Adaptive Pricing Mechanisms for On-Demand Mobility

Maciej Drwal*, Enrico Gerding*, Sebastian Stein*,
Keiichiro Hayakawa◇, Hironobu Kitaoka•
*University of Southampton, Southampton, United Kingdom, {M.K.Drwal, eg, S.Stein}@soton.ac.uk
◇Toyota Central R&D Labs., Inc., Aichi, Japan, kei-hayakawa@mosk.tytlabs.co.jp
•Toyota Motor Corporation, Aichi, Japan, hironobu_kitaoka_ab@mail.toyota.co.jp

ABSTRACT

We consider on-demand car rental systems for public transportation. In these systems, demands are often unbalanced across different parking stations, necessitating costly manual relocations of vehicles. To address this so-called "deadheading" effect and maximise the operator's revenue, we propose two novel pricing mechanisms. These adaptively adjust the prices between origin and destination stations depending on their current occupancy, probabilistic information about the customers' valuations and estimated relocation costs. In so doing, the mechanisms incentivise drivers to help rebalance the system and place a premium on trips that lead to costly relocations. We evaluate the mechanisms in a series of experiments using real historical data from an existing on-demand mobility system in a French city. We show that our mechanisms achieve an up to 64% increase in revenue for the operator and at the same time up to 36% fewer relocations.

Keywords

optimal mechanism design, electric vehicles, optimisation, on-demand mobility, carsharing

1. INTRODUCTION

We consider the problem of optimal pricing in the emerging market of on-demand electric vehicle rentals for urban settings. On-demand mobility is contributing to the improvement of quality of municipal transportation by reducing car ownership and, in effect, by decreasing road congestion and offering convenient alternative commuting choices [14]. In particular, we consider a one-way rental setting, where users are free to return a rented vehicle to any available station and this may be different from the departure station. While providing flexibility for users, a major issue is the "deadheading effect", which occurs when too many users drop their car off in the same location and a shortage of vehicles occurs in other locations [13]. This results in loss of bookings and significant manual relocation costs. To address this issue and improve revenue, we investigate novel adaptive pricing mechanisms to incentivise users to choose different locations, and help rebalance the system. This is a challenging problem since different prices can be set for each origin-destination pair, and these prices need to be updated given the current state of the system.

In the setting considered in this paper, a company operates a network of vehicle stations. Each station may hold a limited number of vehicles that are available for clients to be rented for a certain period of time (usually limited to a few hours). A client is free to return the vehicle to any station of her choice (declared beforehand through a reservation system, typically a mobile application). However, a payment is charged per unit of use time. The operator's revenue depends both on the prices charged, as well as on the variable operational costs, of which the main contributing factor is the "deadheading effect". This occurs when too many clients drop their cars at the same station in a short period of time. Such vehicles need to be consecutively relocated to different stations, especially to the ones that have a high number of incoming clients and low vehicle availability. Otherwise, the operator incurs losses due to service unavailability. However, relocations require a dedicated crew of employees, which utilises both manpower and additional fuel (or electricity).

To address the significant costs of deadheading, we consider adaptive pricing mechanisms, in which rental prices depend on the origin and declared destination of the journey, as well as on the state of the system. This approach allows the operator to incentivise clients to choose destination stations so that deadheading is reduced or avoided, and the numbers of vehicles at all stations are kept balanced.

In more detail, our adaptive pricing mechanisms enable both an improvement in revenue and a reduction in the number of necessary relocations. This is achieved by utilising probabilistic information about the clients' varying preferences for different destination stations when setting prices. These different preferences may arise, for example, from flexibility in the clients' final destination or in their willingness to walk further to their final destination given a price discount. As different levels of information may be available about individual clients in practice, we propose two pricing mechanisms that vary in their requirements for information.

In the first approach, we assume that only general probabilistic information about clients is available. Thus, the preferred destination station of a given client is unknown and prices are shown to the client before she chooses her destination. In the second approach, we consider a two-stage process where, first, the client chooses both the origin and her preferred destination. This information is then used to compute prices for alternative destination stations, which are then presented to the client. For both settings, we present an algorithm for optimising prices, and we evaluate the mecha-

Appears in: *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.*

nisms using real data from an existing on-demand car rental system in France.

To summarise, we advance the state of the art in the following ways:

- We propose a new adaptive pricing mechanism for on-demand mobility based on a principled probabilistic mechanism design framework. Unlike existing work, this considers both the income from vehicle rentals as well as the costs from relocations, which vary dynamically depending on the state of the system.
- We design a two-stage version of the mechanism, which enables it to use more fine-grained information about a client's preferences and thereby achieve a higher revenue.
- We conduct an experimental study of both mechanisms using real data from an on-demand mobility system in France, showing that our mechanisms achieve an up to 64% increase in revenue while requiring up to 36% fewer relocations at the same time.

The paper is organised as follows. In Section 2 we briefly present related work. In Section 3 the problem formulation is given. That section starts with a description of the system and then formalises the optimisation problems we consider. Next, in Section 4 we describe the implementation of our two pricing mechanisms. Section 5 contains the experimental results, and finally, Section 6 concludes the paper.

2. RELATED WORK

Due to the recent rise of popularity of on-demand mobility services, a number of studies in the area have been published. Solution approaches for dealing with "deadheading" and the relocation problem have been considered for both bike-sharing systems [22, 8, 19], as well as car-sharing systems [13, 18, 6]. Many of the existing approaches involve mathematical programming for relocation planning, as well as mechanisms for extracting private information about the clients' transportation preferences. However, there has been no attempt to combine revenue maximising mechanisms with the challenge of reducing variable relocation costs in the on-demand mobility setting so far. Thus, our approach is the first to address both revenue maximisation and relocation minimisation in one principled and concise framework. Moreover, the presented approach seamlessly integrates with existing booking systems and relocation policies.

The mechanisms considered in this paper are based on an underlying optimal mechanism design problem called the monopolist pricing problem, which considers how to set revenue-maximising prices for stochastic client types. The problem dates back to the work of Mussa and Rosen [17], and it has been considered by mathematical economists for a long time [20, 15, 2]. In the past few years, it has also attracted the attention of the computer science community [5, 11, 4]. Many algorithmic and complexity-theoretic results concerning different variants of this problem were developed. Although for the discrete set of valuations and discrete set of possible prices it is easy to formulate linear programs, these are exponential in size, and in general the multidimensional pricing problem is NP-hard [9]. This makes existing work unsuitable for the on-demand mobility settings we consider here, as they are potentially large in scale (with dozens of stations, large numbers of feasible prices and hundreds or more of clients per day). Moreover, the settings are dynamic, with clients arriving over time, and there are variable costs associated with relocations, which also change with the state of the system. In extending the monopolist pricing problem to on-demand mobility systems and in proposing scalable algorithms, we are the first to simultaneously address all of these challenges.

3. THE OPTIMAL PRICING PROBLEM

In this section we introduce the notation, definitions and basic assumptions of our model. We start with a high-level description of the on-demand mobility system before detailing the agents and the pricing problem. Our model consists of an operator and a sequence of clients. The system consists of M stations. A client approaches station i (called *origin* station), and wishes to rent a vehicle, in order to travel to station j (called *destination* station). The operator receives payment p for the service, but also needs to cover the costs of providing it. Costs depend on the state of the system (which we assume is fully known to the operator) and on the client's choice of origin and destination pair. In our model we consider only the variable costs caused by the need for relocating vehicles after the client terminates the journey (all other costs are considered fixed, and are omitted from the model). We assume that the cost for each origin-destination pair and for each time instant is given. In general, these costs would depend on the way that the operator schedules the relocations (we discuss this problem later in Section 4.3).

3.1 Model of Agents

In order to book a journey from the origin station, the client declares her choice of destination station. We assume that the client can be fully described by an index of her origin station and by a vector $\mathbf{x} \in \Omega \subset \mathbb{R}_+^M$, called her type. The set Ω is a fixed space of all allowed types. For a type vector $\mathbf{x} = [x_1, \ldots, x_M]$, element x_j represents the valuation per unit time of rental associated with jth destination choice. Moreover, it is assumed that the client's type is drawn from a probability distribution with density f_i, for each origin station i, which is known to the operator.[1]

The operator posts a set of prices for all possible destinations that the client can choose to declare for returning the car. The client's utility is given as follows:

$$u(\mathbf{x}, \mathbf{q}, p) = \mathbf{x} \cdot \mathbf{q} - p. \tag{1}$$

Scalar $p \in \mathbb{R}$ is the payment per time unit that the client makes to the system, and vector $\mathbf{q}$ indicates the choice of station. Note that, since both the valuation x and the payment p are normalised per unit of time, we can ignore travel duration in the utility function. We denote by $Q = \{\mathbf{q} \in \{0,1\}^M : \sum_{i=1}^M q_i \leq 1\}$ the set of feasible choices (that is, we assume that each client books at most one destination).

It is assumed that clients are individually rational, and thus would refuse any offer that results in a utility of less than zero. Furthermore, we assume that they seek to maximise their utility (1). Finally, from the operator's point of view, types of clients are i.i.d. random variables, distributed according to the density f_i, and clients approach the system sequentially

[1] Note we assume that each client has one fixed origin station, i, which is were the booking is initiated, but is flexible about the destination station, as expressed through the type $\mathbf{x}$.

(there is no competition between them). Each client entering the system results in a *transaction* and we denote by $\tau(t) \in \{1, \ldots, M\}$ a random variable indicating the origin station of the client in the t-th transaction.

3.2 Pricing Problem Formulation

It is assumed that the operator is a monopolist, and so we neglect the potential effects of competition with providers of similar services. The problem faced by the operator is to determine a price schedule $\mathbf{p}(t) = [p_1(t), \ldots, p_M(t)]$ that maximises the expected revenue from a sequence of T transactions, $t = 1, \ldots, T$. By $C(j, t)$ we denote the cost of relocating a car that would be generated if the client selects destination j from origin $\tau(t)$ in transaction t.

Under the above assumptions, given probability distribution densities f_i supported on Ω, the distribution of origin station τ, and cost function $C(j, t)$, we can determine the prices that maximise the expected revenue by solving the following optimisation problem:

maximise:

$$R(\mathbf{p}) = \sum_{t=1}^{T} \int_{\Omega} \left[\mathbf{q}(\mathbf{x}) \cdot \mathbf{p}(t) - \sum_{j=1}^{M} C(j,t) q_j(\mathbf{x}) \right] f_{\tau(t)}(\mathbf{x}) d\mathbf{x} \tag{2}$$

subject to:

$$\forall_t \ \forall_{\mathbf{x} \in \Omega} \ \ U_t(\mathbf{x}) = \max_{\mathbf{q} \in Q} \left\{ \mathbf{q} \cdot (\mathbf{x} - \mathbf{p}(t)) \right\}, \tag{3}$$

$$\forall_t \ \forall_{\mathbf{x} \in \Omega} \ \ U_t(\mathbf{x}) \geq 0, \tag{4}$$

where $\mathbf{q}(\mathbf{x}) \in Q$ is referred to as an *allocation* function and denotes the (unique) function that maximises $\mathbf{q} \cdot (\mathbf{x} - \mathbf{p}(t))$ for a given type $\mathbf{x}$. This vector can be seen as a function of type $\mathbf{x}$ which, along with price schedule $\mathbf{p}(t)$, is an unknown variable in the optimisation problem. The objective function in (2) is the expected value of the difference between the payment and the cost of the most preferred allocation $\mathbf{q}(\mathbf{x})$, being sold to client of type $\mathbf{x}$.

Equation (3) is called the incentive compatibility (IC) constraint. This both defines the utility and, at the same time, ensures that the allocation always maximises the user's utility. The constraint can be written as:

$$\forall_{\mathbf{x} \in \Omega} \ \ \forall_{\mathbf{q} \in Q} \ \ \mathbf{x} \cdot \mathbf{q}(\mathbf{x}) - \mathbf{p} \cdot \mathbf{q}(\mathbf{x}) \geq \mathbf{x} \cdot \mathbf{q} - \mathbf{p} \cdot \mathbf{q}.$$

Equation (4) ensures individual rationality (IR):

$$\forall_{\mathbf{x} \in \Omega} \ \ \mathbf{x} \cdot \mathbf{q}(\mathbf{x}) - \mathbf{p} \cdot \mathbf{q}(\mathbf{x}) \geq 0.$$

There is also a constraint on the choice of allocation function $\mathbf{q}(\mathbf{x}) \in Q$ embedded in the constraint (3), which, by our definition of Q, guarantees that up to one destination is selected.

Note that this problem can be seen as a bi-level optimisation, where an optimisation sub-problem appears in the constraint (3) (it can be also seen as a two-player game between the client choosing strategy $\mathbf{q}(\mathbf{x})$ and the seller choosing strategy $\mathbf{p}$). In the next section we show how to reformulate this problem into a single-level optimisation problem.

3.3 Solving the Pricing Problem

The problem (2)–(4) is difficult to solve in its general form. Specifically, the cost function $C(j, t)$ is usually not known in advance, but its values for any transaction t can be accessed only at the time of that transaction. Prediction of these costs is a challenging problem, as it involves predicting future demand and station occupancy levels. Instead, here we use a heuristic cost function (see Section 4.3) and decompose the problem into separate optimisation problems for each transaction $t = 1, \ldots, T$. This results in a myopic approach to the revenue-maximisation problem, but, in return, allows us to work with optimisation problem instances of realistic sizes.

Furthermore, the IC constraint (3) involves quantification over the entire set of types Ω. However, it is known that this constraint is equivalent to the requirement that $U(\mathbf{x})$ is a convex continuous function on Ω [15]. Considering this, the formulation of the pricing problem for a single transaction can be written as an optimisation problem over unknown utility functions U:

maximise:

$$R(U) = \int_{\Omega} \left[\mathbf{x} \cdot \nabla U(\mathbf{x}) - U(\mathbf{x}) - C(\mathbf{x}) \right] f_{\tau(t_0)}(\mathbf{x}) d\mathbf{x} \tag{5}$$

subject to:

$$U \text{ is convex continuous on } \Omega, \tag{6}$$

$$C(\mathbf{x}) = \sum_{j=1}^{M} C(j, t_0) \nabla U_j(\mathbf{x}), \tag{7}$$

$$\forall_{\mathbf{x} \in \Omega} \ \ U(\mathbf{x}) \geq 0, \tag{8}$$

$$\forall_{\mathbf{x} \in \Omega} \ \ \nabla U(\mathbf{x}) \in Q. \tag{9}$$

Here, the value t_0 is an index of the considered transaction. Given an optimal U^*, the optimal prices can be computed using the definition of utility:

$$p^*(\mathbf{q}) = \mathbf{x} \cdot \mathbf{q} - U^*(\mathbf{x}). \tag{10}$$

This "dual" formulation constitutes a problem of calculus of variations with the "non-standard" constraint that the unknown function must be convex. Since even for standard probability distributions f_i analytical solutions are not known, in practice, problems of this kind are solved by means of discretisation and approximation [7, 16]. However, the efficiency of general solution methods (such as the finite difference method or the finite element method) suffer from the high dimensionality of the solution space. The dimensionality can sometimes be reduced by disregarding stations of low valuations in the clients' type, but can still be considerably high in practice.

Fortunately, the dual formulation (5)–(9) allows for applying a very efficient solution method that is based on discretisation and sampling. Instead of solving the problem for an unknown continuous function U, we construct a discretised problem by sampling N points $\mathbf{x}_i \in \Omega$ from the client's type distribution, and restricting the solution only to the values $U(\mathbf{x}_i)$, $i = 1, \ldots, N$. The sampling is carried out using Markov Chain Monte Carlo (e.g., Metropolis-Hastings algorithm). Given the samples, we construct a linear program, which can be solved efficiently by standard methods. In the discretised problem, the constraint (6) becomes (13), and is easy to handle [10, 3].

Thus, the discretised problem is the following:

maximise:

$$\frac{1}{N} \sum_{i=1}^{N} \left[(\mathbf{x}_i - \mathbf{c}_i) \cdot \nabla U(\mathbf{x}_i) - U(\mathbf{x}_i) \right] \tag{11}$$

subject to:

$$\forall_{i=1,\dots,N} \;\; U(\mathbf{x}_i) \geq 0, \;\; \nabla U(\mathbf{x}_i) \geq 0, \tag{12}$$

$$\forall_{i,j=1,\dots,N} \;\; \nabla U(\mathbf{x}_j)(\mathbf{x}_i - \mathbf{x}_j) \leq U(\mathbf{x}_i) - U(\mathbf{x}_j), \tag{13}$$

$$\forall_{i=1,\dots,N} \;\; \nabla U(\mathbf{x}_i) \in Q, \tag{14}$$

where $\mathbf{c}_i = [C(1,t_0), \dots, C(M,t_0)]$. In this problem, the unknown (scalar) variables are $U(\mathbf{x}_i)$ and $\nabla U(\mathbf{x}_i)$, for $i = 1, \dots, N$.

The above formulation is useful if we can sample from distribution f_i. We will also consider a simple special case of the problem (2)–(4) in which the distribution of clients' types is discrete, and for each of a finite number of possible valuation vectors $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(K)}$, we have $\rho_k = \mathbf{Pr}[\mathbf{x} = \mathbf{x}^{(k)}]$, $k = 1, \dots, K$. In this case, we may discretise the original formulation directly. Let $\hat{\phi}_j$ be the price associated with the jth destination. Let $\hat{q}_{kj}$ be the variable denoting allocation of jth station to the client of type $\mathbf{x}_k$ (i.e., $\hat{\mathbf{q}}_k = \mathbf{q}(\mathbf{x}_k)$ as a vector), for $k = 1, \dots, K$. We obtain:

maximise:

$$\sum_{k=1}^{K} \rho_k \sum_{j=1}^{M} \hat{q}_{kj} \left(\hat{\phi}_j - C(j,t_0) \right) \tag{15}$$

subject to:

$$\forall_k \;\; \sum_{j=1}^{M} \hat{q}_{kj} \left(x_j^{(k)} - \hat{\phi}_j \right) \geq 0, \tag{16}$$

$$\forall_k \; \forall_l \;\; \sum_{j=1}^{M} \hat{q}_{kj} \left(x_j^{(k)} - \hat{\phi}_j \right) \geq x_l^{(k)} - \hat{\phi}_l, \tag{17}$$

$$\forall_k \;\; \sum_{j=1}^{M} \hat{q}_{kj} \leq 1. \tag{18}$$

In order to solve this mixed-integer problem (MIP) using standard methods we apply appropriate linearisation [1] of the mixed terms $\hat{z}_{kj} = \hat{q}_{kj}\hat{\phi}_j$. We obtain:

maximise:

$$\sum_{k=1}^{K} \rho_k \sum_{j=1}^{M} (z_{kj} - C(j,t_0)\hat{q}_{kj}) \tag{19}$$

subject to:

$$\forall_k \forall_j \;\; z_{kj} - B\hat{q}_{kj} \leq 0, \tag{20}$$

$$\forall_k \forall_j \;\; \phi_j + B\hat{q}_{kj} - z_{kj} \leq B, \tag{21}$$

$$\forall_k \forall_j \;\; z_{kj} - \phi_j \leq 0, \tag{22}$$

$$\forall_k \;\; \sum_{j=1}^{M} \hat{q}_{kj} x_{kj} - \sum_{j=1}^{M} z_{kj} \geq 0, \tag{23}$$

$$\forall_k \forall_l \;\; \sum_{j=1}^{M} \hat{q}_{kj} x_{kj} - \sum_{j=1}^{M} z_{kj} \geq x_{kl} - \phi_l, \tag{24}$$

$$\forall_k \;\; \sum_{j=1}^{M} \hat{q}_{kj} \leq 1. \tag{25}$$

Here B is the maximum price allowed to be posted. Unknown variables z_{ij}, as well as ϕ_j, are nonnegative continuous, while $\hat{q}_{kj}$ are a binary variables.

Alternatively, we may reformulate (5)–(9) for discrete distribution of types. By substituting the formula for utility $u_k = \hat{\mathbf{q}} \cdot (\mathbf{x} - \hat{\phi})$ into (15)–(18) we derive:

maximise:

$$\sum_{k=1}^{K} \rho_k \sum_{j=1}^{M} (\hat{q}_{kj} x_{kj} - \hat{q}_{kj} C(j,t_0) - u_k) \tag{26}$$

subject to:

$$\forall_i \forall_j \;\; u_j - u_i \geq \sum_{l=1}^{M} \hat{q}_{il}(x_{jl} - x_{il}), \tag{27}$$

$$\forall_k \;\; \sum_{j=1}^{M} \hat{q}_{kj} \leq 1. \tag{28}$$

This formulation, with new nonnegative continuous variables u_k, is very concise and allows for computing solution very fast, which is suitable for real-time applications. In the realistic settings considered in Section 5, the solution is typically found in less than 3 seconds on a standard laptop.

4. MECHANISM IMPLEMENTATIONS

The total expected revenue of the operator is the sum of expected revenues (2) from transactions within the considered planning horizon. However, as we have stated in the previous section, costs of transactions are usually not known to the operator before the previous transactions are completed. This is because these costs, being caused by the need for relocating vehicles between stations, depend on both the instantaneous vehicle supply at all stations, and on the origin station of the current client, as well as on the types of the next clients to come. The supply, in turn, is influenced by the outcomes of previous transactions, that depend on prices that were previously posted. The actual costs of relocation clearly depend also on the relocation schedule used by the system's operator. Since the exact costs are hard to determine in advance, we assume in our model that they are provided in an online fashion for incoming transactions.

In the following, we propose two specific mechanisms for adaptively computing prices. They vary in the information they use about individual clients, and both are parameterised by the costs of the current transaction. This is followed by a discussion of how we estimate relocation costs in practice.

4.1 One-stage mechanism

Our first adaptive pricing mechanism uses only general probabilistic knowledge about the next client, as given by f_i. In this mechanism, the client receives a set of posted prices, one for each destination. Prices are computed using the solution of the monopolist pricing problem from the previous section, based on the current cost estimates. The mechanism is myopic, in the sense that it computes prices based on a relocation cost heuristic which depends on the state of the system. This assumption is justified by the difficulty of reliably estimating costs far into the future, but it also allows for applying an arbitrary relocation policy in the system (which is likely to have a major influence on the costs). The way that prices are computed does not depend on the relocation scheme. Instead, the prices simply adapt to the current state of the system, which is partially influenced by the underlying relocation scheme. It is assumed that we have access to the estimated probability distributions of clients' types, based on the payments made so far for each origin–destination pair of stations. These parameters can be updated in an on-line fashion as well, as new clients participate in the system.

We call this mechanism one-stage, since the client makes only a single choice of their return destination. A solution of the monopolist pricing problem contains the information of revenue-maximising prices, given the distribution of the client's willingness to pay and the projected costs of delivering the services. Depending on the type of probability distribution we have available, we would choose a different solution method. If the distribution is discrete over selected valuation vectors $\mathbf{x}$, we are able to use either the MIP (19)–(25) or (26)–(28). If the distribution is continuous, and we are able to sample from it (using, e.g., a Monte Carlo method), then we assume that we have pre-sampled points $\{\mathbf{x}_i\}_{i=1}^N$, and we would use LP (11)–(14). All of these scale to problems of realistic sizes.

4.2 Two-stage Mechanism

The revenue generated by a pricing mechanism can be significantly improved if more information about a client's type is available. In on-demand mobility systems, we are dealing with many different classes of clients (e.g., commuters, shoppers, leisure users), whose respective distributions of valuations could be substantially different. However, in the one-stage mechanism, the general distribution f_i is used to compute prices, and so they are adjusted to an "expected" client, where any potential class-related information is lost. A potential improvement can be achieved if we are allowed to extract the "class" information of a client. A reasonable option to model this (and one that can potentially be learnt from historical data) is to use type distributions $f_{i,k}$, which define probabilities of valuations of a client departing from station i and with a most preferred destination k (in the absence of prices).

This can be realised in several ways. A simple approach is to ask the client to declare her destination before showing the prices. Such a mechanism, however, will be susceptible to strategic manipulations. A client may choose to misreport the true highest-valued destination, when the cost associated with that station is higher than the cost associated with another station. The mechanism may then post a lower price for the latter station, that is far below the true client's willingness to pay for it. However, in order for the client to act strategically, she would need to know the actual costs associated with each station, which is private information of the mechanism. A slightly better approach is to offer price discounts for switching after a client has already chosen a return destination. An initial client's choice is made given the prices computed by the one-stage mechanism, described in the previous subsection. Such commitment reveals the client's class k, as the client does not know if she will be offered different prices. However, at some later point in time, if the distribution of vehicles has changed, a client would be presented with a new set of prices. They may then choose a lower price, on the condition that they agree to switch to a different destination.

We call this mechanism two-stage, as it involves two decisions on the client side to be passed to the mechanism: first, the declaration of a preferred destination station, either declared directly (which is what we assume in our experiments in Section 5) or based on an initial set of prices, with the opportunity to switch to another destination, once the final prices are shown. As previously noted this mechanism is not always incentive compatible. However, instead of asking directly, such data could be derived from the history of use.

4.3 Modelling Relocation Costs

In this section, we now explore how to model relocation costs given the current state of the system. The problem of determining an optimal relocation policy can, in its simplest form, be modelled as a classic transportation problem [21]. However, this does not consider the impact of slight imbalances in the system, before any relocations are required. Hence, in order to estimate the costs used in our mechanisms, we use the following approach. Let $n_{i,t}$ be the current occupancy at station i, i.e., the current number of cars at the station. For each station i, there is a target occupancy interval, defined by an upper bound $\overline{n}_i^*$ and a lower bound $\underline{n}_i^*$.

For example, it may be desirable to have about 50% of each station capacity used most of the time (in order to keep room for incoming cars, as well as to provide high availability). We would then set $\underline{n}_i^* = \overline{n}_i^* = \frac{1}{2}n_{i,\max}$, where $n_{i,\max}$ is the ith station's capacity (number of installed charging slots). The operator might want also to set higher $\underline{n}_i^*$ for popular departure stations, and lower $\overline{n}_i^*$ for popular destination stations.

Let us denote by $\tilde{C}(i,j,t)$ the cost[2] of relocating one vehicle from station i to station j. We define this cost as follows:

$$\tilde{C}(i,j,t) = C_{\text{source}}(i,t) + C_{\text{dest}}(j,t),$$

where:

$$C_{\text{source}}(i,t) = \begin{cases} -\gamma & \text{if} n_{i,t} > \overline{n}_i^*, \\ \delta & \text{if} n_{i,t} < \underline{n}_i^*, \\ 0 & \text{otherwise}, \end{cases}$$

$$C_{\text{dest}}(j,t) = \begin{cases} -\gamma & \text{if} n_{j,t} < \underline{n}_j^*, \\ \delta & \text{if} n_{j,t} > \overline{n}_j^*, \\ 0 & \text{otherwise}, \end{cases}$$

Parameters γ and δ need to be selected appropriately. We address this problem in Section 5.

It should be noted that, in practice, a good relocation policy requires considering many factors, which are not addressed in this paper. These factors include scheduling relocations for appropriate time periods and assigning staff members to selected locations [23], [12]. However, our pricing mechanism is not specific to a particular relocation policy and so can co-exist with arbitrary policies.

5. EXPERIMENTAL EVALUATION

We have validated our model in a series of simulations, using data from an on-demand mobility system operating in the French city Grenoble. The data covers five months (October 2015 – February 2016) of service provided at 28 rental stations, as shown in Figure 1, and it consists of 4,139 transactions. More specifically, we used this data to estimate the parameters of a probabilistic model of the clients' types. The goal of the experiments was to compare the generated income (i.e., the money collected from clients for the rental service) and number of performed relocations, using different pricing mechanisms, including a flat pricing mechanism that is current in use in Grenoble.

[2]We then have $C(j,t) = \tilde{C}(\tau(t),j,t)$.

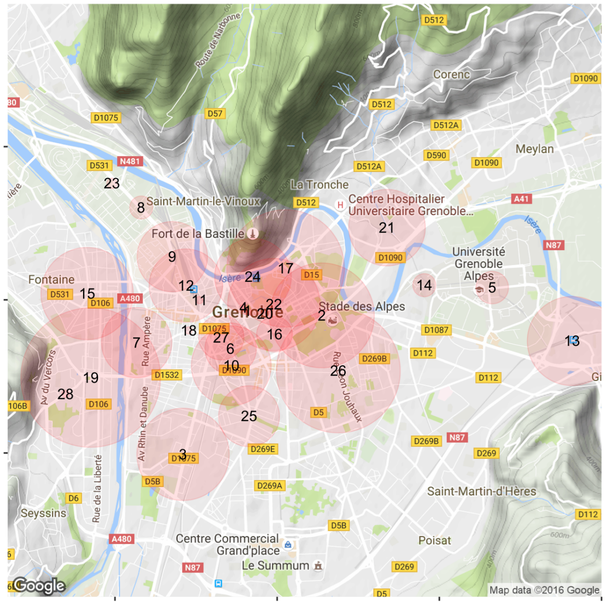

Figure 1: Map of Grenoble with stations of the existing car sharing system marked by numbers from 1 to 28. The radius of each circle indicates the demand level for the service, as estimated from past data.

5.1 Distribution of Client Types

We used a mixture of multivariate normal distributions for modelling the clients' types generated in our simulations. Each of the mixture's component distributions corresponds to a "class" of clients. Each class is associated with a certain destination, which is valued the most by the members of that class. We index classes by origin-destination pair (i, k). For each class (i, k) we define a multivariate normal distribution with means $\boldsymbol{\mu}^{(i,k)} = (\mu_1^{(i,k)}, \ldots, \mu_n^{(i,k)})$, where each $\mu_j^{(i,k)}$ is the mean valuation of a service that departs from the ith station and terminates at the jth station. For each class, we set a maximum mean valuation $\mu_k^{(i,k)} := \hat{\mu}_{ik}$ (estimated from the "flat price" scheme, see below), as the kth station is valued the most by the members of that class of clients, on average. The remaining mean valuations $\mu_j^{(i,k)}$, for $j \neq k$, diminish according to the distance d_{jk} between stations j and k:

$$\mu_j^{(i,k)} = \hat{\mu}_{ik} \exp\left(-d_{jk}^2/\alpha\right), \tag{29}$$

where α is the parameter controlling how the valuation diminishes over distance. The greater the value of α, the less the valuation is diminished, i.e., clients are more willing to switch their destinations. Figure 2 shows some examples for how the mean valuation diminishes for various stations that are at increasing distances from a preferred destination station (10 in this example) and for different values of α.

This way we obtain a multivariate normal distribution with mean $\boldsymbol{\mu}^{(i,k)}$ for each class. We denote by $f_{i,k}$ the density function of each such distribution. The multivariate normal model has distribution given by $f_i = \sum_{k=1}^{M} b_{i,k} f_{i,k}$, where $b_{i,k}$ is the probability that a random client belongs to class (i, k) (estimated from the data of departures). The journey durations were modelled by random variables following an exponential distribution, with expected duration parameters estimated from data.

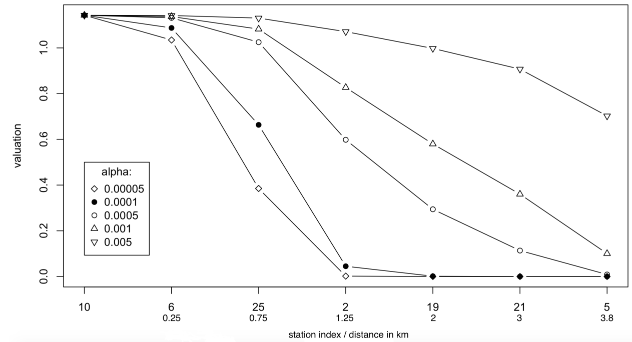

Figure 2: Assumed change in valuations of alternative destinations for random client who wants to reach station with index 10, assuming different values of parameter α in (29) (see Figure 1 for locations of the stations).

In order to estimate an average price that clients are willing to pay for the services, we used records of real transactions in the Grenoble car sharing system. These datasets contain transactions conducted with a "flat price" scheme where of €1 per unit of rental time (15 minutes), except the first unit of rental time, which was charged €2. For each pair of origin and destination stations, we calculated the average rental time, in order to obtain the mean price for a unit of rental time of each service. In our study we denote this price by $\hat{\mu}_{ik}$, for each service with origin-destination pair (i, k). Note that this could be seen as a pessimistic estimate of valuations, since real valuations might be higher than the fixed price. As a result, the potential income and reduction in relocations when using our adaptive pricing mechanisms may be higher in practice.

For solving the optimal pricing problems, we used MIP (26)–(28). For the input data we used a discrete probability distribution $\rho_{(i,j)} = \mathbf{Pr}[\text{client books journey from } i \text{ to } j]$. In the one-stage mechanism we used the overall distribution of client types from a given station i, while in the two-stage mechanism we used the conditional distribution, given that we know the client's preferred destination.

5.2 Experimental Setup

In the simulations we keep track of the number of vehicles available at every station between each journey. We start each run from the fully balanced state, where every station has 50% of capacity filled. The simulation then consists of the following steps:

1. Sample a random client's origin station i, type $\mathbf{x}$, and journey duration.
2. Determine the estimated relocation costs associated with each destination choice j.
3. Compute prices ϕ_j^* for each destination j, based on the relocation costs.
4. Given the prices, compute client's utilities $u_j = x_j - \phi_j^*$ of using each destination j.
5. Select destination j with maximum utility u_j, or cancel the journey if all utilities are less than zero.
6. Perform the necessary relocations, whenever the capacity of station j is exceeded, or the capacity of the station i reaches zero.

Table 1: Experiment with no relocations.

"flat" prices		1-stage		2-stage	
cancelled	income	cancelled	income	cancelled	income
52 %	530.4	39 %	988.46	37 %	864.59

The above scheme is repeated for a selected number of sampled clients. Note that at a given time there may be many concurrent journeys taking place. The relocation procedure is triggered whenever the capacity of any station is violated. The pricing mechanism, described in the Section 4, is applied as a part of step 3 in the above scheme.

For all settings we consider, we test three different pricing mechanisms: the "flat" pricing mechanism (originally used in the existing car sharing system), and our one-stage and two-stage mechanisms, which use the solutions of the optimal pricing problem, as described in Section 4. Solutions of generated MIPs were obtained using CPLEX 12.6 software, all results are averaged over 30 runs of 100 consecutive transactions and they are shown with 95% confidence intervals (similar results are obtained for experiments with more consecutive transactions).

5.3 Experimental Results

First, we consider a setting with no relocations at all (and we fix $\alpha = 0.0001$, $\overline{n}_i^* = 0.6$, $\underline{n}_i^* = 0.4$, $\gamma = \delta = 1.0$). Here, vehicles are only moved by the clients, and any transactions from an empty origin station are cancelled,[3] while full destination stations are not available. Table 1 shows the results in this setting, demonstrating that our adaptive pricing mechanisms result in significantly fewer cancelled journeys, i.e., more balanced vehicle availability, and in a higher income. This is because they incentivise clients to balance the system. Note that the one and two-stage mechanisms provide different results since they are using different valuation distributions resulting in different prices and user choices. Interestingly, the two-stage mechanism sometimes results in lower revenue. This is because the system is myopic and sets prices to maximise the revenue for each transaction, without considering how this might affect future transactions.

Next, we consider more settings with relocations. Here, we explore different parameter values for α (controlling the willingness to switch), $\overline{n}_i^*$, $\underline{n}_i^*$ (target occupancy levels), γ and δ (modelling the relocation costs). Specifically, Figures 3–5 show settings with $\overline{n}_i^* = 0.6$, $\underline{n}_i^* = 0.4$, i.e., the target occupancy level is between 40%–60% full, and we vary the parameters of our cost estimation function from $\gamma = \delta = 0.7$ to $\gamma = \delta = 2.0$. In Figure 6, we show a setting with a wider target occupancy level, $\overline{n}_i^* = 0.75$, $\underline{n}_i^* = 0.25$ and $\gamma = \delta = 1.0$.

Overall, it is apparent that all settings related to the cost function estimation lead to similar performance characteristics, indicating that the mechanisms are robust to choosing different parameter values. Otherwise, several general trends emerge from the figures.

First, in almost all cases, relocations are reduced significantly by using our adaptive pricing mechanisms. This is because our pricing mechanisms take into account the estimated costs of relocations and, as a result, increase the prices of stations that are outside their target occupancy levels. However, when clients are almost indifferent regarding their destination (the highest setting of α in our simulations), they would not necessarily be redirected properly to alternative stations in cases of deadheading, which is noticeable in a slight increase of the number of relocations.

Second, average income from our mechanisms rises significantly as clients are more flexible in their final destination. This is due to several reasons: clients are more likely to be re-directed to stations that are beneficial to the system, but the system is also generally able to post higher prices, as it is likely that clients will still be able to find a destination station that results in non-negative utility. For very high values of α, our adaptive mechanisms offer an over 200% improvement in income over the flat pricing scheme (albeit at a slight increase in relocations). In more realistic settings, e.g., with $\alpha = 0.0005$, our adaptive schemes can result in an increase in average income of about 64%, while simultaneously reducing relocations by up to 36%. Note, however, that when clients are very inflexible ($\alpha \leq 0.00005$), our adaptive mechanisms perform worse than flat pricing. This is likely due to the estimated (rather than accurate) cost function and the myopic decisions of our pricing mechanisms, which here result in sub-optimal prices. However, these inaccuracies are compensated by the benefits of our mechanism as the clients' flexibility increases.

Third, the two-stage mechanism generally results in significantly higher income than the one-stage mechanism. This is not surprising, given that it uses more information about the clients, and is therefore able to post higher prices for the clients' preferred stations, thus extracting more revenue.

In conclusion, our mechanisms work well in a variety of settings and for problems of realistic sizes. However, they are sensitive to the clients' willingness to switch, i.e., on the distribution of valuations of alternative destination stations. Yet, as soon as clients display some flexibility in their destination stations, there is a clear benefit to using adaptive pricing mechanisms.

6. CONCLUSIONS

In this paper, we proposed two revenue-maximising adaptive pricing mechanisms for on-demand mobility systems, with different assumptions regarding our probabilistic knowledge of the clients' valuations. These are the first mechanisms that combine both the problem of setting prices to maximise revenue with the problem of dealing with costly relocations. We validated the mechanisms in a series of simulations based on real data, and we showed that revenue can be increased significantly as compared to the simple "flat" pricing scheme. The mechanism at the same time helps to diminish the deadheading effect, reducing the number of relocations by up to 36%.

Future work will focus on methods for more accurately and non-myopically predicting relocation costs and future demands. We will also consider how to learn the distribution of clients' types from historical bookings. Moreover, we intend to perform user studies to elicit a client's willingness to switch to a nearby station to improve the user model. Finally, we will extend our mechanisms to providing incentives not only for switching to different destinations, but also to different origin stations or to postpone a journey to a later time.

[3]Note that the system never cancels any bookings already made. Cancellations refer to journeys where a client approaches the system, but does not book due to either unavailability or negative utility.

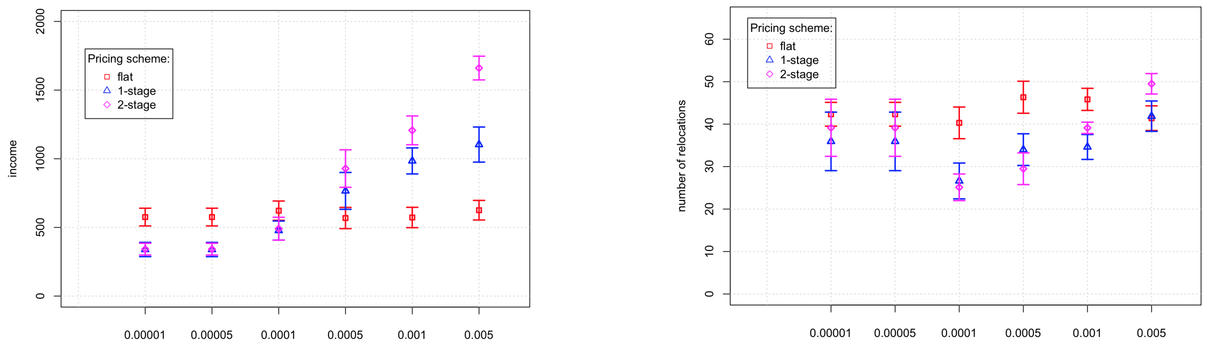

Figure 3: Performance of mechanisms for different values of parameter α, with $\overline{n}_i^* = 0.6n_{i,\max}$, $\underline{n}_i^* = 0.4n_{i,\max}$, $\gamma = \delta = 0.7$.

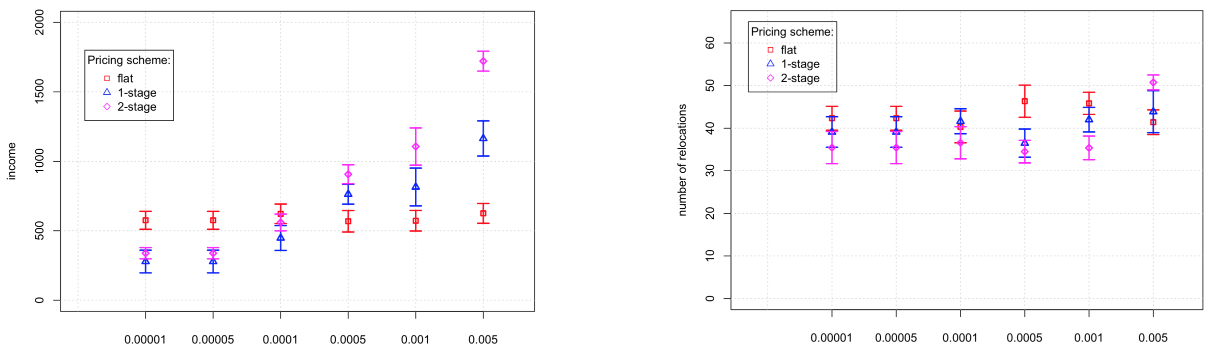

Figure 4: Performance of mechanisms for different values of parameter α, with $\overline{n}_i^* = 0.6n_{i,\max}$, $\underline{n}_i^* = 0.4n_{i,\max}$, $\gamma = \delta = 1.0$.

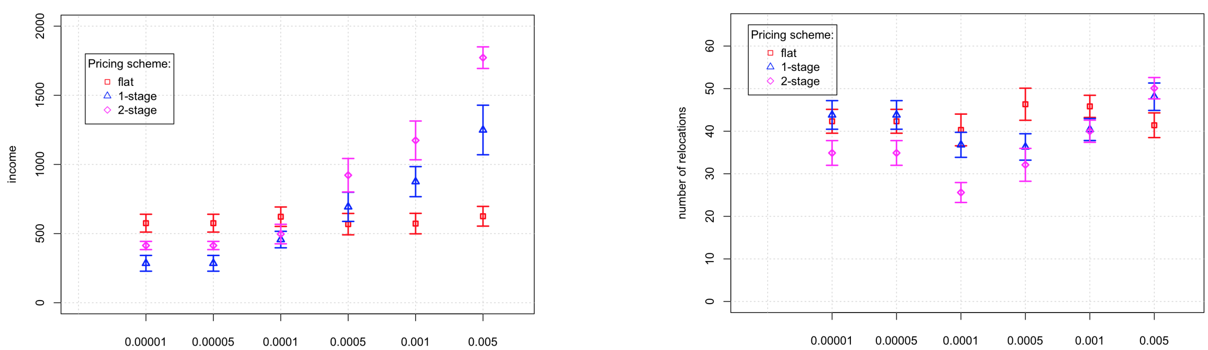

Figure 5: Performance of mechanisms for different values of parameter α, with $\overline{n}_i^* = 0.6n_{i,\max}$, $\underline{n}_i^* = 0.4n_{i,\max}$, $\gamma = \delta = 2.0$.

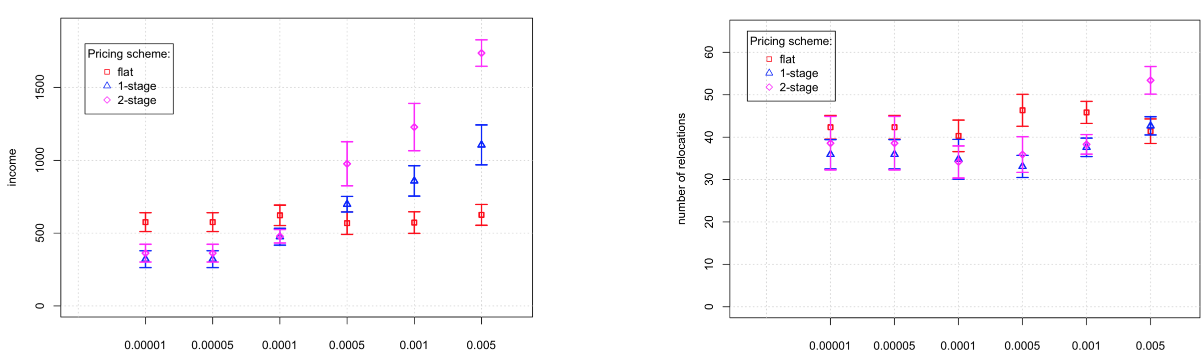

Figure 6: Performance of mechanisms for different values of parameter α, with $\overline{n}_i^* = 0.75n_{i,\max}$, $\underline{n}_i^* = 0.25n_{i,\max}$, $\gamma = \delta = 1.0$.

REFERENCES

[1] W. P. Adams and R. J. Forrester. A simple recipe for concise mixed 0-1 linearizations. *Operations Research Letters*, 33(1):55–61, 2005.

[2] S. Basov. Hamiltonian approach to multi-dimensional screening. *Journal of Mathematical Economics*, 36(1):77–94, 2001.

[3] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

[4] P. Briest, S. Chawla, R. Kleinberg, and S. M. Weinberg. Pricing lotteries. *Journal of Economic Theory*, 156:144–174, 2015.

[5] Y. Cai, C. Daskalakis, and S. M. Weinberg. Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In *53rd Annual Symposium on Foundations of Computer Science*, pages 130–139, 2012.

[6] A. Carlier, A. Munier-Kordon, and W. Klaudel. Optimization of a one-way carsharing system with relocation operations. In *10th International Conference on Modeling, Optimization and SIMulation MOSIM 2014*, 2014.

[7] G. Carlier, T. Lachand-Robert, and B. Maury. A numerical approach to variational problems subject to convexity constraint. *Numerische Mathematik*, 88(2):299–318, 2001.

[8] D. Chemla, F. Meunier, and R. W. Calvo. Bike sharing systems: Solving the static rebalancing problem. *Discrete Optimization*, 10(2):120–146, 2013.

[9] X. Chen, I. Diakonikolas, D. Paparas, X. Sun, and M. Yannakakis. The complexity of optimal multidimensional pricing. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1319–1328, 2014.

[10] A. d'Aspremont. *Shape Constrained Optimization*. PhD thesis, Stanford University, 2004.

[11] S. Hart and N. Nisan. Approximate revenue maximization with multiple items. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 656–656, 2012.

[12] S. Jian, D. Rey, and V. Dixit. Dynamic optimal vehicle relocation in carshare systems. In *Transportation Research Board 95th Annual Meeting*, number 16-4592, 2016.

[13] D. Jorge, G. H. Correia, and C. Barnhart. Comparing optimal relocation operations with simulated relocation policies in one-way carsharing systems. *IEEE Transactions on Intelligent Transportation Systems*, 15(4):1667–1675, 2014.

[14] R. Katzev. Car sharing: A new approach to urban transportation problems. *Analyses of Social Issues and Public Policy*, 3(1):65–86, 2003.

[15] A. M. Manelli and D. R. Vincent. Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic Theory*, 137(1):153–185, 2007.

[16] Q. Mérigot and E. Oudet. Handling convexity-like constraints in variational problems. *SIAM Journal on Numerical Analysis*, 52(5):2466–2487, 2014.

[17] M. Mussa and S. Rosen. Monopoly and product quality. *Journal of Economic Theory*, 18(2):301–317, 1978.

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

[19] J. Pfrommer, J. Warrington, G. Schildbach, and M. Morari. Dynamic vehicle redistribution and online price incentives in shared mobility systems. *IEEE Transactions on Intelligent Transportation Systems*, 15(4):1567–1578, 2014.

[20] J.-C. Rochet and P. Chone. Ironing, sweeping, and multidimensional screening. *Econometrica*, 66(4):783–826, 1998.

[21] A. Schrijver. *Combinatorial optimization: polyhedra and efficiency*. Springer Science & Business Media, 2002.

[22] A. Singla, M. Santoni, G. Bartók, P. Mukerji, M. Meenen, and A. Krause. Incentivizing users for balancing bike sharing systems. In *AAAI*, pages 723–729, 2015.

[23] S. Weikl and K. Bogenberger. Relocation strategies and algorithms for free-floating car sharing systems. *IEEE Intelligent Transportation Systems Magazine*, 5(4):100–111, 2013.