, dependence of the hire duration on price, and size of the customer population. The results provide optimal strategies for pricing of car sharing and inform strategic managerial decisions such as whether to use time- or state-dependent pricing and optimizing the fleet size.

Keywords: car sharing; dynamic pricing; Markov chains; optimal pricing; queuing theory.

1 Introduction

Greater awareness of climate change and significant increases in fuel costs offer an opportunity to shift behavior toward shared mobility. Shared car services support the transition to a more sustainable transport system (Ampudia-Renuncio et al., 2018). They decouple ownership from use, offering the accessibility and flexibility of private car usage while reducing their cost and carbon footprint. They are however complex systems: judicious business models, demand management, and pricing optimization are key to their financial viability. Their sustainability depends on financial, ecological (Hartl et al., 2018), and technological advances (Litman, 2000; Roblek et al., 2021). Most importantly their popularity is affected by cost, distance to vehicle, savings (Krueger et al., 2016), convenience (Costain et al., 2012; Bojković et al., 2019), confidence in using the service, and safety (Maas and Attard, 2020; Rahimi et al., 2021; Jie et al., 2021). Better access to shared cars and more competitive prices would encourage the adoption of shared transport schemes (Namazu et al., 2018).

Pricing is a critical instrument in car sharing systems due to its potential to manage demand. Unlike other shared mobility modes, availability control mechanisms cannot usually be used by car sharing services due to demand spontaneity, making pricing a critical lever to optimize (Soppert et al., 2024). Current trends in car sharing research reveal a lack of techniques for pricing of car sharing journeys, in particular those related to round trips (Ferrero et al., 2018). The absence of analytic research persists (Giorgione et al., 2020) despite the sensitivity of demand to users' costs, distance to bays, and waiting times. Existing techniques fail to integrate realistic business practices and inherent complexities of car sharing, especially concerning pricing (Golalikhani et al., 2021a).

Optimal pricing of car sharing services would make them more accessible and attractive to users, encouraging them to voluntarily shift their behavior to substitute private cars with shared cars for some journeys. In turn, this behavioral shift will gradually reduce the need for private car ownership and support the transition toward net zero. We assume that a successful car sharing service that is widely used will have environmental benefits (cf. CoMoUK (2022) for examples of such benefits.). Our focus is on developing dynamic pricing methods that optimize expected revenues, making such services more commercially sustainable. Here, we describe a fast exact method for the optimal pricing of car sharing and the estimation of key performance indicators.

1.1 Problem Description

We consider a car sharing service with a fixed fleet of cars, centrally managed and offered to individual customers. Customer requests are assumed to arrive at the time of hire and customers will decide whether or not to hire the car based on the hourly hire price being charged. Arrivals follow a random process that can have a constant or time-dependent expected arrival rate. Unlike traditional car rental models where bookings are commonly made in advance, bookings in car sharing tend to arrive close to when the car is needed. A further key difference between car sharing and car rental schemes is the duration of hire, with customers using a shared car for a shorter period, e.g., for the weekly grocery shop or to travel somewhere with poor public transport. The duration is also not always known at the time of hire. Here, we assume that the hire duration is a random variable and consider two different scenarios for its probability distribution: first, that it is fixed and known; and second, that it depends on the hourly hire price.

Even though free-floating systems (where cars can be picked-up and dropped-off anywhere in a given operating area) are substantially studied in the car sharing literature, station-based systems are increasingly drawing interest (Huang et al., 2020; Chen and Liu, 2023). Round trips are the most common business model in the UK (e.g. Enterprise Car Clubs, Hiyacar, Co-Wheels), where cars are picked up and dropped off in the same location. London is the only UK location currently offering one-way hire, with only 9% of users opting exclusively for one-way hire, where cars can be dropped off at a different location (CoMoUK, 2022). Discussions with UK providers at recent workshops (Currie and M'Hallah, 2024; Currie et al., 2025) suggest that this is because the one-way hire business model is neither environmentally sustainable nor financially viable. Even though round trip systems (also referred to as two-way station-based car sharing) are widely used in practice, they have received limited attention in the literature, an exception being recent work by Ströhle et al. (2019). As a result, in this research, we focus on pricing of round trips.

Two main operating models can be used to run a round trip car sharing service: subscription plus hire fee (Becker et al., 2017) or hire fee only (Golalikhani et al., 2021b). Where a subscription is paid, there may be a desire for higher car availability and a reduced hire fee compared with the second model, which is more similar to traditional car rental. We allow for both business models here by assuming either a finite population of customers to describe the case where customers sign up to the company in advance, or an infinite population where no subscription is paid. We do not consider a mixed business model composed of both subscribed and non-subscribed customers.

1.2 Models

We develop fast, tractable models that optimize the hourly hire prices charged to car sharing customers. Three pricing strategies are considered. (i) A fixed price strategy, where hourly hire prices are independent of the number of cars on hire; (ii) a state-dependent dynamic pricing strategy, where the hourly hire price varies with the number of cars on hire; and (iii) a time-dependent pricing strategy where the hourly hire price changes with the time of day. The second strategy allows pricing to vary dynamically dependent on how busy the system is, while the third varies the prices based on forecast demand during set time intervals.

Assuming that the state of the system can be described by the number of vehicles on hire, we describe the system's evolution using a Markov chain. The structure is similar for the two prevalent business models: subscriptions or pay-as-you-go. Therefore, we use this structure to derive analytical results for both cases. When the mean hire duration is independent of the hourly hire price, the Markov chain simplifies to a queuing model when the price is either fixed at a single fare or is the state-dependent. This allows us to find exact solutions for the expected revenue using analytical models that are very fast to compute for most cases. The only exception is when the hire duration depends on the hourly hire price under the state-dependent pricing strategy and here, we need to use simulation optimization to find the optimal prices. We also use this simulation optimization to determine the optimal hire price when we assume time-dependent pricing or non-exponential service time distributions.

The queuing models and the simulation optimization allow us to estimate the steady-state expected revenue, car availability, and number of cars available for the system under different values for the hourly hire price. While the expected revenue is our key target, expected availability and number of cars available are also important. The expected number of cars available reflects the fleet's utilization, which is key to providing an economically viable service, while the expected availability reflects customers' satisfaction with the service. These indicators are particularly critical to car clubs, where customers are paying a subscription in addition to the price per hour for each hire. As the models we develop are very fast, we can estimate the optimal prices to charge using standard numerical solvers. Similarly, the models allow us to estimate the expected revenue gains in increasing the fleet size; thus, enabling strategic planning as well as dynamic pricing.

1.3 Contributions

Our key contribution is a simple yet effective method for modeling car sharing, which can be used to optimize prices in real-time. The Markov models that we present complement previous large, complex simulation models that are excellent for describing system behavior and examining operational questions, but are too computationally expensive to use for dynamic pricing. The proposed models capture the problem structure in an innovative, realistic and straightforward way, supported by a strong analytical foundation. The observation that the Markov model can be simplified to a queuing model where the hire duration is independent of price speeds up the computation significantly, allowing for a rapid evaluation of key management indicators, such as expected revenues, fleet availability, and fleet utilization, for different pricing strategies and fleet sizes. More importantly, they allow for scaling up the problems tackled.

In addition to the innovative models, we provide key managerial insights. We analyze the impact of different price response functions, finite and infinite population assumptions, and relevant pricing strategies such as state-dependent fares. The results provide optimal operational strategies and inform strategic managerial decisions for different settings, such as different levels of traffic intensity and fleet size.

1.4 Outline

The remainder of the paper is organized as follows. Section 2 reviews the literature on car sharing, with a focus on pricing. The methodology is presented in Section 3, progressing from a single fare model to multi and two fare models where a different hourly hire price can be charged for each possible number of cars on hire, allowing for very fine-grained pricing. We also introduce the simulation model that is used for price optimization for the time-dependent pricing strategy and state-dependent pricing where the hire duration is dependent on the hourly hire price. Section 4 analyzes the results and exhibits the general trends in behavior for different model assumptions, and how the models could be used for both operational and strategic decision-making. Finally, we conclude in Section 5, summarizing our findings and indicating some future directions of research.

2 Literature Review

We begin by reviewing the use of optimization approaches for car sharing systems in Section 2.1, with a focus on its use in pricing. Methods using queuing theory and Markov chains are also considered in more detail. Section 2.2 describes the use of simulation modeling of car sharing, and how simulation optimization has been used to find optimal fleet sizes, prices and station capacities. Finally, Section 2.3 discusses the main contributions of this paper and how it fits within the existing literature.

2.1 Optimization of Car Sharing

Much of the optimization research in car sharing focuses on the problem of one-way hires. Typically, it uses mathematical programming to balance supply and demand or (re)allocate cars, as Illgen and Höck (2019) highlight in their review article on vehicle relocation for car sharing.

Deterministic optimization is also used for dynamic pricing of car sharing schemes. Jorge et al. (2015) modeled the pricing of one-way car sharing trips as a mixed integer linear program and solved it approximately using an iterated local search. They used non-linear programming for dynamic pricing of one-way car sharing journeys, with the objective of balancing car availability across bays. A network of 75 hypothetical stations in Lisbon (Portugal) was used to test their approach. By better balancing the supply and demand across the network, under the assumption that demand is sensitive to price, they changed a daily deficit into a profit. The benefits to the car sharing service provider came at the expense of reduced satisfaction and higher cost for the consumer. Simulation results showed how dynamic pricing incentivizes users to balance the supply and demand at different bays and increase service providers' profit.

Origin-based pricing strategies, where the price depends on the origin and destination, emanate from practical considerations such as different parking costs at both locations. Soppert et al. (2022) proposed a mathematical optimization model and an approximate dynamic programming algorithm for a car sharing system that differentiates prices temporally and spatially. When applied to a case study in Florence, Italy, their pricing strategy increased the profits obtained by business practices. Alternatively, Huang et al. (2020) found that incentives based on drop-off location increase both profits and quality of service. Müller et al. (2023) proposed a customer-centric approximate dynamic pricing approach, based on the user's location and choice behavior, and real-time vehicle availability data.

To address the vehicle imbalance issues inherent to one-way systems, there has been a trend in the literature to integrate pricing and fleet management in car sharing. Xu et al. (2018), who focused on one-way hires using electric cars, proposed a mixed-integer non-linear and non-convex model that integrates decisions on fleet size, trip pricing, and relocation. They tested their solution approach in a case study in Singapore. Pantuso (2022) integrated pricing and relocation while considering the uncertainty in customer preferences. He developed a two-stage stochastic programming model that he evaluated on artificial instances built for the city of Milan, Italy. Focusing on autonomous electric cars, Chen and Liu (2023) introduced an integrated optimization framework that addresses both long-term charging facility deployment decisions and short-term operational decisions, such as vehicle assignment, relocation, and charging. The authors proposed a two-stage stochastic integer program and an accelerated two-phase Benders decomposition-based algorithm, which were tested and validated on data from Shanghai City.

Dynamic pricing is often addressed using Markov decision processes, although not directly in the car sharing literature. Aviv and Pazgal (2005) modeled dynamic pricing problems faced by retailers of fashion-like goods using a stylized partially observed Markov decision process framework. Their model, which accounts for demand uncertainties, was approximately solved via an active-learning heuristic pricing policy. Zhu et al. (2019) used dynamic pricing as a driver to matching electricity supply and demand in smart grids. They developed a social welfare maximization model for real-time pricing of smart home appliances based on Markov decision processes. They further divided the optimization into a subproblem for the user and a subproblem for the supplier and built a heuristic for each. For a revenue management application in the air cargo industry, Han et al. (2010) used a Markov model to study capacity allocation and accept/reject decisions. They used a bid-price control policy to determine whether to accept a booking request or hold the space for future reservations, and derived optimal solutions by maximizing a reward function of a Markov chain.

Markov chains and queuing models are emerging techniques in shared mobility systems (Braženas and Valakevičius, 2023). They are a promising tool for describing, analyzing, and optimizing the dynamic nature of these systems. In car sharing, they allow for the analytical optimization of different system features, such as fleet size or station capacity, yet their exploration in car sharing remains limited. Benjaafar et al. (2022) aimed to find the minimal fleet size that guarantees satisfaction of a given service level for a one-way car sharing system with random demand, rental duration, and car availability. They formulated the problem as a closed queuing network model with two types of queues: pick-up queues in each station and transit queues.

Queuing theory is also used to model specific components of car sharing. Guo and Kang (2022) modeled the charging process of electric shared cars as multi-server queues that are embedded in a joint framework for optimal pricing, charging, and rebalancing of a one-way car sharing system. Nakamura et al. (2022) presented a new mobility service that combines car and ride sharing. They modeled the service as a queuing model based on a Markov chain, and used the model to determine the optimal number of passengers to share a ride and guarantee a minimum level of system viability.

Banerjee et al. (2022) developed an approximation framework to optimize pricing and balance shared vehicle systems, including car sharing networks. They employed steady-state Markovian models to capture the dynamics of shared vehicle systems and developed approximation algorithms for the pricing problem under different system objectives and constraints on service-level or social welfare. The modelling of arrivals and service durations matches our assumptions, but the inclusion of one-way hires and multiple depots means that their Markov Chain is much more complex. As a result, it is only possible to solve a small case exactly and standard-sized problems approximately. By focusing only on round trip hires, and using results derived in Section 4.3, we are in contrast able to find optimal prices to exact problems even for relatively large fleet sizes.

A related body of literature deals with the pricing of reusable resources. Motivated by applications like cloud computing, Doan et al. (2020) focused on dynamic pricing of reusable resources under uncertain demand and service time. They used robust optimization to model the pricing problem and proposed deterministic approximation models and heuristic fixed-price policies. The dynamic pricing problems for car sharing and cloud computing share common ground in terms of optimizing limited reusable resources under uncertain demand. However, there are key differences that make the car sharing problem distinct. In car sharing, vehicles are a limited resource with fixed availability, usually in a much smaller scale than cloud computing, which makes resource management discrete and highly dependent on current limited stock levels. In cloud computing, resource scalability is higher compared to the physical limitations of vehicle availability in car sharing. Balseiro and Ma (2023) considered the effect of using “two-price” policies in dynamic pricing of reusable resources where there is variation in customers’ willingness to pay and usage duration. The authors proposed and analyzed stock-dependent pricing policies. Considering two prices allows for maintaining computational simplicity while providing improved performance over static pricing. They minimized regret rather than maximizing revenue. They further ignored some of the complexities we consider here such as the possibility of a finite customer population and the possibility of service times being price-dependent. It is also more difficult to see directly how their optimal pricing could be implemented in a real system and to assess the likely impacts that their pricing would have on key performance indicators such as availability and utilization.

2.2 Simulation of Car Sharing

Simulation modeling of problems related to car sharing has commonly relied on agent based modeling (ABM) to tackle demand uncertainty. Multi-agent simulation can grasp the microscopic nature and high temporal resolution of car sharing and enables the simulation of large-scale scenarios while offering a completely disaggregated representation of car sharing operations and utilization. However, most of the focus is on one-way trips, with some literature mimicking the experience and insights gained from bike sharing. For example, Ciari et al. (2013) predicted demand using an activity-based micro-simulation that included availability and price differentiation as parameters. Their simulation used 276 bay stations and 160,000 agents scattered in the Greater Zurich area, Switzerland. Ciari et al. (2014) compared free floating and round trip car sharing for the city of Berlin using an ABM. Their simulation models are computationally intensive, with 14 hours of run time for 40 replications. Giorgione et al. (2020) compared availability-based and time-based dynamic pricing for round trip car sharing from both the users' and service providers' perspective. They investigated the interaction of demand and supply using ABM, where car sharing users (from Berlin, Germany) are segmented according to their income and time utility. They found that dynamic pricing increases revenue, helps balance vehicle supply, and supports larger fleet sizes. However, their ABM, which has a large number of agents (over 280,000) and vehicles (85), was computationally expensive to run, requiring around 30 hours per scenario.

Discrete-event simulation provides an alternative to ABM for modeling car sharing. Li et al. (2021) developed a discrete-event simulation for one-way electric car sharing. Their model incorporated road congestion and travel speed, which are essential for the charging process in electric vehicles. They investigated the optimal fleet size, station capacities and pricing. Their simulation-optimization framework, which uses simultaneous perturbation stochastic approximation, returns stable results after around 10,000 replications; thus, requires a long runtime. Fanti et al. (2014) addressed the optimal sizing of a fleet of shared electric cars for a multi-location system with the objective of maximizing revenue. They modeled the problem as a discrete event system in a closed queuing network with each station having three queues of fully-charged, partially-charged and uncharged vehicles. Their model uses an approximation to reduce computation times. Zhou et al. (2023) used discrete simulation-based optimization to optimize the spatial allocation of a car sharing fleet. They used a mixed integer program within the simulation optimization algorithm and applied their model to Zipcar data from Boston for round trip car sharing. The run times of their offline simulation optimization were quite significant (10 - 60 hours).

2.3 Research Gaps

Simulation models have played a critical role in addressing the complexities and challenges of pricing in car sharing, particularly in the context of uncertainties in demand and usage. They have offered valuable insights into fleet dynamics and customer behavior. Incorporating optimization with simulation has supported informed decision-making but at the expense of long computation times. Existing simulation optimization models have also put a strong focus on one-way car sharing, while the predominant business model in the UK is to offer only round trip hires.

Analytical models that can quickly support optimal price and fleet sizing decisions for car sharing are still missing, although some approximate models exist for pricing of one-way hires. We bridge this gap by proposing fast models that account for uncertainty in arrivals and hire durations and enable pricing to vary with the number of cars on hire, thus facilitating dynamic pricing strategies. Where price affects the duration of hire or where hourly hire prices are assumed to be time-dependent, simulation models are more effective. Following the design of the Markov models also enables the simulation models to run quickly enabling relatively fast simulation optimization routines. The paper demonstrates general results about car sharing and the impact of different system factors on the optimal pricing and the optimal revenues but, most importantly, it enables a fast calculation of optimal prices for relatively complex pricing strategies, and the estimation of expected values of key

Table 1: Scenarios considered in the modeling

Customer population	Pricing Strategies	Effect of Price
Finite: subscription model	Flat fare structure with one hourly hire price	Purchase probability: linear or logistic price response models
Infinite: pay-as-you-go model	Full and reduced state-dependent pricing	Hire duration: linear or no dependence
	Time-dependent pricing	

performance indicators.

3 Methodology

We use continuous time Markov chains to describe how individuals enter the system, interact with it and then leave having either hired and then returned a car or having left because the price is too high. In what follows we describe the simplest model, where the price is fixed at a single fare, before we consider pricing strategies in which the hourly hire price depends on how many cars are available at the time of hire or where the hourly hire price depends on the time of hire. Table 1 describes the different scenarios that we consider.

We assume that customers either accept a hire price with a probability dependent on their price-response function, which we assume is homogeneous across the population, or alternatively have a stochastic hire duration with a mean value $1/\mu(r)$ that depends on r, the hourly hire price. Typically, the price-response functions would be estimated from historical booking data or surveys. We introduce two widely used price response functions for probability of purchase in the next section and our assumed model for the effect of hourly hire price on hire duration.

3.1 Price-Response Functions

Customers are assumed to react to prices in one of two ways: they either decide whether to purchase the car hire or determine their hire duration according to price. We begin by considering the first of these, which is seen more often in the literature.

We consider two common price-response functions: a linear price-response function, which corresponds to a uniform distribution of reservation prices; and a logit price-response function, which assumes a logistic distribution for the reservation prices (see e.g. Phillips (2005) for an introduction to these functions). The probability that a person drawn randomly from the population will pay a price r is the probability that their reservation price is greater than or equal to r. The distribution of reservation prices is assumed to have a cumulative distribution function $F(r)$; therefore the probability of purchase at price r, which we denote by $\omega(r)$, is equal to $1 - F(r)$. Below we give expressions for $\omega(r)$ for the two price response functions we consider here.

1. **Linear price-response function**: the demand varies linearly with price. In this case, reservation prices are uniformly distributed between a and b, $U(a, b)$ and we can write the probability of purchase at price r as
$$\omega(r) = \frac{b - r}{b - a}. \tag{1}$$

2. **Logit price-response function**: close to the market price, small shifts in price result in large changes in demand; further away from the market price, demand changes more slowly with price. The probability of purchase at rate r is

$$\omega(r) = (1 + e^{-m/s})\frac{e^{-(r-m)/s}}{1 + e^{-(r-m)/s}}, \tag{2}$$

where m is the mean and mode of the distribution of reservation prices for the population and s is set such that $\pi s^2/3$ is equal to the variance of the reservation process. The first term normalizes the probabilities so that the probability of purchasing at $r = 0$ is equal to 1. The probability of purchase tends to zero for large values of r and equals $0.5(1 + e^{-m/s})$ when $r = m$.

Less is known about how the price of hire might affect its duration. Here, we make the reasonable assumption that the higher the hourly hire price, the shorter the duration of hire. We experiment with a simple model where there is a linear dependence between the expected hire duration and the hourly hire price such that

$$\mu(r) = \frac{\mu_0}{1 - cr}, \tag{3}$$

where $1/\mu_0$ is the expected duration when the price is set to its minimum value, and c describes the price-dependence of the duration on r. We assume that c ranges between 0 and $1/\max\{r\}$ to avoid negative $\mu(r)$.

3.2 Single Fare Pricing

We can view a car sharing service with a fleet size of n cars and a constant hourly hire price of r as a queuing model with a state $X(t)$ representing the number of vehicles on hire. We assume that car hire requests arrive following a Poisson process with an hourly rate $\lambda > 0$, such that on average λ customers will consider hiring a car during a one hour period. When all n cars are on hire, demand is turned away. Assuming arrivals follow a Poisson distribution is standard, as it is a good approximation for many systems, while enabling the application of known results regarding steady-state behavior. We consider **two** alternative models for how customers react to the hourly hire price. In the **first** case, a customer who enters the system either accepts the offered hourly hire price r with probability $\omega(r)$ or rejects r with probability $1 - \omega(r)$. This acts to *thin* the arrivals yielding an effective arrival rate of $\lambda\omega(r)$, which we write as $\lambda(r)$ in what follows. In the **second** case, a customer's hire duration follows an exponential distribution with mean $1/\mu(r)$, as shown in Equation 3. The assumption that service durations follow an exponential distribution allows us to apply standard results in the analytical modeling that follows. We assume that the hourly hire price must be within an allowed range. This does not change the analytical model significantly but fits with practical situations where the price is bounded above and below.

There are **two** options for the customer population: infinite or finite. These reflect different business models in the car sharing industry. The **first** option represents a system open to any user on a pay-as-you-go basis, without requiring previous subscription, thus being represented by an infinite population. The **second** option represents a system that operates on a subscription basis, where potential users are registered with the company and can thus be represented by a finite population. We address each of these **two** options, presenting models to describe the evolution of

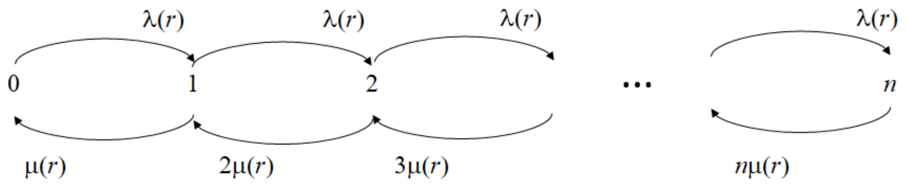

Figure 1: State transition diagram for a single price and an infinite customer population

the state of the system and providing expressions for key statistics such as the expected revenue, the probability of finding a car when the system is in a steady-state, and the expected number of cars available. A stochastic system is said to be in a steady-state when the probability of finding the system in a given state is independent of time.

3.2.1 Infinite Customer Population

Under the assumption of an infinite customer population, the rate of arrivals into the system is unaffected by the number of vehicles already on hire. The state transition diagram is given in Figure 1. For this model, the steady-state probability that there are i cars on hire is

$$\pi_i = \pi_0 \rho(r)^i / i!, \tag{4}$$

where $\rho(r) = \lambda(r)/\mu$ in the case that the customer's probability of purchase depends on the hourly hire price and $\rho(r) = \lambda/\mu(r)$ when only the customer's hire duration depends on the hourly hire price. In this case customers always accept an offered price but reduce their hire duration when prices are high. Where the linear price response function is used for the probability of purchase $\omega(r)$, we can see that $\rho(r)$ has the same linear dependence on r in both cases. The probability of there being no cars on hire is given by

$$\pi_0 = \left[\sum_{i=0}^{n} \rho(r)^i / i! \right]^{-1}. \tag{5}$$

Using these steady state probabilities, we can write an expression for the expected hourly revenue in the steady-state as

$$E[R] = r \sum_{i=1}^{n} i \pi_i(r) = r \pi_0 \sum_{i=1}^{n} \frac{\rho(r)^i}{(i-1)!}. \tag{6}$$

For simplicity, we will refer to the expected hourly revenue as expected revenue throughout this paper.

Substituting the expression for π_0 in Equation 5 into the sum in Equation 6 and rewriting, optimizing the expected revenue with respect to r can then be written as

$$\max_r \left\{ r \rho(r) \left(1 - \frac{\rho(r)^n / n!}{\sum_{i=0}^{n} \rho(r)^i / i!} \right) \right\}. \tag{7}$$

This expected revenue can be optimized for different expressions for $\omega(r)$ and different values for n.

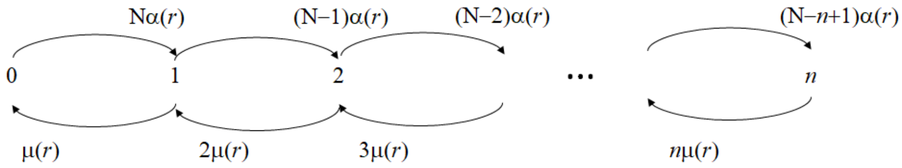

Figure 2: State transition diagram for a single price and a finite customer population of size N.

3.2.2 Finite Population

With a finite number N of car club members, the arrival rate into the system changes with the number of cars on hire. The queuing model can be represented as in Figure 2, where $\alpha(r)$ is the arrival rate of an individual into the system and is approximated by $1/N$ of $\lambda(r)$, the arrival rate used in the infinite population system, such that when there are N people wishing to hire a car, the system arrival rate is equal to $\lambda(r)$. The steady-state probability that there are i cars available is

$$\pi_i = \pi_0 \frac{\rho(r)^i N!}{(N-i)!i!}, \tag{8}$$

where $\rho(r) = \alpha\omega(r)/\mu$ or $\rho(r) = \alpha/\mu(r)$, as discussed in the previous section, and

$$\pi_0 = \left[\sum_{i=0}^{n} \frac{\rho(r)^i N!}{(N-i)!i!}\right]^{-1}. \tag{9}$$

The expected revenue is then equal to

$$E[R] = r\sum_{i=1}^{n} i\pi_i = r\pi_0 \sum_{i=1}^{n} i\frac{\rho(r)^i N!}{(N-i)!i!}. \tag{10}$$

We can see that the expected revenue for the finite population case will be bounded above by the expected revenue for the infinite population. For $N \gg n$, $\frac{N!}{(N-i)!} \to 1$ and π_0 and $E(R)$ tend to the expressions in Equations 6 and 7, respectively. The results section includes comparisons of results for finite and infinite customer populations and a description of the effect of N on the behavior of the system.

3.3 Multi Fare Pricing

We now assume that the hourly hire price depends on the state of the system when the hire begins. A person investigating hiring a car when there are i cars available will be offered an hourly hire price of r_i for the whole duration of their hire. We consider the two cases where the hire duration is dependent and independent of r respectively, showing that in the second case we can simplify the complex Markov chain to a queuing model.

As with the fixed rate, the arrival rate of customers checking the price and availability of hire cars is constant at λ for an infinite population and α per person for a finite population, but is multiplied by the probability of paying a particular fare $\omega(r_i)$, and a function of N and i for case of a finite customer population. For convenience, we write the final arrival rate as λ_i when the system is in a state where there are i cars on hire and the final hire return rate as μ_i for a customer who is paying an hourly hire price of r_i. Expressions for λ_i are given in Section 3.3.1 for the case of an infinite customer population and in Section 3.3.2 for the case of a finite customer population.

We now assume a state space $\mathbf{X}(t) = (X_0(t), X_1(t), X_{n-1}(t))$, where $X_i(t)$ is the number of customers paying r_i at time t. The expected revenue equals the sum over all feasible states of the probability of being in that state multiplied by the revenue being received. We can write a general expression for the steady-state expected revenue

$$E[R] = \sum_{S_n} \pi_{\mathbf{x}} \sum_{j=0}^{n-1} x_j r_j, \tag{11}$$

where S_n is the set of feasible states when there are n cars in the fleet, $\mathbf{x} = (x_0, x_1, \ldots, x_{n-1})$ describes the value of the state-space variable $\mathbf{X}$, and $\pi_{\mathbf{x}}$ is the steady-state probability of being in a particular state $\mathbf{x}$. For general $n \in \mathbb{N}$, we let $S_n = \{\mathbf{x} : x_i \in \mathbb{N}, x_i \in \{0, \ldots, i+1\}, x_0 + \ldots + x_i \leq i+1 \text{ for } i \leq n-1\}$. Transitions between states can also be generalized. Consider state $\mathbf{x} \in S_n$ such that $x_0 + \ldots + x_{n-1} = k \in \{0, \ldots, n\}$.

- When $k \neq n$, the system can move from state $\mathbf{x}$ to a new state $\mathbf{x}'$ via the hire of a car, which happens at a rate λ_k such that $\mathbf{x}' = \mathbf{x}$ except for x'_k which becomes $x_k + 1$.
- When $k \neq 0$, the system can move from state $\mathbf{x}$ to a new state $\mathbf{x}'$ via the return of a hired car. This transition can be to one of at most k states (as many states as there are non-zero x_i entries in state $\mathbf{x}$). For each $x_i \in \mathbf{x} : x_i \neq 0$, the system moves to state $\mathbf{x}'$ at a rate $\mu^{(k)} = x_i \mu(r_i)$ resulting in $\mathbf{x}' = \mathbf{x}$ except for x'_i which becomes $x_i - 1$.

We consider a small example where $n = 3$ to help explain the structure of this model. Figure 3 gives the Markov chain describing transitions between the different states where we account for different fares being paid. This shows the complexity of the Markov chain model.

The number of feasible states increases at a greater than linear rate. As the number of feasible states increases, the computation time of steady-state statistics becomes significantly longer. Fortunately, we can show that estimates of the expected revenue, availability, and cars available can be obtained from a linked queuing model when the hire duration is independent of the hourly hire price such that $\mu(r) = \mu$.

With a constant service rate μ, it is possible to use the structure of the problem to simplify the calculation of the expected revenue. Essentially, we show below that we can use a one-dimensional state-space with transitions that follow a much simpler queuing model to describe the problem and we can use results calculated with this model to find the expected revenue and other important performance indicators.

Define a new state space $Z(t) \in (0, 1, \ldots, n)$ to be a one-dimensional state indicating the number of cars on hire at time t such that $Z(t) = \sum_{i=0}^{n-1} X_i(t)$. Transitions from $Z(t) = m$ to $Z(t) = m+1$, corresponding to a car being hired, happen at a rate λ_m, while transitions from $Z(t) = m$ to $Z(t) = m-1$, which correspond to a car being returned, happen at a rate of $m\mu$, $m = 1, \ldots, n$. The state $Z(t) = m$ combines all of the states in which $\sum_{i=0}^{n-1} X_i(t) = m$, and each of the original states has the same transition rate to the set of states where the number of cars on hire is $m+1$ or $m-1$.

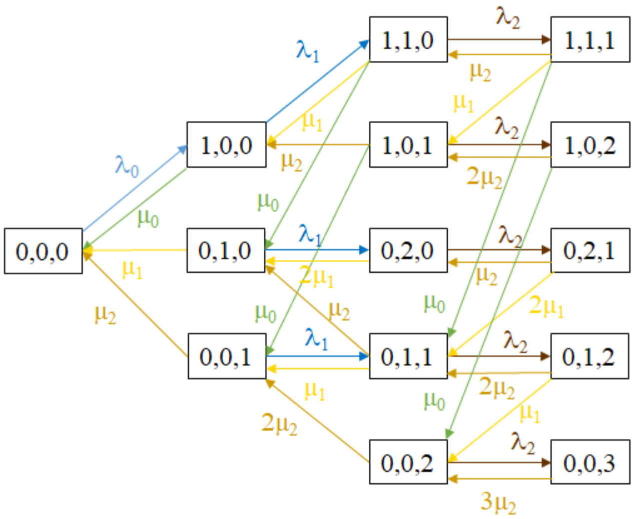

Figure 3: Transition diagram for state-dependent pricing and a fleet size n of three cars. Each state is represented by three values, the number of cars on hire paying r_0, r_1 and r_2. For ease of display we write $\lambda(r_i)$ and $\mu(r_i)$ as λ_i and μ_i respectively.

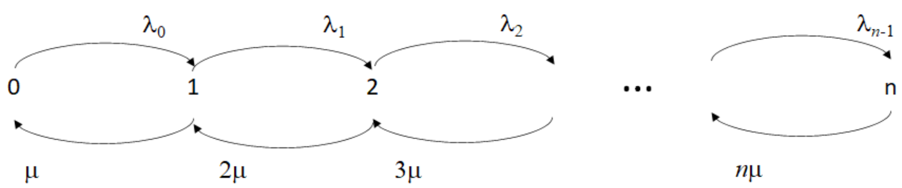

Figure 4: Queuing theory representation of the state-dependent pricing case for an infinite customer population.

Having established that we can describe the number of cars on hire using a queuing theory model, as shown in Figure 4, we determine the steady-state probabilities π_i that the system is in a state where i cars are on hire as

$$\pi_i = \frac{\lambda_{i-1}}{i\mu}\pi_{i-1},\ i = 1, \ldots, n \tag{12}$$

and

$$\pi_i = \frac{\prod_{j=1}^{i-1}\lambda_{j-1}}{i!\mu^i}\pi_0,\ i = 1, \ldots, n. \tag{13}$$

The probability that a customer arrives and finds a car available is the probability that the system is not in state n,

$$p = P[\text{Car available}] = 1 - \pi_n. \tag{14}$$

The expected number of cars on hire is

$$E[\text{Number of cars on hire}] = \sum_{i=1}^{n} i\pi_i; \tag{15}$$

thus, the expected number of cars available is

$$v = E[\text{Number of cars available}] = \sum_{i=1}^{n}(n-i)\pi_i. \tag{16}$$

It is less obvious that this model can be used to find the expected revenue because of the complication that when there are i cars on hire, $i > 0$, it is impossible to distinguish between states $\mathbf{x}^{(1)} = (x_0^{(1)}, x_1^{(1)}, \ldots, x_{n-1}^{(1)})$ and $\mathbf{x}^{(2)} = (x_0^{(2)}, x_1^{(2)}, \ldots, x_{n-1}^{(2)})$, where $\sum_{j=0}^{n} x_j^{(1)} = \sum_{j=0}^{n} x_j^{(2)} = i$. While each state has the same number of cars on hire, they may produce a different hourly revenue if $x_j^{(1)}$ is not equal to $x_j^{(2)}$ for some $j = 1, \ldots, n-1$. We show below that the structure of the problem is such that when the average hire duration $1/\mu$ is constant, this is not important and we can just consider the total number of cars on hire in determining the expected revenue in the steady-state.

Anyone who arrives when there are i cars in the system will pay an hourly rate of r_i. The expected income per hour from the system is therefore equal to

$$E[\text{Revenue}] = \sum_{i=0}^{n-1} r_i\pi_i\lambda_i(1/\mu), \tag{17}$$

the sum over all possible numbers of cars on hire of the hourly hire price r_i, multiplied by the proportion of time that the system is in a state where i cars are on hire π_i, multiplied by the number of people who start a hire when the price is r_i, λ_i, multiplied by the average duration of the hire $1/\mu$. From Equation 12, this can be written as

$$E[\text{Revenue}] = \sum_{i=0}^{n-1}(i+1)r_i\pi_{i+1}. \tag{18}$$

Reparameterizing the sum, this is equivalent to

$$E[\text{Revenue}] = \sum_{j=1}^{n} jr_{j-1}\pi_j, \tag{19}$$

which matches the expression you would obtain by assuming that all cars in state j pay a price r_{j-1}. The fact that this is true means that we can quickly evaluate the expected steady-state revenue using a simple queuing model rather than a Markov chain with a very large number of allowed states.

Such a simplification is not valid if the hire price is assumed to affect the duration of hire. In this case, the computation time and memory requirements of the Markov model are prohibitively large, and we use simulation to describe transitions between different states in the Markov chain model, implementing simulation optimization to find the optimal prices to charge. The simulation model is described in Section 3.5.

3.3.1 Infinite Customer Population

For multiple fares, under the assumptions of an infinite customer population and a hire duration that is independent of r, the arrival rate is state-dependent and can be written as

$$\lambda_i = \lambda\omega(r_i). \tag{20}$$

3.3.2 Finite Customer Population

Where the customer population is finite, under the assumption of a hire duration that is independent of r, we have an additional multiplier in the expression in λ_i such that

$$\lambda_i = \alpha(N-i)\omega(r_i). \tag{21}$$

To obtain similar levels of traffic for the finite population and infinite population cases, we set α to be λ/N. While $1/\alpha$ is the expected time between two requests of an individual in the system, $1/\lambda$ is the expected time between two requests from the whole population in the system, as discussed previously.

3.4 Two Fare Pricing

In practice, a car sharing service may wish to offer fewer than n prices to reduce the complexity of the offering to the customer. Implementing a full state-dependent pricing strategy results in frequently changing prices, which can be challenging to implement and may cause customers to become frustrated with the system. Consequently, a pricing strategy with fewer price points is likely to be an attractive alternative. We consider here a special case of how the problem can be adapted if only two prices are offered, acknowledging that a provider may decide to use more than two, but fewer than n prices.

We assume that the provider offers only two fares: r_0 where the number of cars on hire is less than a threshold n' and r_1 where the number of available cars is greater than or equal to n'. For the case where the hire duration is independent of the hourly hire price, Equations 20 and 21 can then be used to set the λ_i, which can be input into Equation 17 to determine the expected hourly revenue in the steady-state. The expected revenue can be written as

$$E[R] = \sum_{i=1}^{n'} r_0\pi_i + \sum_{i=n'+1}^{n} r_1\pi_i, \tag{22}$$

where the π_i can be found using Equation 12. The optimization routine must now select the optimal value for n', $0 < n' < n$, as well as the optimal prices r_0^*, r_1^* to maximize revenue. Where the hire duration depends on the price, we use the simulation model to describe the behavior of the system.

3.5 Simulation Model

Algorithm 1 summarizes the main steps of one replication of a discrete event simulation of the car sharing system. Its inputs are the number n of shared cars, the arrival rate λ, the hire return rate μ, the price-dependent willingness-to-pay distribution $\omega(\cdot)$, the coefficient c describing the price dependence of the hire duration, and the price vector $\mathbf{r}$. Its output is an estimate $\hat{R}$ of the expected revenue $E[R]$.

Lines 1-6 initialize the car sharing system. Line 1 fixes the replication length T to the length of the simulation period. Line 2 sets the simulation clock t_{Now} to 0. Line 3 sets the list $\mathcal{L}$ of simulation events to the empty set. Lines 4 and 5 initialize the number of cars on hire η and the revenue $\hat{R}$ to 0 as the system is empty at the beginning of the simulation. Finally, Line 6 schedules the first event at time $t_{Now} = 0$. This event is $\{(1, 0, 1)\}$. It corresponds to the first car rental request. Events are defined via a triplet (ℓ_1, ℓ_2, ℓ_3), where $\ell_1 \in \mathbb{N}^*$ denotes the request number, $\ell_2 \geq 0$ the time it occurs, and $\ell_3 \in \{1, 2\}$ its type with $\ell_3 = 1$ if the event is a car hire request and $\ell_3 = 2$ if the event is the end of a car hire.

Lines 7-27 constitute the iterative step, which stops when the simulation clock reaches or exceeds the replication length T. Line 8 checks whether the event defined by the triplet (ℓ_1, ℓ_2, ℓ_3) corresponds to an arrival. When this is the case, Line 9 checks the availability of a free car for hire by comparing η, the number of cars on use, to the fleet size n. When $\eta < n$ cars are in use, Line 10 checks the willingness of the customer to pay the fare price $r(\eta)$ for the requested journey. This is achieved by generating a random number, labeled $Willingness(\ell_1)$, from the Uniform[0,1] and comparing it to $\omega(r_\eta)$. When client ℓ_1 accepts the price offered, s/he hires the car. Line 11 generates a hire duration d_{ℓ_1} from an exponential distribution with parameter $\mu(r)$ (Equation 3). For experiments testing the impact of using an alternative distribution for hire duration, d_{ℓ_1} is generated from a gamma distribution. Line 12 calculates the end of the hire f_{ℓ_1} as the sum of t_{Now} and d_{ℓ_1}. Line 13 defines the event $(\ell_1, f_{\ell_1}, 2)$ as the end of the hire event and appends it to the set $\mathcal{L}$ of events. Line 14 updates the estimate $\hat{R}$ of the revenue by incrementing $\hat{R}$ by the product of the hire duration d_{ℓ_1} and of the accepted price r_η. Finally, this hire increases the number of cars on hire; thus, Line 16 increments η by 1. If the client is not willing to pay the fare price r_η (Line 17) or there is no car available for hire (Line 19), the algorithm disposes of the request entity.

Regardless of the success of the hire request, Line 20 plans the next arrival of a car hire request (i.e., request $\ell + 1$) by randomly generating an inter-arrival time δ according to an exponential distribution of parameter λ. Line 21 calculates a, the time of the next request arrival as the sum of this inter-arrival time δ and t_{Now}. Subsequently, Line 22 inserts this new arrival event $(\ell_1 + 1, a, 1)$ into the event list $\mathcal{L}$. Because the events of the simulation happen in ascending chronological order, Line 23 sorts this list.

If on the other hand, the event defined by the triplet (ℓ_1, ℓ_2, ℓ_3) corresponds to the end of a car hire, Line 25 decreases the number η of cars on hire by 1 and disposes of the request entity.

Lines 26 and 27 set up the parameters of the next iteration. Line 26 retrieves the first occurring event from the sorted list $\mathcal{L}$ of events and defines it as the new (ℓ_1, ℓ_2, ℓ_3) event. Line 27 advances the simulation clock, setting t_{Now} to the time ℓ_2 of this new event.

Algorithm 1: Pseudo code of a discrete event simulation of the car sharing system

Data: $n \in \mathbb{N}^*, \lambda, \mu, \omega(r), c, \mathbf{r}$ `// fleet size, arrival rate, hire rate, willingness-to-pay distribution, price-dependence of hire duration, and the price vector`
Result: R `// revenue`
`/* Initialization. */`
1 $T \leftarrow 100$ `// replication length`
2 $t_{Now} \leftarrow 0$ `// current simulation time`
3 $\mathcal{L} \leftarrow \emptyset$ `// Initializing the list of events`
4 $\eta \leftarrow 0$ `// number of cars on hire`
5 $R \leftarrow 0$ `// revenue`
6 $(\ell_1, \ell_2, \ell_3) \leftarrow \{(1, 0, 1)$ `// first event`
`/* Iterative step continues while simulation did not end. */`
7 **while** $t_{Now} < T$ **do**
    `/* Event is arrival of journey request` ℓ_1 `*/`
8   **if** $\ell_3 = 1$ **then**
        `/* A car is available */`
9   **if** $\eta < n$ **then**
            `/*` $Willingness(\ell_1) \sim U(0,1)$ `*/`
10   **if** $Willingness(\ell_1) > \omega(r_\eta)$ **then**
11   Generate hire duration d_{ℓ_1}
12   End of hire $f_{\ell_1} \leftarrow t_{Now} + d_{\ell_1}$
13                 $\mathcal{L} \leftarrow \mathcal{L} \cup \{(\ell_1, f_{\ell_1}, 2)\}$
14                 $R \leftarrow R + d_{\ell_1} * r_\eta$
15                 $\eta \leftarrow \eta + 1$
16   **else**
17   Dispose of ℓ_1
18   **else**
19   Dispose of ℓ_1
        `/* Schedule the next arrival. */`
20   Generate inter-arrival time δ between car hire request ℓ_1 and $\ell_1 + 1$
21         $a \leftarrow t_{Now} + \delta$ `// Calculate arrival time` a `of` $\ell_1 + 1$
22         $\mathcal{L} \leftarrow \mathcal{L} \cup \{(\ell_1 + 1, a, 1)\}$; `// Appending next arrival event to` $\mathcal{L}$
23   Sort $\mathcal{L}$;
24   **else**
        `/* Event is end of hire of` ℓ_1 `*/`
25         $\eta \leftarrow \eta - 1$ `// Update number of cars on hire`
26   Dispose of ℓ_1
27   Remove first event (ℓ_1, ℓ_2, ℓ_3) of $\mathcal{L}$
28     $t_{Now} \leftarrow \ell_2$

In the case where the probability of purchase is independent of the price being charged, Lines 10 and 16-17 are removed. When the price-response mechanism is reflected by the hire duration, Line 11 generates the hire duration using an exponential distribution whose parameter is $\frac{\mu}{1-cr(\eta)}$. When it is independent of the hire duration, c is set to 0. Finally, if the population size is finite, Line 20 generates the inter-arrival time using an exponential distribution of parameter $\frac{\lambda(N-\eta)}{N}$.

As described, Algorithm 1 is fed with different fares that depend on the number of cars on hire. It can be adapted to the single fare case by setting all the entries of $\mathbf{r}$ equal to a single fare r. For the two fare case, the first n' entries of $\mathbf{r}$ are set to r_0 and the last $n - n'$ entries to r_1.

Inter-arrival times, hire durations, and willingness-to-pay are generated using independent random streams. When validating the results of the simulation model against the queuing theory results,

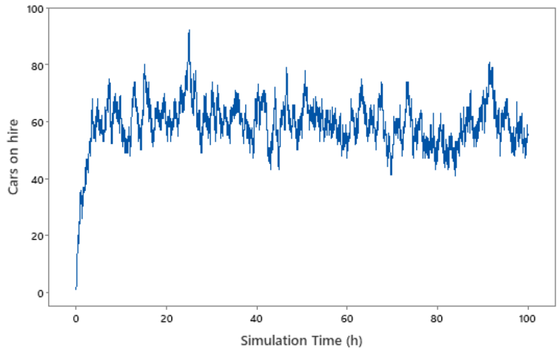

Figure 5: Cars on hire over time (single fare, $\lambda = 200$, $\mu = 1$)

we start collecting the simulation statistics after the system reaches a steady state. For example, for the simulation of Figure 5, statistics are only collected after 3 hours of simulation time. In addition, each replication is long enough to capture enough variability. For instance, the simulation of Figure 5 is run for $T = 100$. However, when mimicking a real life system, we run the simulation without a warm up period, starting with all n cars available, and for $T = 24$ hours.

Algorithm 1 provides a single observation of the statistics of interest, such as the mean revenue R. We use the average of a sample of observed statistics over n_{rep} replications as the point estimate of R. This is the output that we optimize using the *fmincon* function of Matlab or OptQuest within Arena, where constraints limit the fares to being within a given range.

4 Results

This section presents the main computational results and corresponding insights. The objective of the computational investigation is fourfold: (i) to compare prices and key performance indicators for finite and infinite customer populations and assess the effect of population size N on the behavior of the system; (ii) to show how the optimal prices and expected steady-state revenue vary with the size of the customer population N; (iii) to assess the impact of single, multi and two fare pricing on key performance indicators and (iv) to determine the effectiveness of time-dependent pricing. We present results for a range of performance measures, reflecting the competing objectives of car share systems. These are introduced in Section 4.1 before we discuss the results for single fare pricing in Section 4.2, multi fare pricing in Section 4.3, and for the special case of two fare pricing in Section 4.4. Two fare pricing can be a realistic and relevant trade-off for practical implementation. We also analyze the sensitivity of revenue to the fleet size in Section 4.5. Finally, we discuss the results for time-dependent demand in Section 4.6. We ran simulation optimization experiments using gamma distributions with different shape parameters to describe service duration. These suggested that the optimal prices obtained are almost indistinguishable from those for the exponential distribution and as a result we only include these results in the supplementary material. Throughout this section, the computational results are presented rounded to two decimal places.

4.1 Performance Measures

A car-sharing service can only exist if it has high utilization and high availability, whilst being economically attractive to service providers. As a result, our main objective is the optimal steady-state expected revenue, given by Equations 6 and 10 for the fixed price infinite and finite population cases, respectively; Equation 19 for state-dependent prices; and Equation 22 for the case of two prices. We report two additional key performance indicators in the steady state: the expected availability and the expected number of cars available. The expected availability, defined by Equation 14, is the probability of a car being available and reflects users' perceptions of service quality. The expected number of cars available, given by Equation 16, reflects the fleet utilization and can be thought of as a service provider's performance measure.

When measuring the impact of how busy the system is, we use $\lambda/n\mu$, the unconstrained traffic intensity. This would be the traffic intensity of the system if customers always accepted the price being offered and the hire duration was independent of price.

Where acceptance of hire depends on the hourly hire price we present results for the linear price response function, but show for each case (single fare, multi fare, two fare) that the results have a similar pattern for the logistic price response function. When we only display results for a linear price response function in the main text, the corresponding results for the logistic function are provided in the supplementary material.

4.2 Single Fare Pricing

We consider initially how the expected revenue is affected by price for the different price response functions. Here, and in the rest of the numerical results, we limit the prices to the range $[0, 1]$. Figure 6 assumes $\mu = 1$ and $n = 100$ and presents results for a linear price response function on the left (Figure 6(a)) and a logistic price response function on the right (Figure 6(b)). The top chart in Figure 6(a) shows how the optimal fare moves from 0.50 toward higher fares approaching 1.00 as demand increases. At lower unconstrained traffic intensities, the expected revenue curve is close to being symmetric about its maximum at $r = 0.50$. In this case, the fleet size is large relative to demand and we essentially have no limits imposed by it being finite. This result reproduces standard economic theory for a problem with unlimited supply and a linear price-response function, (e.g. see Phillips (2005)). As demand rates increase, the expected revenue curve loses its symmetric shape and becomes right skewed with a higher expected revenue and the optimal price r^* tends toward 1.00. A similar relationship can be seen in the top chart in Figure 6(b) for the logistic price response function. The optimal fare increases as the unconstrained traffic intensity increases, and the curve becomes more skewed.

As the charts in Figure 6 illustrate for the linear and logistic cases, not only does the unconstrained traffic intensity impact the optimal fare and expected revenue, but it also affects the expected availability and the expected number of cars available. The expected availability is highest when the demand is low and decreases as the arrival rate increases, while the expected availability increases as the price r increases. The price that optimizes the expected availability is different from the price that maximizes the expected revenue.

Similarly, the expected number of cars available decreases as the demand increases, reaching n when $r = 1$. The inflection point of the expected number of cars available occurs at a price that differs from r^* and from the price that optimizes the expected availability.

Figure 7 illustrates the dependence of the optimal price r^*, optimal expected revenue, corresponding expected availability and number of cars available on fleet size and the unconstrained traffic intensity for a fleet size of n cars for n up to 100 and a linear price response function. The figure shows that the higher the unconstrained traffic intensity, the higher the optimal price r^*. When the unconstrained traffic intensity equals 2, the price decreases as the number of cars increases, but this trend reverses for higher values. The general behavior of the curves is similar for a linear and a logistic price response function (cf. the supplementary material).

We can also see in Figure 7 that the expected revenue per car increases as the fleet size n increases, with the largest marginal increase obtained for values of n up to around $n = 20$, as suggested by the inflection points of the curves. Evidently, the larger the unconstrained traffic intensity, the larger the expected revenue, but with the largest marginal increment observed for an increase in the unconstrained traffic intensity from 2 to 4. Figure 7 also indicates that the expected availability increases as the fleet size increases and as traffic intensity decreases.

As expected, the ratio of the expected number of cars available to fleet size decreases as the unconstrained traffic intensity increases and as fleet size increases, with the largest marginal decrease observed for n around 10. This measure can also be thought of as the expected proportion of the fleet that is available.

Figure 8 reports the same information as Figure 7 but for a finite population of size N and a fleet size $n = 100$. We can see that for $N > 2000$ there is little variation from the results obtained for an infinite population case. For lower values of N, we see an increase in the optimal price and expected revenue as N increases and a decrease in the probability of there being a car available and the expected number of cars on hire.

To validate the theoretical results, we use a simulation model that reproduces results for the case where the probability of purchase is dependent on price for the single fare case, with similar results (see Figure 21 in the supplementary material). Point and 95% confidence interval estimates of the mean paired difference between the simulated and the theoretical revenues are 0.06 and $(0.03, 0.08)$, respectively.

Figure 9 displays the results of an experiment with a linear price-dependence on the hire duration compared with a linear price-dependence on the probability of accepting a hire price. If the hire duration is price-dependent with $c = 1$, but we assume everyone accepts the hourly hire price, we end up with an identical model to the case where we assume the duration is independent of time but the probability that someone accepts the hire is dependent on price. This is clearly illustrated by the coinciding blue and green curves.

Figure 10 shows the effect of the hire duration price-dependence coefficient c on the expected revenue. As c decreases, and the hire duration is less affected by the hourly hire price, the optimal price and the expected revenue both increase.

4.3 Multi Fare Pricing

This section presents results for the full state-dependent pricing. We describe the trend in price between the different states and give plots for the key performance indicators as for the single fare case.

Figure 11 shows how the pricing varies among states for different unconstrained traffic intensities,

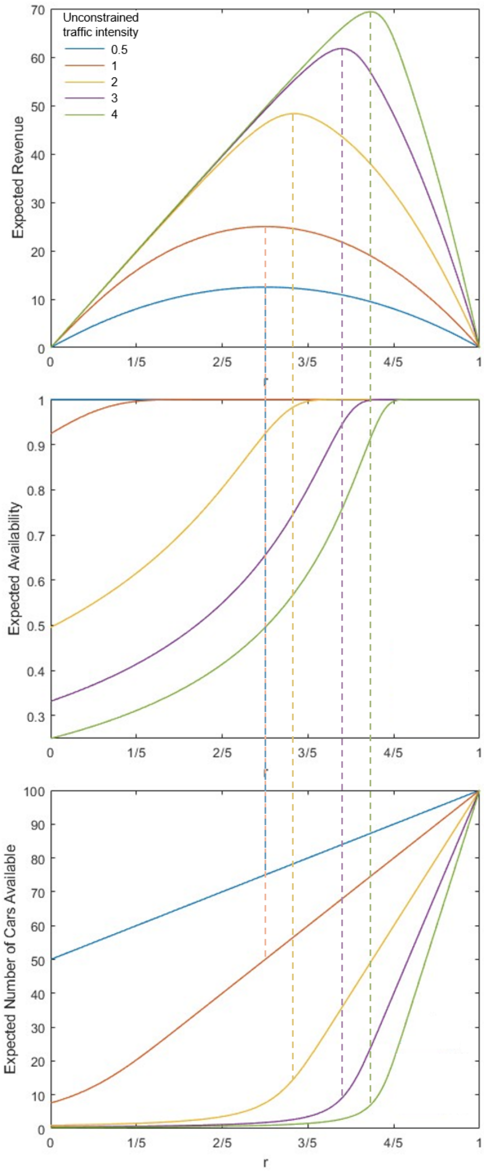

(a) Uniform

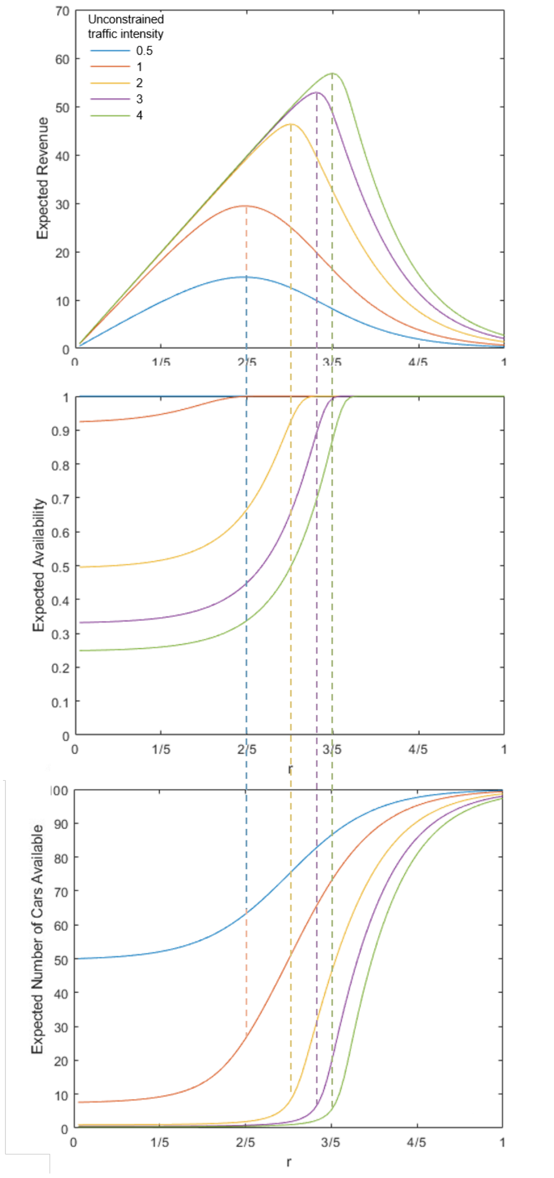

(b) Logistic

Figure 6: Single fare pricing: expected revenue, availability and number of cars available for different unconstrained traffic intensities $\frac{\lambda}{n\mu}$ where $\mu = 1$ and fleet size $n = 100$ for (a) linear price response function, $a = 0, b = 1$ and (b) logistic price response function, $m = 0.5, s = 0.1$.

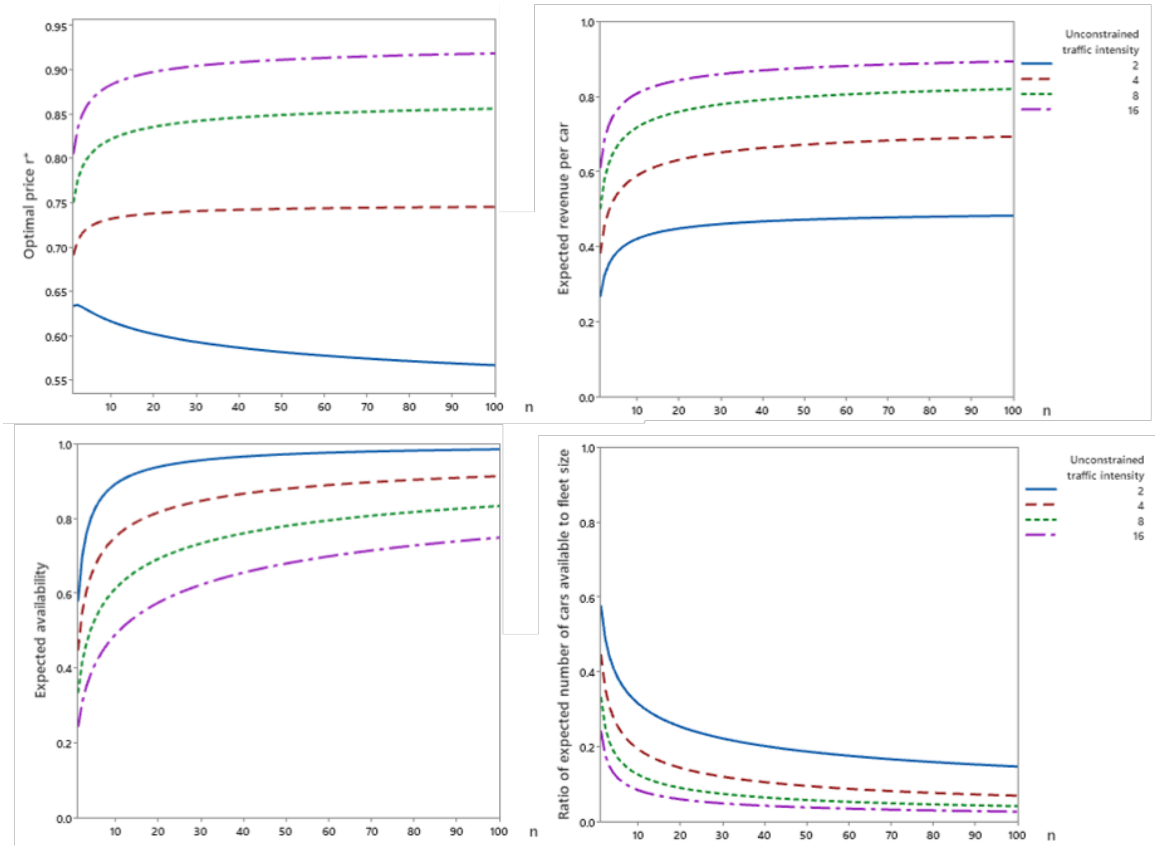

Figure 7: Single fare pricing: optimal price r^*, expected revenue per car, availability and the proportion of the fleet that is available when the price is set at r^* versus versus fleet size n for different unconstrained traffic densities $\frac{\lambda}{n\mu}$ where $\mu = 1$. The price response function is assumed to be linear $a = 0, b = 1$.

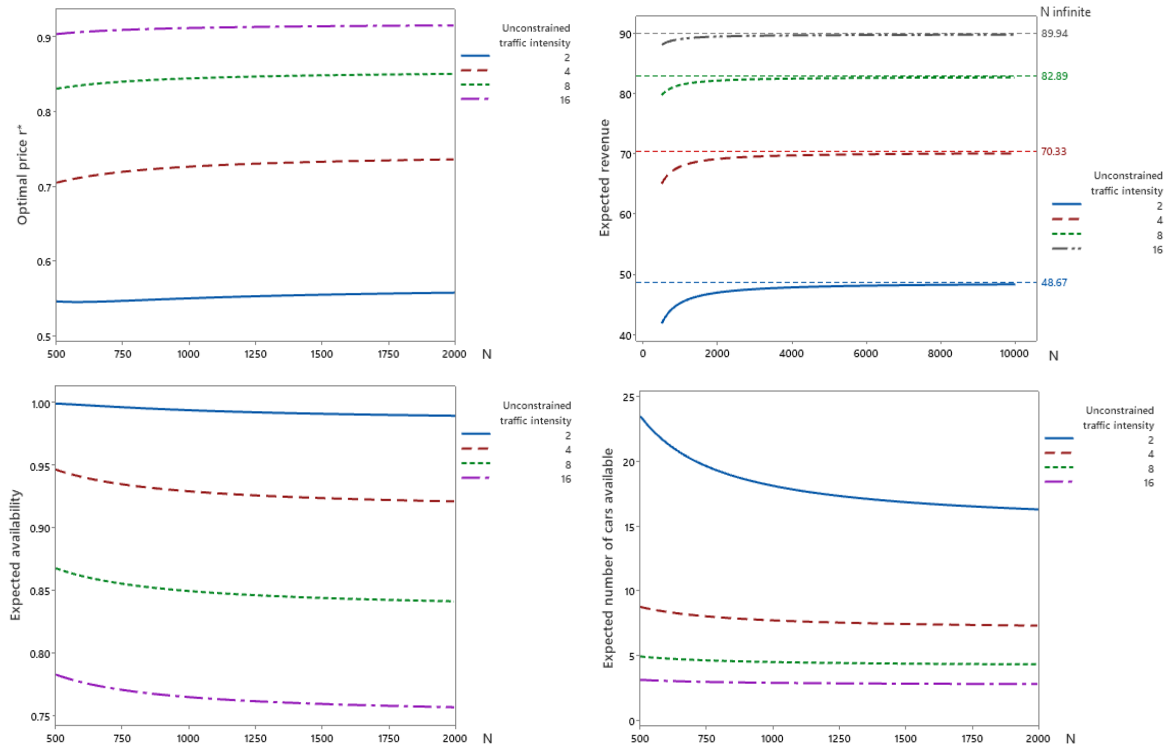

Figure 8: Single fare pricing: optimal price r^*, expected revenue, availability and number of cars available versus the size of the customer population N for $n = 100$, different unconstrained traffic densities $\frac{\lambda}{n\mu}$ where $\mu = 1$. The price response function is assumed to be linear $a = 0, b = 1$.

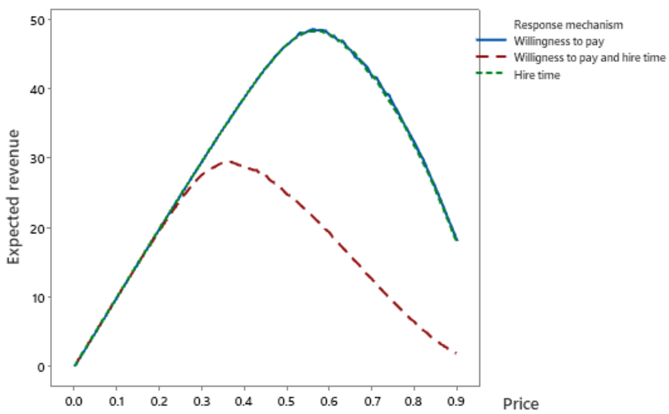

Figure 9: Single fare pricing: simulated expected revenue versus r for $n = 100, \lambda = 200, \mu = 1$ for linear-dependent willingness-to-pay or linear-dependent hire duration. The price response function is assumed to be linear $a = 0, b = 1$, and the price hire duration response function is also assumed to be linear with $c = 1$.

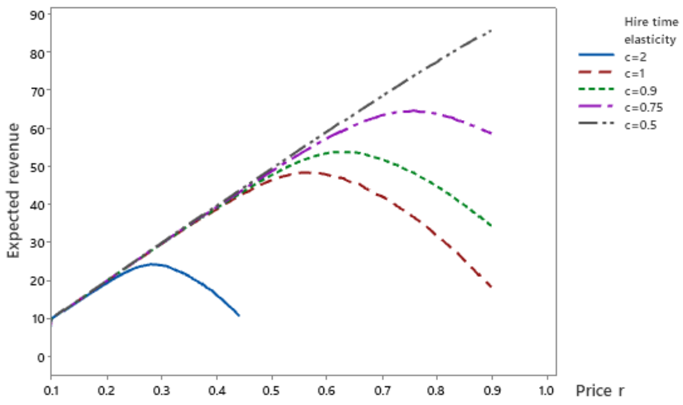

Figure 10: Single fare pricing: simulated expected revenue versus price r for $n = 100, \lambda = 200, \mu = 1$ and linearly-dependent hire durations with c between 0.50 and 2.00.

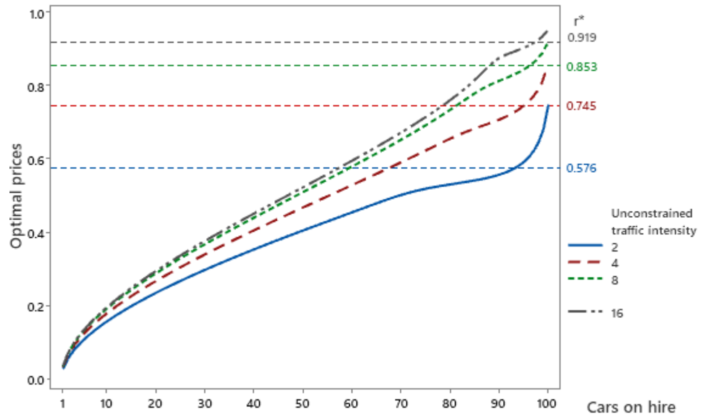

Figure 11: Multi fare pricing: optimal prices for $n = 100$ for different unconstrained traffic intensities where $n = 100, \mu = 1$. The horizontal dashed lines correspond to the optimal price r^* of the single fare scheme. We assume a linear price response function, $a = 0, b = 1$.

where $n = 100$. The curves show that the optimal fare increases as the number of cars on hire increases with steep increases at high values of the number on hire. We also see higher optimal fares for higher values of the unconstrained traffic intensity, although the difference is much smaller for states where only a small number of cars are on hire. The dashed lines on the chart show the single-fare optimal prices to charge, signaling that the optimal state-dependent fares will be below the single-fare prices for the majority of states.

Figure 12 displays the optimal expected revenue, and corresponding expected availability, and number of cars available for a fleet size $n = 2$ to $n = 100$ and different unconstrained traffic intensities. The expected revenue per car increases as the fleet size increases but the marginal increases get smaller as n gets larger. Increasing the unconstrained traffic intensity increases the expected revenue per car but the relationship is non-linear. The expected availability and ratio of the expected number of cars available to fleet size are highest when the unconstrained traffic intensity is 2 and decrease as it increases. The expected availability increases as the fleet size increases. A larger fleet size also results in a lower proportion of cars available, with the ratio of expected cars

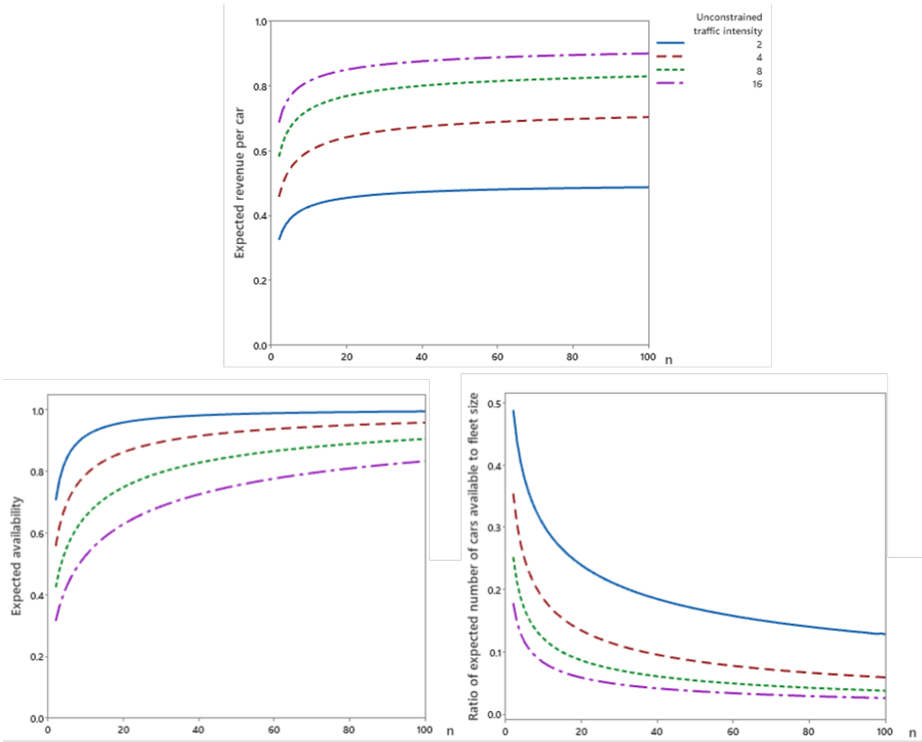

Figure 12: Multi fare pricing: optimal expected revenue versus unconstrained traffic intensities for $n = 100, \mu = 1$ and corresponding expected availability and number of cars available. We assume a linear price response function, $a = 0, b = 1$.

available to fleet size decreasing with increasing n. This results in a higher utilization of the fleet, which is attractive to service providers.

Many car sharing services use a subscription service and consequently have a finite population of members. Figure 13 shows the convergence of the expected revenue, availability and number of cars available for different values of N for $n = 100$ under different unconstrained traffic intensities. The horizontal lines on the chart indicate the values of the key performance indicators under an infinite population case. Figure 13 confirms the results of Figure 12 (for infinite N), and suggests that the infinite population case is a good approximation to the finite population model for high customer populations relative to fleet size.

4.4 Two Fare Pricing

We focus initially on the optimal prices being charged under two fare pricing. Figure 14 shows that the optimal prices r_0^* and r_1^* generally increase with n but with stepped decreases at particular values of n. These discontinuities correspond to the difference between the optimal threshold value n' and the number of cars in the fleet n increasing as we move to higher values of n, which will happen only at certain values of n. The threshold value n' is typically very close to n, taking a value of $n - 1$ for low values of n and the unconstrained traffic intensity, and moving further from n as the fleet size increases and the system becomes busier.

Figure 15 compares the optimal two fare prices r_0^* and r_1^* to r^* for different unconstrained traffic intensities and $n = 100$. This shows that the optimal $r_0^* < r^* < r_1^*$. Again we see some discontinuities

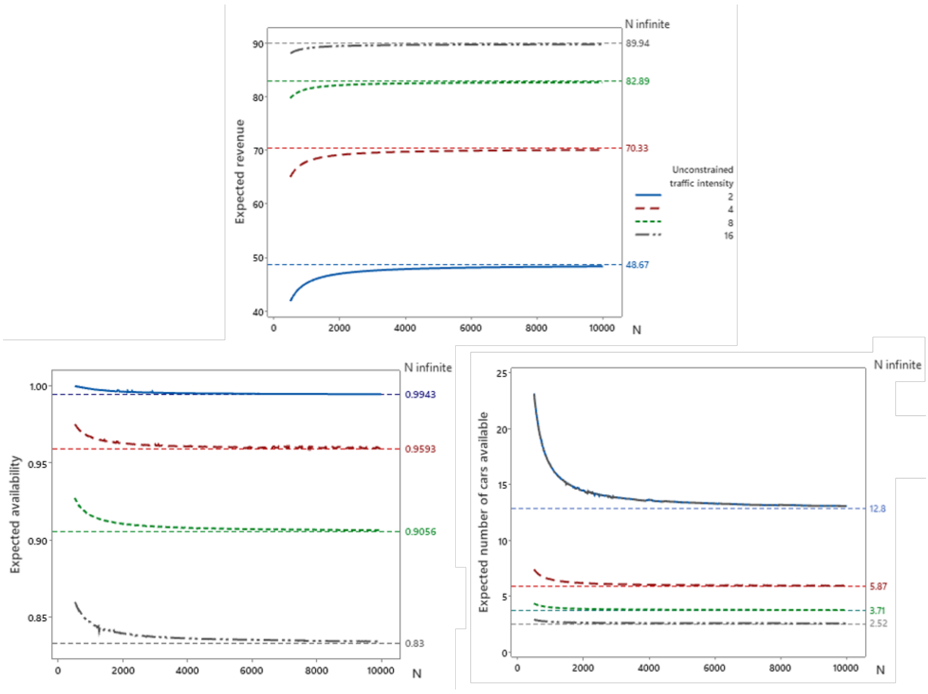

Figure 13: Multi fare pricing: finite versus infinite population optimal expected revenue, availability and number of cars for $n = 100, \mu = 1$ for different sizes of the customer population N. We assume a linear price response function, $a = 0, b = 1$.

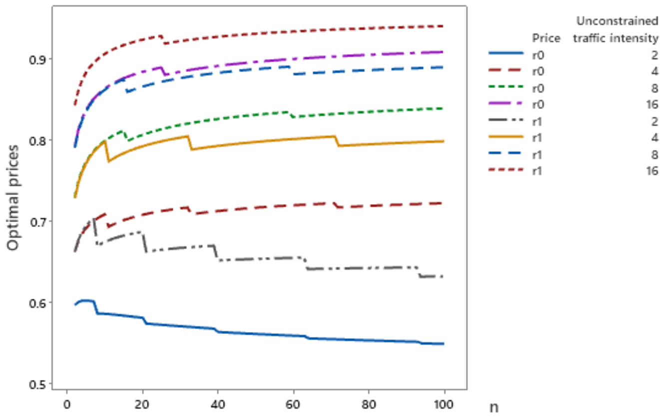

Figure 14: Two fare pricing: variation of the optimal prices r_0^* and r_1^* with n under different unconstrained traffic intensities. We assume a linear price response function, $a = 0, b = 1$ and display results that use the optimal value of n' for each n.

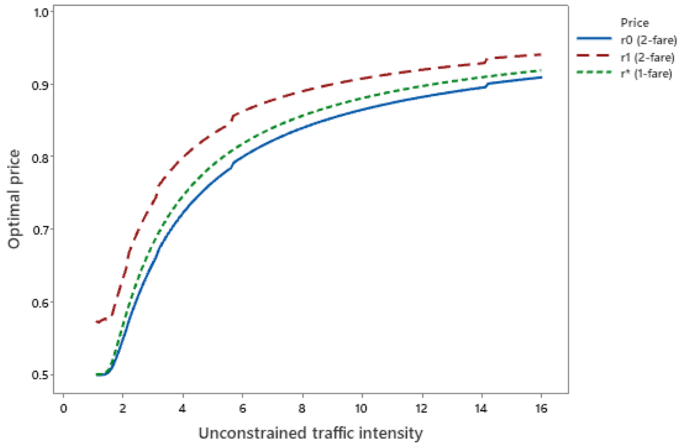

Figure 15: Two fare pricing: comparison of the optimal r_0^*, r_1^* for the two fare case to the optimal fare r^* for the single fare case for different unconstrained traffic intensities and $n = 100, \mu = 1$, infinite population. We assume a linear price response function, $a = 0, b = 1$.

where the optimal n' changes as the unconstrained traffic intensity increases.

Table 2 compares the multi fare with the single and two fare pricing strategies where hire duration is independent of price, focusing on the expected revenue $E[R]$, availability p, and number of cars available v for different values of unconstrained traffic intensities $\frac{\lambda}{n\mu}$, $n = 100, \mu = 1$ when the population is infinite or of finite size N and the price response function is linear $a = 0, b = 1$. It shows that the expected revenue is highest under the multi fare strategy; the two fare strategy still allows users to obtain cheaper prices when the number of cars on hire is low because consistently $r_0^* < r^* < r_1^*$. It also has higher expected revenues than the single fare strategy. As the population size N increases, the three key performance indicators approach those of an infinite population.

When the hire duration is dependent on the hire price, it is possible to solve the problem analytically for single, two and multi fare schemes. However, when the fleet size n increases, it becomes prohibitively expensive to do so in terms of runtime. Table 3 presents the main results for this situation, for a fleet size of $n = 8$, including the runtime. The results for $n < 8$ are available in the supplementary material (Table 9), as well as the multi fare optimal prices (Table 10).

Due to this limitation of the analytical approach, for large values of fleet size n we apply the simulation optimization procedure presented previously. Tables 4 and 5 compare the single fare to the two fare pricing strategies where the hire duration is linearly dependent on the hourly hire price r with elasticity constant c, focusing on the expected revenue $E[R]$, availability p, and number of cars available v for different values of unconstrained traffic intensities $\frac{\lambda}{n\mu}$, $n = 100, \mu = 1$ when the population is infinite or of finite size N. They show that the expected revenue is higher with the two fare scheme, and that this allows users to obtain cheaper prices when the number of cars on hire is low because consistently $r_0^* < r^* < r_1^*$. As the population size N increases, the three key performance indicators approach those of an infinite population. This is confirmed in Figure 16, which shows how the expected revenue when hire duration is linearly dependent on price.

Table 2: Comparison of expected revenue $E[R]$, availability p and number of cars available v under different multi, single, and two fare schemes for $n = 100$ and $\mu = 1$ where the hire duration is independent of the hourly hire price. We assume a linear price response function, $a = 0, b = 1$. r^*, r_0^* and r_1^* represent optimal prices and n' the optimal threshold value for two fares.

		Multi fare			Single fare				Two fares					
N	$\lambda/n\mu$	$E[R]$	p	v	r^*	$E[R]$	p	v	n'	r_0^*	r_1^*	$E[R]$	p	v
∞	2	48.68	0.99	12.80	0.57	48.32	0.98	14.85	94	0.55	0.63	48.57	0.99	13.39
	4	70.33	0.96	5.87	0.75	69.35	0.91	6.96	96	0.72	0.80	70.07	0.95	6.13
	8	82.89	0.91	3.71	0.86	82.06	0.83	4.17	97	0.84	0.89	82.68	0.89	3.81
	16	89.94	0.83	2.52	0.92	89.38	0.75	2.72	98	0.91	0.94	89.80	0.81	2.55
500	2	41.84	1.00	23.14	0.55	41.84	1.00	23.44	96	0.54	0.61	41.85	1.00	23.25
1000	2	45.23	1.00	16.70	0.55	45.11	0.99	8.76	97	0.68	0.76	64.83	0.97	7.72
2000	2	46.97	1.00	14.46	0.56	46.74	0.99	4.93	98	0.81	0.87	79.58	0.91	4.45
10000	2	48.34	0.99	13.07	0.57	48.01	0.99	3.13	98	0.89	0.92	88.02	0.84	2.96
500	4	65.03	0.98	7.38	0.70	64.31	0.95	18.08	95	0.54	0.61	45.20	1.00	17.14
1000	4	67.90	0.97	6.50	0.73	67.03	0.93	7.72	97	0.70	0.78	67.66	0.96	6.78
2000	4	69.16	0.96	6.15	0.74	68.24	0.92	4.50	98	0.83	0.88	81.28	0.90	4.09
10000	4	70.10	0.96	5.91	0.74	69.14	0.92	2.90	99	0.90	0.94	88.99	0.82	2.71
500	8	79.79	0.93	4.32	0.83	78.98	0.87	16.27	95	0.54	0.62	46.90	0.99	15.01
1000	8	81.49	0.92	3.98	0.84	80.66	0.85	7.31	97	0.71	0.79	68.91	0.95	6.43
2000	8	82.22	0.91	3.84	0.85	81.39	0.84	4.33	98	0.83	0.89	82.01	0.89	3.94
10000	8	82.76	0.91	3.73	0.86	81.93	0.84	2.81	99	0.91	0.94	89.42	0.82	2.63
500	16	88.17	0.86	2.86	0.90	87.57	0.78	15.11	95	0.55	0.63	48.24	0.99	13.68
1000	16	89.14	0.85	2.68	0.91	88.56	0.76	7.03	97	0.72	0.80	69.84	0.95	6.18
2000	16	89.56	0.84	2.60	0.92	88.99	0.76	4.20	98	0.84	0.89	82.55	0.89	3.84
10000	16	89.87	0.83	2.54	0.92	89.31	0.75	2.74	99	0.91	0.94	89.73	0.81	2.57

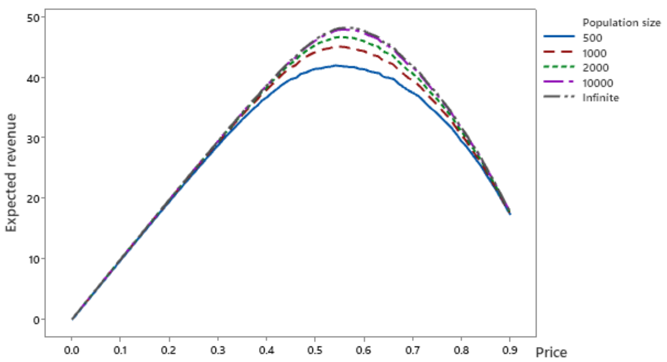

Figure 16: Single fare: simulated expected revenue versus price when hire duration is linearly dependent on price for different population sizes where hire duration elasticity $c = 1, \mu = 1$, and $\lambda = 200$.

Table 3: Results for pricing dependent on hire duration under single, two fare and multi fare schemes for fleet size $n = 8$: optimal prices r^*, r_0^*, r_1^*, expected revenue $E[R]$, availability p, number of cars available v, and runtime (RT) in seconds for different values of requests arrival rate λ (with hire return rate $\mu = 1$) and c, which describes the price-dependence of the hire duration on r, as described in Equation 3.

		Single fare					Two fares							Multi fare			
λ	c	r^*	$E[R]$	p	v	RT	n'	r_0^*	r_1^*	$E[R]$	p	v	RT	$E[R]$	p	v	RT
2	2.00	0.25	0.25	1.00	7.00	543	1	0.25	0.25	0.25	1.00	7.00	3399	0.25	1.00	7.00	15069
4	2.00	0.25	0.50	1.00	6.01	773	1	0.25	0.25	0.50	1.00	6.01	1317	0.50	1.00	6.01	7248
8	2.00	0.26	0.97	0.98	4.31	821	1	0.26	0.26	0.97	0.98	4.31	1701	0.97	0.98	4.31	7455
16	2.00	0.31	1.65	0.87	2.69	1048	1	0.29	0.31	1.65	0.87	2.69	2092	1.65	0.88	2.69	9373
2	1.00	0.50	0.50	1.00	7.00	819	1	0.50	0.50	0.50	1.00	7.00	1267	0.50	1.00	7.00	15911
4	1.00	0.50	1.00	1.00	6.01	641	1	0.50	0.50	1.00	1.00	6.01	1145	1.00	1.00	6.01	8898
8	1.00	0.53	1.95	0.98	4.31	729	1	0.51	0.53	1.95	0.98	4.31	1820	1.95	0.98	4.31	5740
16	1.00	0.62	3.29	0.87	2.69	1820	1	0.58	0.62	3.29	0.87	2.69	2853	3.30	0.88	2.69	12233
2	0.90	0.56	0.56	1.00	7.00	593	1	0.56	0.56	0.56	1.00	7.00	1416	0.56	1.00	7.00	19730
4	0.90	0.56	1.11	1.00	6.01	823	1	0.56	0.56	1.11	1.00	6.01	1875	1.11	1.00	6.01	10021
8	0.90	0.59	2.16	0.98	4.31	821	1	0.57	0.59	2.16	0.98	4.31	1445	2.16	0.98	4.31	8091
16	0.90	0.69	3.66	0.87	2.69	1139	1	0.65	0.69	3.66	0.87	2.69	2569	3.66	0.88	2.69	14291
2	0.75	0.67	0.67	1.00	7.00	683	1	0.67	0.67	0.67	1.00	7.00	1153	0.67	1.00	7.00	17892
4	0.75	0.67	1.33	1.00	6.01	636	1	0.67	0.67	1.33	1.00	6.01	1498	1.33	1.00	6.01	9779
8	0.75	0.70	2.59	0.98	4.31	682	1	0.69	0.70	2.59	0.98	4.31	2085	2.60	0.98	4.31	10204
16	0.75	0.83	4.39	0.87	2.69	1138	1	0.78	0.83	4.39	0.87	2.69	3221	4.40	0.88	2.69	19927
2	0.50	1.00	1.00	1.00	7.00	1228	1	1.00	1.00	1.00	1.00	7.00	1811	1.00	1.00	7.00	19897
4	0.50	1.00	2.00	1.00	6.00	1315	6	1.00	1.00	2.00	1.00	6.00	2631	2.00	1.00	6.00	16362
8	0.50	1.00	3.88	0.97	4.12	951	3	1.00	1.00	3.88	0.97	4.12	1996	3.88	0.97	4.12	8319
16	0.50	1.00	6.12	0.76	1.88	680	4	1.00	1.00	6.12	0.76	1.88	1269	6.12	0.76	1.88	11460

Table 4: Comparison of estimated optimal prices r^*, r_0^*, r_1^*, expected revenue $E[R]$, availability p and number of cars available v, under different population sizes for a linear price-dependent hire duration with elasticity $c = 1$, and $\mu = 1, n = 100$.

		Two fares						Single fare			
λ	N	n'	r_0^*	r_1^*	$E[R]$	p	v	r^*	$E[R]$	p	v
2	∞	94	0.53	0.61	48.14±0.14	0.98	12.64±0.19	0.57	47.90±0.14	0.98	14.70±0.26
	500	97	0.56	0.79	41.76±0.14	1.00	26.13±0.24	0.57	41.52±0.14	1.00	26.65±0.24
	1000	94	0.55	0.64	44.79±0.13	1.00	19.26±0.22	0.57	44.72±0.13	1.00	21.21±0.23
	2000	94	0.55	0.58	46.54±0.13	0.99	16.03±0.24	0.56	46.51±0.13	1.00	19.26±0.22
	10000	92	0.56	0.64	47.86±0.14	0.99	15.78±0.24	0.56	47.73±0.13	0.98	15.25±0.24
4	∞	94	0.74	0.76	69.46±0.07	0.90	6.44±0.07	0.75	69.08±0.08	0.92	7.65±0.11
	500	93	0.70	0.80	64.10±0.09	0.95	8.93±0.10	0.71	64.05±0.08	0.96	9.78±0.12
	1000	95	0.70	0.75	66.69±0.67	0.92	6.98±0.06	0.73	66.76±0.96	0.94	8.72±0.12
	2000	93	0.73	0.78	68.37±0.87	0.95	8.98±0.09	0.73	67.98±0.66	0.91	7.04±0.09
	10000	94	0.74	0.78	69.22±0.82	0.94	8.67±0.08	0.75	68.72±0.82	0.94	8.67±0.08

Table 5: Comparison of estimated optimal prices r^*, r_0^*, r_1^*, expected revenue $E[R]$, availability p and number of cars available v, under different values of hire duration elasticity c on price, where c is defined in Equation 3, $\lambda = 2, \mu = 1, n = 100$.

	Two fares						Single fare			
c	n'	r_0^*	r_1^*	$E[R]$	p	v	r^*	$E[R]$	p	v
1.00	94	0.53	0.61	47.91±0.13	0.98	12.64±0.19	0.57	48.03±0.14	0.98	14.81±0.26
0.90	95	0.64	0.70	53.87±0.18	0.99	17.03±0.26	0.68	52.13±0.19	1.00	23.49±0.27
0.75	99	0.70	0.78	63.34±0.15	0.97	11.17±0.17	0.75	63.93±0.18	0.98	14.39±0.25
0.50	100	1.00	1.00	91.68±0.14	0.92	7.87±0.14	0.85	80.74±0.06	0.83	4.46±0.08

4.5 Sensitivity of Revenue to the Fleet Size

Car sharing companies may want to assess the loss in revenue associated with reducing the size of the fleet (e.g. for maintenance) or the potential gains in increasing the size of the fleet to meet demand. We conduct experiments with the single fare and two fare strategies to determine the incremental revenue in increasing the fleet size by one car as n varies for different levels of λ. Figure 17 shows results for both cases and suggests that the incremental additional revenue gained by increasing the fleet size is positive and declines with n, but with a decreasing gradient.

The model can also be used to find the optimal value of n for a given customer population N. Figure 18 shows how the expected revenue per car varies with the fleet size n in the single fare case and suggests that this has a peak, with the peaks (indicated on the plot by vertical dashed lines) becoming more defined for lower values of N.

4.6 Time-dependent Pricing

Where demand for car sharing is time-dependent, a company may decide to implement time-dependent pricing, i.e., the price is higher in periods that are known to be busy. In this section, we compare state-dependent pricing with three time-dependent pricing methods, labeled *Queue*, *Sim*, and *BM*.

- *Queue* determines the optimal price to charge by applying the analytical single-fare pricing queueing model using the estimated arrival rate for each time period.
- *Sim* uses simulation optimization to find the optimal prices to charge in each period of the day; and
- *BM* uses a MILP-based dynamic pricing model adapted from Soppert et al. (2022).

The *BM* method acts as a benchmark for our methods. It is an adaptation of the original model in Soppert et al. (2022) but shifted towards round-trips and revenue maximization (instead of the original one-way setting with marginal contribution maximization). The resulting adaptation, which is a time-dependent dynamic pricing model for car sharing, is presented in the appendix. Given that this is a deterministic approach, we use the average data described in this section for trip duration (μ), fleet size (n) and base demand at each period of the day (λ_t). The dynamic pricing MILP discretizes prices, and the pricing decision consists of choosing one price point among a set of potential price points, each associated with a demand level. To provide a fine grid of potential prices, we divide the interval considered ($[0.5, 1]$) into intervals of 0.01 (because preliminary tests

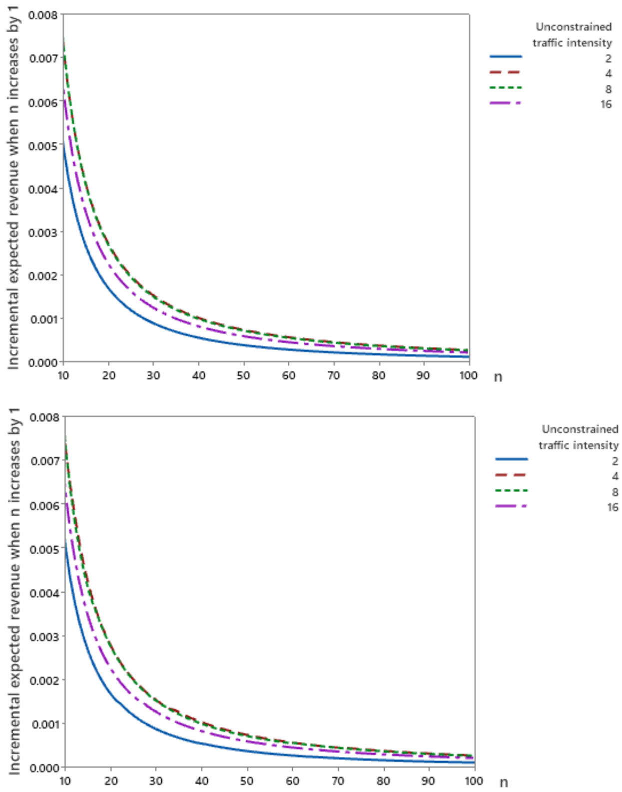

Figure 17: Sensitivity to the fleet size: incremental value of expected revenue when the fleet size n increases by one car under single fare pricing (top figure), and two fare pricing (bottom figure) for different unconstrained traffic intensities, $\frac{\lambda}{n\mu}$, where $\mu = 1$, the customer population is infinite and the price response function is linear, with $a = 0, b = 1$.

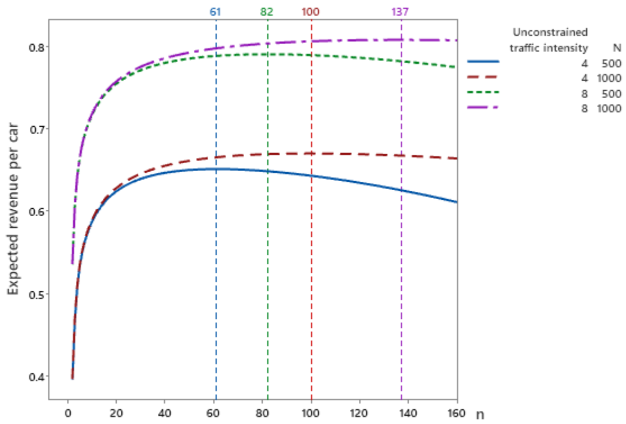

Figure 18: Single fare: expected revenue per car as a function of fleet size for a finite population with different values of N and the unconstrained traffic intensity, $\lambda/n\mu$, where $\mu = 1$ and a linear price response function, with $a = 0, b = 1$.

showed that all resulting prices were above 0.5) and derive the corresponding demand levels from Equation (1). The benchmark MILP model was solved using Gurobi Optimizer 12.0.1.

We apply the three time-dependent methods and state-dependent pricing on an example where the arrival rates are estimated using data on the number of vehicles by time of day on all roads in the UK during 2022 (Department for Transport, 2023). Table 6 reports the prices for the time-dependent methods: column 2 lists the arrival rates; columns 3-5 report the hourly prices r_t for each of the three time-dependent approaches *Queue*, *Sim*, and *BM*, respectively.

Table 6: Average arrival rates and time-dependent pricing for $\mu = 1, n = 100$. *Queue* indicates the results from the single price queuing model, *Sim* the results from the simulation model, and *BM* the results from the adapted benchmark model.

			r_t	
$[t, t+1]$	λ_t	Queue	Sim	BM
$[0, 6]$	0	-	-	-
6	182	0.50	0.74	0.53
7	333	0.70	0.71	0.70
8	378	0.74	0.75	0.73
9	301	0.70	0.69	0.67
10	293	0.67	0.68	0.66
11	305	0.68	0.70	0.67
12	317	0.70	0.72	0.68
13	319	0.60	0.65	0.69
14	342	0.69	0.70	0.71
15	384	0.77	0.70	0.74
16	414	0.78	0.79	0.76
17	408	0.71	0.71	0.75
18	312	0.68	0.69	0.68
19	219	0.50	0.73	0.55
20	153	0.53	0.73	0.50
21	113	0.46	0.70	0.50
22	82	0.48	0.72	0.53
23	49	0.49	0.73	0.50

The application of a full state-dependent multi fare scheme is cumbersome in practice. We therefore test a four fare scheme in a car sharing system assuming that the effective arrival rate is constant at $\lambda = 272.22$ during the 18 hours that the system is operational. The optimized four-fare state-dependent prices are given in Table 7.

Table 7: State-dependent pricing for $\mu = 1, n = 100$, where η is the number of cars on hire.

η	r_η
0 - 90	0.63
91 - 95	0.69
96 - 98	0.73
99 - 100	0.80

We test each of the pricing strategies using the simulation model with time-dependent arrival rates. Estimates of the key performance indicators of the time-dependent prices (Table 6) and state-dependent prices (Table 7) are reported in Table 8. The results show that the optimal prices obtained using the time-dependent approaches, namely the single fare queuing model, the simulation optimization and the benchmark MILP model are very similar and the corresponding expected

Table 8: Key performance indicators for time-dependent and state-dependent pricing strategies where demand is time-dependent with $\mu = 1, n = 100$. *Queue* indicates the results from the single price queuing model, *Sim* the results from the simulation model, and *BM* the results from the adapted benchmark model.

Pricing Strategy	$E(R)$	p	v
Queue	50.48 ± 0.31	0.55	18.63 ± 0.53
Sim	50.43 ± 0.38	0.50	27.40 ± 0.57
BM	51.11 ± 0.27	0.48	18.97 ± 0.51
State-dependent approach	52.20 ± 0.35	0.96	20.57 ± 0.50

revenues are consequently not significantly different. The state-dependent pricing scheme with four fares, assuming a constant, average arrival rate over the whole day, achieves the highest expected revenue increases of 52.20. The expected availability of 0.96 is also the highest value reached. Finally, the expected number of cars available is 20.57, a value which is only surpassed by the simulation model. These findings suggest that state-dependent pricing offers similar or better performance indicators than time-dependent pricing and that the steady state single-fare pricing (*Queue*) can be used to set the prices in each time period and achieve similar results to simulation optimization. Additionally, the comparison with the benchmark shows that our analytical approach has a good relative performance besides being fast and substantially more scalable.

5 Conclusion

The lack of a theoretical framework for dynamic pricing for car sharing has hindered its application in practice. This article bridges this gap and describes fleet usage via Markov chain models that optimize the prices charged to customers purchasing round trip car sharing. A key enabler in these Markov chain models is the modeling of price acceptance as a thinning of the customer arrival process. The Markov chain models return price recommendations reasonably fast for single, state-dependent, and time-dependent hire prices, and both finite and infinite customer populations. State-dependent pricing allows the hourly rental price to vary dependent on the number of cars available to hire when the rental period begins, enabling prices to react dynamically to the current state of the system and increasing the expected revenues by around 1% with respect to single fare pricing.

Where the **hire duration is independent of price**, the Markov chain models further reduce to much simpler and faster queuing models that enable fast price optimization even for large car sharing fleets of more than 100 vehicles. The speed and ease of returning price recommendations make dynamic pricing a practical proposition and enhances the chances of its implementation in real situations. Where the **hire duration is dependent on price**, the Markov chain model has a prohibitively large state-space for a realistic fleet size; thus, it is best solved using simulation optimization. Numerical experiments show that the analytic and simulation optimization perform comparably for small fleet sizes. Where **demand is time-dependent**, state-dependent pricing performs the best when compared with a benchmark from the literature, simulation optimization, and a time-dependent pricing structure where pricing is optimized using a series of queueing models, each set up for the current customer arrival rate. We show that assuming an infinite population is a good approximation to the finite population case for high customer populations relative to fleet size, suggesting that it can be used for setting prices even under a subscription model.

Additional complexities such as advance bookings, a business model that mixes subscription and non-subscription customers, a heterogeneous fleet or a heterogeneous customer population could be incorporated into future models. These complexities could result in higher revenues at the expense of higher computational times. Where including the complexities significantly improves the solutions, the models described in this article could be used as part of a multi-fidelity optimization approach in which an initial set of price recommendations are produced quickly in a relatively simple low fidelity model and used to speed up the optimization of a more complex high fidelity stochastic model.

We assumed round trip car sharing, where cars are picked up and returned at the same place. This is the modus operandi of UK car sharing clubs operating outside London and a more environmentally sustainable model than one-way or free-floating car sharing schemes. A natural extension would consider price incentivization to overcome operational difficulties associated with one-way hires and encourage customers to return cars to stations or bays that suit the car sharing service provider. However, a model that optimizes the operation of a connected system of round trip car sharing hubs across a city would be more useful for many car sharing schemes used in practice. Providing optimal dynamic pricing and re-allocation of cars that react dynamically to demand would have great practical benefits.

Acknowledgments

We are grateful for funding from the Engineering and Physical Research Council UK under grant numbers EP/Y008014/1 and EP/Y008324/1 which enabled this research to take place. We would also like to thank CoMoUK and the Centre for Urban Science and Progress (CUSP) London for their valuable contributions to this research.

References

Ampudia-Renuncio, M., Guirao, B., and Molina-Sanchez, R. (2018). The impact of free-floating carsharing on sustainable cities: analysis of first experiences in Madrid with the university campus. *Sustainable Cities and Society*, 43:462–475.

Aviv, Y. and Pazgal, A. (2005). A partially observed markov decision process for dynamic pricing. *Management Science*, 51(9):1400–1416.

Balseiro, S. R. and Ma, W. (2023). Dynamic pricing for reusable resources: The power of two prices. *arxiv [math.OC]*, pages 1–42.

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

Becker, H., Ciari, F., and Axhausen, K. W. (2017). Comparing car-sharing schemes in Switzerland: User groups and usage patterns. *Transportation Research Part A: Policy and Practice*, 97:17–29.

Benjaafar, S., Wu, S., Liu, H., and Gunnarsson, E. B. (2022). Dimensioning on-demand vehicle sharing systems. *Management Science*, 68(2):1218–1232.

Bojković, N., Jeremić, V., Petrović, M., and Tica, S. (2019). Preferences for car sharing service attributes among university students: Evidence from an emerging market. *Journal of East European Management Studies*, 24:636–653.

Bražėnas, M. and Valakevičius, E. (2023). Approximation of Non-Markovian Car Sharing Systems Models by Markovian One. *Lecture Notes in Networks and Systems*, 543 LNNS:458–474.

Chen, Y. and Liu, Y. (2023). Integrated Optimization of Planning and Operations for Shared Autonomous Electric Vehicle Systems. *Transportation Science*, 57(1):106–134.

Ciari, F., Bock, B., and Balmer, M. (2014). Modeling station-based and free-floating carsharing demand: Test case study for berlin. *Transportation Research Record*, 2416(1):37–47.

Ciari, F., Schuessler, N., and Axhausen, K. W. (2013). Estimation of carsharing demand using an activity-based microsimulation approach: Model discussion and some results. *International Journal of Sustainable Transportation*, 7(1):70–84.

CoMoUK (2022). Comouk annual car club report london. Technical report.

Costain, C., Ardron, C., and Habib, K. N. (2012). Synopsis of users' behaviour of a carsharing program: A case study in Toronto. *Transportation Research Part A: Policy and Practice*, 46:421–434.

Currie, C. S. and M'Hallah, R. (2024). Notes from a workshop on future directions in modelling car share schemes. https://doi.org/10.5258/SOTON/D3414.

Currie, C. S., M'Hallah, R., and Oliveira, B. B. (2025). Notes from a workshop on what areas are good for car share schemes. https://doi.org/10.5258/SOTON/D3413.

Department for Transport (2023). Road traffic estimates in great britain: 2022. `https://www.gov.uk/government/statistics/road-traffic-estimates-in-great-britain-2022` [Accessed April 2024].

Doan, X. V., Lei, X., and Shen, S. (2020). Pricing of reusable resources under ambiguous distributions of demand and service time with emerging applications. *European Journal of Operational Research*, 282:235–251.

Fanti, M. P., Mangini, A. M., Pedroncelli, G., and Ukovich, W. (2014). Fleet sizing for electric car sharing system via closed queueing networks. *2014 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 1324–1329.

Ferrero, F., Perboli, G., Rosano, M., and Vesco, A. (2018). Car-sharing services: An annotated review. *Sustainable Cities and Society*, 37:501–518.

Giorgione, G., Ciari, F., and Viti, F. (2020). Dynamic pricing on round-trip carsharing services: Travel behavior and equity impact analysis through an agent-based simulation. *Sustainability*, 12(17):6727.

Golalikhani, M., Oliveira, B. B., Carravilla, M. A., Oliveira, J. F., and Antunes, A. P. (2021a). Carsharing: A review of academic literature and business practices toward an integrated decision-support framework. *Transportation Research Part E: Logistics and Transportation Review*, 149(February):102280.

Golalikhani, M., Oliveira, B. B., Carravilla, M. A., Oliveira, J. F., and Pisinger, D. (2021b). Understanding carsharing: A review of managerial practices towards relevant research insights. *Research in Transportation Business and Management*, 41(March):100653.

Guo, G. and Kang, M. (2022). Rebalancing and Charging Scheduling With Price Incentives for Car Sharing Systems. *IEEE Transactions on Intelligent Transportation Systems*, 23(10):18592–18602.

Han, D. L., Tang, L. C., and Huang, H. C. (2010). A Markov model for single-leg air cargo revenue management under a bid-price policy. *European Journal of Operational Research*, 200(3):800–811.

Hartl, B., Sabitzer, T., Hofmann, E., and Penz, E. (2018). "Sustainability is a nice bonus" the role of sustainability in carsharing from a consumer perspective. *Journal of Cleaner Production*, 202:88–100.

Huang, K., An, K., Rich, J., and Ma, W. (2020). Vehicle relocation in one-way station-based electric carsharing systems: A comparative study of operator-based and user-based methods. *Transportation Research Part E: Logistics and Transportation Review*, 142(July):102081.

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

Jie, F., Standing, C., Biermann, S., Standing, S., and Le, T. (2021). Factors affecting the adoption of shared mobility systems: Evidence from Australia. *Research in Transportation Business and Management*, 41:100651.

Jorge, D., Molnar, G., and de Almeida Correia, G. H. (2015). Trip pricing of one-way station-based carsharing networks with zone and time of day price variations. *Transportation Research Part B: Methodological*, 81:461–482.

Krueger, R., Rashidi, T. H., and Rose, J. M. (2016). Preferences for shared autonomous vehicles. *Transportation Research Part C: Emerging Technologies*, 69:343–355.

Li, Y., Chen, S., Hu, L., Liang, Z., Jiang, Y., and Tang, Y. (2021). Simulation-optimization for station capacities, fleet size, and trip pricing of one-way electric carsharing systems. *Journal of Cleaner Production*, 321(September):129035.

Litman, T. (2000). Evaluating carsharing benefits. *Transportation Research Record*, 1702(1):31–35.

Maas, S. and Attard, M. (2020). Attitudes and perceptions towards shared mobility services: Repeated cross-sectional results from a survey among the maltese population. *Transportation Research Procedia*, 45:955–62.

Müller, C., Gönsch, J., Soppert, M., and Steinhardt, C. (2023). Customer-Centric Dynamic Pricing for Free-Floating Vehicle Sharing Systems. *Transportation Science*, 57(6):1406–1432.

Nakamura, A., Phung-Duc, T., and Ando, H. (2022). Queueing analysis of a Car/Ride-Share system. *Annals of Operations Research*, 310(2):661–682.

Namazu, M., MacKenzie, D., Zerriffi, H., and Dowlatabadi, H. (2018). Is carsharing for everyone? Understanding the diffusion of carsharing services. *Transport Policy*, 63:189–99.

Pantuso, G. (2022). Exact solutions to a carsharing pricing and relocation problem under uncertainty. *Computers and Operations Research*, 144(March):105802.

Phillips, R. L. (2005). *Pricing and Revenue Optimization*. Stanford University Press.

Rahimi, N., Rohani, M., and Asman, N. (2021). The effects of fuel price variations on students' travel patterns in higher education institutions. *Kajian Malaysia*, 39:203–232.

Roblek, V., Meško, M., and Podbregar, I. (2021). Impact of car sharing on urban sustainability. *Sustainability*, 13(2).

Soppert, M., Oliveira, B. B., Angeles, R., and Steinhardt, C. (2024). On the benefit of combining car rental and car sharing. *Journal of Business Economics*. (in press) DOI: 10.1007/s11573-024-01204-9.

Soppert, M., Steinhardt, C., Müller, C., and Gönsch, J. (2022). Differentiated pricing of shared mobility systems considering network effects. *Transportation Science*, 56(5):1279–1303.

Ströhle, P., Flath, C. M., and Gärttner, J. (2019). Leveraging customer flexibility for car-sharing fleet optimization. *Transportation Science*, 53:42–61.

Xu, M., Meng, Q., and Liu, Z. (2018). Electric vehicle fleet size and trip pricing for one-way carsharing services considering vehicle relocation and personnel assignment. *Transportation Research Part B: Methodological*, 111:60–82.

Zhou, T., Fields, E., and Osorio, C. (2023). A data-driven discrete simulation-based optimization algorithm for car-sharing service design. *Transportation Research Part B*, 178(102818).

Zhu, H., Gao, Y., Hou, Y., Wang, Z., and Feng, X. (2019). Real-time pricing considering different type of smart home appliances based on Markov decision process. *International Journal of Electrical Power and Energy Systems*, 107(December 2018):486–495.

Appendix

The appendix presents the adaptation of the model in Soppert et al. (2022) to the round-trip setting with no variable costs considered, based on a time-dependent approach using hourly demand rates (and prices). The adaptation of the original model specifically consists of (1) the shift from one-way to round-trip, by dropping the location indexes in all decision variables and the parameters related to demand and starting number of vehicles, and (2) the consideration of revenue maximization, which consisted of removing from the objective function the marginal cost per trip fulfilled (which was a parameter). The notation presented here is parallel to the notation in the original model, except for these modifications.

Equation A1 presents the revenue-maximizing objective function, where the number of realized rentals in a given period and at a given price point is multiplied by the corresponding value per time period and its duration. The fact that the decision variable tracks the number of realized rentals depending on the price charged increases the number of decision variables yet allows the objective function to be linear. Constraints A2 to A4 track the number of available vehicles at each period. This value is related to the stock of vehicles (in the current and previous period) and the realized rentals that use those cars (Constraints A2 and A3), and its value at the beginning of the time horizon is parametrized (Constraint A4). Constraints A5 ensure only one price point is selected in each time period. Constraints A6 limit the realized rentals to the demand, which depends directly on the price charged. Then, Constraints A7 to A10 enforce demand as a lower bound on the realized rentals, to ensure that all demand is fulfilled if cars are available (to avoid the model artificially saving capacity for later, more profitable rentals). Constraints A7 and A8 force the binary variable q_t to take the value 1 if demand exceeds supply and 0 otherwise. If demand exceeds supply, Constraints A10 ensure all cars are rented. If supply exceeds demand, Constraints A9 guarantee that all demand is fulfilled. Finally, Constraints A11 to A15 define the domain of the decision variables.

Indexes and parameters:

$\mathcal{T}$	Set of time periods, indexed by t.
$\mathcal{M}$	Set of price points, indexed by m.
d_t^m	Base demand for trips in period t if price point m is charged.
l	Average rental duration of a trip.
$\hat{a}_0$	Starting number of vehicles.
p_m	Value of price point m.
$\bar{M}$	Very large coefficient.

Decision variables:

r_t^m	Number of realized rentals in period t at price point m.
s_t	Stock of idle vehicles at period t.
a_t	Available vehicles at the beginning of period t.
y_t^m	$= 1$ if price point m is charged for a trip starting at period t ($= 0$ otherwise).
q_t	binary auxiliary variables indicating whether there is more demand (or more supply) in at period t.

$$
\begin{aligned}
\max_{y,q,r,a,s} \quad & \sum_{t \in T} \sum_{m \in M} r_t^m \cdot l \cdot p^m && && \text{(A1)} \\
\text{s.t.} \quad & a_t = \sum_{m \in M} r_t^m + s_t, && \forall t \in \mathcal{T} && \text{(A2)} \\
& \sum_{m \in M} r_t^m + s_t = a_{(t+1)}, && \forall t \in \mathcal{T} \setminus \{T\} && \text{(A3)} \\
& a_0 = \hat{a}_0 && && \text{(A4)} \\
& \sum_{m \in M} y_t^m = 1, && \forall t \in \mathcal{T} && \text{(A5)} \\
& r_t^m \leq d_t^m \cdot y_t^m, && \forall t \in \mathcal{T}, m \in \mathcal{M} && \text{(A6)} \\
& \sum_{m \in M} d_t^m \cdot y_t^m - a_t \leq \bar{M} \cdot q_t, && \forall t \in \mathcal{T} && \text{(A7)} \\
& \sum_{m \in M} -d_t^m \cdot y_t^m + a_t \leq \bar{M} \cdot (1 - q_t), && \forall t \in \mathcal{T} && \text{(A8)} \\
& \sum_{m \in M} d_t^m \cdot y_t^m \leq \sum_{m \in M} r_t^m + \bar{M} \cdot q_t, && \forall t \in \mathcal{T} && \text{(A9)} \\
& s_t \leq \bar{M} \cdot (1 - q_t), && \forall t \in \mathcal{T} && \text{(A10)} \\
& y_t^m \in \{0, 1\}, && \forall t \in \mathcal{T}, m \in \mathcal{M} && \text{(A11)} \\
& q_t \in \{0, 1\}, && \forall t \in \mathcal{T} && \text{(A12)} \\
& r_t^m \in \mathbb{R}_0^+, && \forall t \in \mathcal{T}, m \in \mathcal{M} && \text{(A13)} \\
& s_t \in \mathbb{R}_0^+, && \forall t \in \mathcal{T} && \text{(A14)} \\
& a_t \in \mathbb{R}_0^+, && \forall t \in \mathcal{T} && \text{(A15)}
\end{aligned}
$$