



HAL
open science

Hybrid Mechanisms for On-Demand Transport

Malcolm Egan, Nir Oren, Michal Jakob

► **To cite this version:**

Malcolm Egan, Nir Oren, Michal Jakob. Hybrid Mechanisms for On-Demand Transport. IEEE Transactions on Intelligent Transportation Systems, 2019, pp.1-13. 10.1109/TITS.2018.2886579. hal-01950684

HAL Id: hal-01950684

<https://inria.hal.science/hal-01950684>

Submitted on 11 Dec 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Hybrid Mechanisms for On-Demand Transport

Malcolm Egan, Nir Oren and Michal Jakob

Abstract—Market mechanisms are now playing a key role in allocating and pricing on-demand transportation services. In practice, most such services use posted-price mechanisms, where both passengers and drivers are offered a journey price which they can accept or reject. However, providers such as Liftago and GrabTaxi have begun to adopt a mechanism whereby auctions are used to price drivers. These latter mechanisms are neither posted-price nor classical double auctions, and can instead be considered a hybrid mechanism. In this paper, we describe and study the properties of a novel hybrid on-demand transport mechanism. As these mechanisms require knowledge of passenger demand, we analyze the data-profit tradeoff as well as how passenger and driver preferences influence mechanism performance. We show that the revenue loss for the provider scales with $\sqrt{n \log n}$ for n passenger requests under a multi-armed bandit learning algorithm with beta distributed preferences. We also investigate the effect of subsidies on both profit and the number of successful journeys allocated by the mechanism, comparing these with a posted-price mechanism, showing improvements in profit with a comparable number of successful requests.

Index Terms—On-demand transport, taxis, pricing.

I. INTRODUCTION

An important component of the transportation infrastructure in any city is on-demand or demand-responsive transport. As the name suggests, the timetable and route of these services are not fixed, contrasting with the operation of most trains and buses. As such, on-demand transport (ODT) is useful for last mile transportation—from hubs to residences—and also for areas with low populations or passengers with special needs (e.g., the elderly and disabled) [1].

ODT services can be not-for-profit, municipality funded, or profit driven. Although all types of services are often present in a city, they face different challenges. In particular, the routing and scheduling problems (where to go and when to pick up each passenger) in—typically small-scale—not-for-profit services are typically solved in a centralized manner and aim to minimize costs [1]–[9]. This is due to long-term revenue guarantees from the municipality or local government.

On the other hand, profit-driven ODT services have traditionally been dominated by taxis. Here, revenue from each passenger must be accounted for in routing and scheduling. Moreover, the taxi company and drivers have a restricted ability to optimize pricing, due to regulation [10]–[12]. As such, factors such as how each passenger is priced, how much each driver is paid, where taxis wait, and which passengers they

pick-up play an important role both for minimizing passenger waiting times and fares, and for improving the profits of each driver and the taxi company. Optimizing pricing, scheduling, and routing in this environment is challenging, although significant progress has been made under the assumption of slowly time-varying pricing, where the problem can be reduced to a minimum cost scheduling and routing problem [5], [6], [10]–[18].

Recently, ODT providers such as UberX¹ have been using machine learning techniques to improve routing, scheduling and pricing. In contrast with taxi companies, these ODT providers exploit market mechanisms to price and allocate passengers and drivers. In this case, studied in [19]–[24], up-front prices can be set using information gleaned from large amounts of data (e.g., collected via smartphone apps), and can be dynamically adjusted to account for time, supply and demand, and pick-up and drop-off locations—a significant departure from the static approaches used by traditional taxis.

The use of market mechanisms for pricing and allocation [25]–[27] is raising new challenges for the design of ODT systems. In particular, modern ODT providers are using a range of market mechanisms, not limited to the posted-price approach [16], [28]–[31] currently adopted by Uber, where passengers and drivers are offered prices that they can accept or reject. For example, Liftago² has adopted a hybrid mechanism where the payment received by drivers is determined by an auction and the price paid by passengers is determined by a posted-price mechanism. At present, it is not clear how to determine which mechanism is best suited to each city and scale of provider, beyond financially risky real-world experimentation.

There are two main approaches to study market design for on-demand transport. The first relies on tools from economics in the form of equilibrium analysis. In particular, equilibrium conditions have been derived and solutions studied in work such as [11], [32]–[34]. The main conclusions from this analysis are structural properties of the relationship between market supply and demand in the aggregate; that is, equilibrium behavior of passengers and drivers. While this analysis yields useful high-level insights for regulators, it is not presently possible to use these frameworks to optimize the market mechanism itself for individual providers. However, we note that the simulation-based method to study stochastic user equilibria accounting for within-day operating policies in [34] appears promising in this respect.

The second approach is to develop and compare particular market mechanisms over fixed periods of time, using agent-based modeling approaches. These mechanisms usually take the form of posted-price approaches, and here, several

M. Egan is with the CITI Lab, 69621 Villeurbanne, France which is a joint laboratory in Université de Lyon, INSA-Lyon, and INRIA (e-mail: malcom.egan@inria.fr). During the preparation of this paper he was also with the Faculty of Electrical Engineering, Czech Technical University in Prague. N. Oren is with the Department of Computing Science, University of Aberdeen (e-mail: n.oren@abdn.ac.uk). M. Jakob is with the Faculty of Electrical Engineering, Czech Technical University in Prague (e-mail: michal.jakob@agents.fel.cvut.cz).

¹<https://www.uber.com/>

²<https://www.liftago.com/>

mechanisms have been proposed [10], [13], [14], [16], [19]–[24], [35], [36]. Typically, these approaches either assume passengers will accept all journeys at a given price or rely on passengers and drivers truthfully revealing their preferences in the form of maximum prices, waiting times and other factors such as vehicle type (e.g., [22]–[24]). The assumption that passengers are willing to reveal the maximum price they are prepared to pay is problematic, as often passenger preferences are not well-defined as this is not the way ODTs are currently used. Nevertheless, market mechanisms that assume that passengers and drivers have private price preferences is currently limited to [19]–[21], [35]–[37].

In this paper, we introduce a new hybrid mechanism for ODT services to resolve the tension between supply and demand, where passengers are offered a price they can accept or reject and driver payments and allocations are determined by an auction. That is, like [19]–[21], [35]–[37], passengers and drivers are not required to perfectly reveal their price preferences. Moreover, we do not assume that passengers are prepared to bid (e.g., [19], [20], [37]) nor that drivers receive a salary (as in [21], [35], [36]).

There are two key questions that we address: (1) *how does the availability of data affect the price that should be offered by the service?* and (2) *how do passenger and driver cost and profit preferences affect the performance of the mechanism?* The first question concerns data availability, taking into account the possibility of offering cheaper fares to obtain more information about passenger pricing preferences. It has a direct influence on the effectiveness of the machine learning techniques used to optimize pricing and allocation within the mechanism. This also includes algorithms that are used to estimate demand and trip information [15], [38]–[40] as well as market formation [41], which determines which passengers can be allocated to each driver. In contrast, the second question considers properties which the mechanism cannot control, as such driver and passenger preferences are determined *a priori* due to social, cultural and geographic factors; including the presence of consumer habits [42].

To evaluate our hybrid mechanism, we develop an agent-based model which captures preferences of individual passengers with respect to journey prices, and drivers with respect to their minimum profit requirements per journey. Our model can be viewed as an extension of [19], [36] where both passengers and drivers are self-interested. We formally demonstrate several properties of the hybrid mechanism, including truthfulness of passenger and driver reports, tradeoffs between profit and historical data available to the provider, and the impact of preferences on provider profit.

We also propose two variations on the hybrid mechanism to improve the proportion of passengers that are transported by a driver (which we refer to as the mechanism’s *efficiency*). That is, both the passenger accepts the price offered by the mechanism and there is a driver willing to transport the passenger. These variations are inspired by the social planning perspective in auctions [29] and incorporate variations in the prices passengers are prepared to pay and subsidies to incentivize drivers to transport each passenger.

Ultimately, the performance of a mechanism for an ODT

system depends on the long-term behavior of the mechanism, accounting for the system’s spatial-temporal profile. In particular, the spatial distribution of passenger requests, where drivers without passengers wait, and the passenger drop-off locations all affect key system metrics, namely the profit of the provider, the prices passengers pay, and the daily profits of drivers. As such, we provide a system-level evaluation of our hybrid mechanism and compare it to existing mechanisms for taxis as well as a two-sided posted-price mechanism inspired by the approach adopted by Uber. We show that the hybrid mechanism can significantly outperform a two-sided posted-price mechanism in terms of the profit, while maintaining comparable success rates. We also investigate provider subsidies, which further improve the success rate at the cost of a lower profit.

The remainder of this paper is organized as follows. In Section II, we detail our system model. In Section III, we apply this model to existing ODT services, providing us with a basis for comparison. Section IV, introduces the proposed hybrid mechanism, and in Section V, we analyze the hybrid mechanism to investigate tradeoffs between the quantity of data and imperfect knowledge of system parameters on the profit. In Section VI, two variations on the hybrid mechanism to vary the tradeoff between profit and proportion of passengers successfully transported are proposed. In Section VII, we explain our simulation setup and investigate the effect of system parameters on profit and success rate for the hybrid mechanism, a two-sided posted-price mechanism and the traditional dispatcher taxi model. In Section VIII, we summarize our work and consider future extensions.

II. SYSTEM MODEL

In this section, we develop an agent-based model for the ODT system which forms a basis for our hybrid mechanism in Section IV. We consider an ODT system consisting of a provider (responsible for pricing passengers and matching these with drivers), d vehicles/drivers who each service a single passenger at a time (i.e., ridesharing is not supported), and n passengers that enter the system at different times.

A. Passenger Agent Model

A new passenger enters the system when she makes a journey request to the ODT provider, which consists of a pick-up time, pick-up location and drop-off location. We assume that each passenger wishes to be picked up as soon as possible after she makes her request, with maximum delay of Δ minutes. The pick-up and drop-off locations are represented by elements from the set of vertices V in the directed graph $G = (V, E)$, which models the underlying road network. The set of edges E in the graph represent direct routes between locations in V that can be traversed by the vehicles. Associated to each edge $e \in E$ are a start and end location $u, w \in V$ respectively; a cost $c_e \in [0, \infty)$ for a vehicle to traverse edge e ; and a traversal time $\tau_e \in \mathbb{Z}_+$.

The cost of traversing an edge is incurred due to factors such as fuel consumption and vehicle wear and tear. We assume that the edge cost c_e and traversal time τ_e are computed offline by

solving the shortest path problem between u and w in the underlying road network.

Each passenger has a maximum price she is prepared to pay for her journey, which is private and unknown to the provider. The maximum price is determined by the *maximum price-rate* that the passenger is prepared to pay (in, for example, euros per kilometer) and the distance the passenger seeks to travel. Let R_i denote the distance that passenger i seeks to travel, then the total price she is willing to pay for her journey is $p_{i,\max} = r_{i,\max}R_i$, where $r_{i,\max}$ is the passenger's maximum price-rate. Since R_i is known to both passenger i and the provider, it follows that the passenger's preferences are completely characterized by the maximum price-rate $r_{i,\max}$. As such, if a passenger is offered a journey with price-rate $r_i \leq r_{i,\max}$ it will be accepted and if $r_i > r_{i,\max}$ the offer will be rejected (assuming that they can be picked up for the journey within Δ minutes).

We model the maximum price-rate for each passenger as an independently and identically distributed (i.i.d) random variable, which captures differences in how much each passenger is prepared to pay. Following [36], we assume that $r_{i,\max}$ is beta distributed with probability density function (pdf)

$$f_{r_{\max}}(x) = \frac{1}{\rho_{\max}B(\alpha_r, \beta_r)} \left(\frac{x}{\rho_{\max}}\right)^{\alpha_r-1} \left(1 - \frac{x}{\rho_{\max}}\right)^{\beta_r-1} \quad (1)$$

where $B(\alpha_r, \beta_r) = \frac{\Gamma(\alpha_r)\Gamma(\beta_r)}{\Gamma(\alpha_r+\beta_r)}$ is the beta function and $\Gamma(\cdot)$ is the gamma function. As such, the maximum price-rate is influenced by the three parameters $\alpha_r, \beta_r, \rho_{\max}$, where α_r, β_r affect the shape of the pdf and ρ_{\max} determines the support of the pdf $[0, \rho_{\max}]$.

The choice of the beta distribution to model passenger price preferences can be justified by the flexibility arising from the three parameters and its tractability. We also remark that the beta distribution has also been applied to model revenue in economics (see, e.g., [43]) and that the parameters can be readily estimated using via techniques such as maximum likelihood and method of moments [44]. Variations in the parameters can arise due to changes in perceived supply and demand (i.e., the number of available taxis and the number of passenger requests) or social, cultural and spatial-temporal differences (such as time of day) between passengers.

B. Driver Agent Model

Drivers also have preferences for the minimum payment they are prepared to receive for transporting passengers. Driver j 's profit from transporting passenger i is given by

$$S_{j,i} = (1 - \eta_j)r_iR_i - c_{j,i} \quad (2)$$

where η_j is the proportion of the passenger's payment sent to the provider and $c_{j,i}$ is the cost of transportation. In particular,

$$c_{j,i} = \kappa_j(R_i + R_{j,i}) \quad (3)$$

where κ_j is the cost per kilometer incurred by the driver and $R_{j,i}$ denotes the distance from driver j 's initial location to passenger i 's pick-up location.

Let $S_{j,\min}$ denote the minimum profit per minute driver j is prepared to receive for a journey. We assume that driver j is only willing to transport passenger i if

$$S_{j,i} \geq S_{j,\min}\tau_{j,i} \quad (4)$$

where $\tau_{j,i}$ is the duration of travel from driver j 's initial location and passenger i 's drop-off location. As such, both the journey distance and duration are accounted for by the driver. As for each passenger's price-rate, we assume that $S_{j,\min}$ is an i.i.d beta distributed random variable (albeit with different parameters) with pdf given by

$$f_{S_{\min}}(x) = \frac{1}{\sigma_{\max}B(\alpha_d, \beta_d)} \left(\frac{x}{\sigma_{\max}}\right)^{\alpha_d-1} \left(1 - \frac{x}{\sigma_{\max}}\right)^{\beta_d-1} \quad (5)$$

where α_d, β_d are the shape parameters and the support is $[0, \sigma_{\max}]$. These parameters are affected by variations in perceived supply and demand or spatial-temporal differences between drivers.

We assume that the provider is aware of which drivers are available and which are busy at all times. This information can either be obtained directly from driver reports or a notification from passengers being transported.

C. Provider Agent Model

The ODT provider runs the mechanism that prices and allocates passengers and drivers. In return, the provider receives a proportion of each passenger's payment. In particular, the provider receives $\eta_j r_i R_i$ for passenger i when served by driver j , which corresponds to the fraction of the passenger's payment not given to driver j in exchange for transportation.

To perform pricing and allocation, the provider exploits statistical knowledge of passenger preferences for the price-rate and driver preferences for their minimum profit. This statistical knowledge corresponds to the parameters $\alpha_r, \beta_r, \rho_{\max}$ and $\alpha_d, \beta_d, \sigma_{\max}$ which must be learned from historical transactions. When learned, such parameters will typically be imperfect due to, for example, limited data or non-stationarity of passenger preferences. To distinguish between the actual pdfs for the preferences, $f_{r_{\max}}, f_{S_{\min}}$, and the statistical knowledge of the provider, we denote the estimated pdfs as $\hat{f}_{r_{\max}}$ with parameters $\hat{\alpha}_r, \hat{\beta}_r, \hat{\rho}_{\max}$ and $\hat{f}_{S_{\min}}$ with parameters $\hat{\alpha}_d, \hat{\beta}_d, \hat{\sigma}_{\max}$, respectively.

In the following sections, we overview the evolution of existing ODT mechanisms and introduce a hybrid mechanism which exploits techniques from the theory of online mechanisms to obtain parameters for the preference pdfs.

III. MODELING EXISTING ODT MECHANISMS

In order to motivate the design of our hybrid mechanism, we overview existing mechanisms in ODT. We formalize each mechanism within the context of the model detailed in Section II, which allows us to form a basis for comparison with our proposed hybrid mechanism.

Historically, the dispatcher and Hackney carriage mechanisms were the first mechanisms used to price and allocate passengers in ODT systems. In the absence of smartphone

apps, it was difficult for providers, drivers and passengers to effectively characterize the spatial-temporal properties of the system and to communicate with each other.

In the Hackney carriage mechanism, passengers communicate with drivers by hailing from the side of the road, which only allowed passengers to select from drivers within a distance of approximately 100 meters of their pick-up location. Passengers' prices are set by the driver based on a combination of experience, provider requirements, and region-specific regulations. Payments to a provider (which is often some licensing authority) by drivers are set either in terms of a percentage commission for each journey or a fixed amount paid on a regular basis (e.g., monthly).

The limited choice of drivers for passengers to select is partly resolved by the dispatcher mechanism. In particular, passengers are able to call a provider who then forwards their request to a larger number of drivers within their region, often via a radio network. A driver is then selected based on a heuristic rule implemented by the provider. However, the heuristic does not usually include any information about the price the driver intends to offer the passenger as this information is not available to the provider. Pricing is therefore performed in a similar fashion to the Hackney carriage mechanism.

It is informative to formalize the pricing algorithm used in the Hackney carriage and dispatcher mechanisms in the case where the provider seeks to maximize its profit. Without historical data, the provider can set prices based only on *average* properties of the ODT system, including driver travel distances and durations, how drivers are selected, minimum profit targets, and maximum prices that passengers will pay.

An example of a model based on the average properties of the ODT system has been proposed in [10], [45]. In this model, the average waiting time is given by

$$W = \frac{\omega}{pN_T - \frac{DL}{\gamma}} \quad (6)$$

where ω depends on the density of taxi stands, p is the proportion of available taxis, N_T is the total number of taxis, γ is the average number of passengers per journey, L is the average travel time, and D is the average number of passengers served by the whole taxi system. As such, $pN_T - \frac{DL}{\gamma}$ represents the average number of available taxis. This kind of model forms the basis of optimization problems to maximize revenue and the number of served passengers [10], [45]. In particular, the average monetary and time cost of a trip in time period i is given by

$$D^i(F^i, L^i, W^i) = D^i \exp \left(\beta \left(\frac{F^i}{\gamma} + \psi_1 L^i + \psi_2 W^i \right) \right) \quad (7)$$

where F^i is the average fare price, L^i is the average travel time, W^i is the average waiting time given by (6), $\beta > 0$ is a sensitivity parameter, and ψ_1, ψ_2 are parameters used to convert time costs into monetary costs. Often this approach is studied within the framework of multiperiod pricing [10], [16] however, it is not able to capture the individual private preferences of passengers and drivers.

On the other hand, within our model detailed in Section II, the Hackney carriage and taxi mechanisms can be formalized

as follows. Consider passenger i and driver j , with prices set via the optimization problem

$$\max_{r_i, \eta_j} \mathbb{E}_{R_i, R_{j,i}, \tau_{j,i}, \kappa_j} \left[r_i R_i \eta_j \mathbb{P} \left(\frac{S_{j,\min} \tau_{j,i} + \kappa_j (R_i + R_{j,i})}{1 - \eta_j} \leq r R_i \leq r_{i,\max} R_i \right) \right] \quad (8)$$

Intuitively, this optimization problem maximizes the expected profit from each passenger request over all possible distances and duration $R_i, R_{j,i}, \tau_{j,i}, \kappa_j$. To see this, observe that the probability that both passenger i and driver j accept their prices is given by

$$\mathbb{P} (S_{j,\min} \tau_{j,i} \leq (1 - \eta_j) r_i R_i - \kappa_j (R_i + R_{j,i}), r_i \leq r_{i,\max}) \quad (9)$$

Since the provider receives a profit of $\eta_j r_i R_i$, it follows by rearranging that the expected profit is precisely the objective in (8), where we have accounted for the probability that both the passenger and driver accept the prices (the provider receives nothing if one or both reject their offer).

In general, the problem in (8) cannot be solved directly due to the intractability of the distributions involved and limited knowledge of the distributions of $S_{j,\min}$ and $r_{i,\max}$. As such, one approach to set prices is to decouple the passenger and driver pricing subproblems, which leads to the choices of

$$\begin{aligned} r_{H,i} &= \arg \max_r \mathbb{E}_{R_i} [r R_i \mathbb{P}(r \leq r_{i,\max})] \\ \eta_{H,j} &= \arg \max_{\eta} \mathbb{E}_{R_i, R_{j,i}, \tau_{j,i}, \kappa_j} \left[\eta \mathbb{P} \left(\frac{S_{j,\min} + \kappa_j (R_i + R_{j,i})}{1 - \eta} \leq r_i R_i \right) \right] \end{aligned} \quad (10)$$

As for (8), prices in (10) are averaged over distance and duration parameters. This means that the prices are not dynamically adjusted depending on the real-time state of the ODT system.

We remark that a multiperiod model based on (6) has also been proposed in [13], [14], [22], which also accounts for strategic behavior of drivers, and which considered two additional constraints, namely: the driver does not work for more than n_w periods; and that the driver does not work continuously for more than n_c periods. This leads to a non-cooperative game model, where the Nash equilibrium [25] is obtained via a bilevel optimization problem. Despite accounting for multi-period pricing and driver strategy constraints, it is important to note that the model still relies on the average behavior of passengers and drivers. This is in contrast with our model that accounts for the presence of private individual preferences of each passenger and driver.

With the introduction of smartphone apps, it is now possible to exploit large data sets to price and allocate passengers and drivers. One method is to use a posted-price mechanism approach for both passengers and drivers. In this two-sided posted-price mechanism, the provider sets a single price for all passengers and a single price (in, for example, euros/km) for all drivers in a given region at a given time. Moreover, the mechanism can also account for the locations of passengers and drivers as well as instantaneous supply (availability of drivers) and demand (number of passenger requests).

In contrast with the Hackney carriage and dispatcher mechanisms, the two-sided posted-price mechanism optimizes prices based on the instantaneous properties of the ODT system. In the case that the provider only considers its profits, a natural choice of prices is

$$\begin{aligned} r_{D,i} &= \arg \max_r r R_i \mathbb{P}(r \leq r_{\max}) \\ \eta_{D,j} &= \arg \max_{\eta} \eta \mathbb{P}\left(\frac{S_{j,\min} + \kappa_j(R_i + R_{j,i})}{1 - \eta} \leq r_i R_i\right) \end{aligned} \quad (11)$$

Note that the key difference between (11) and (10) is the absence of expectations over distances and durations. This is due to the fact that the instantaneous and historical data available to the provider allows for the prices to take into account real-time properties of the ODT system. The quality of the solutions in (11) will depend on how effectively the provider has learned the distributions of $r_{i,\max}$ and $S_{j,\min}$.

All of the mechanisms so far have relied on posted-prices, where passengers and drivers can either accept or reject offers. For passengers, it is common to use a posted-price mechanism as each passenger does not make a sufficient number of requests to make informed choices for the bids that are required in auction-based approaches. Moreover, there is a concern that adopting an auction-based approach for passengers would mean that wealthier passengers are prioritized over others. As this is an undesirable feature for ODT services designed for the public, it rules out applying auctions to passenger pricing (with the possible exception of business focused services [19]).

On the other hand, drivers are paid based on a minimum profit criterion, raising the question of whether it is desirable to use auctions to determine driver payments and allocation. There are three reasons that suggest auctions may play a useful role for driver pricing and allocation. First, the two-sided posted-price mechanism in (11) requires estimates of the pdf both for the maximum price-rate passengers are prepared to pay and the minimum profit drivers are willing to receive. If the provider has not been operating for long or is not using an efficient learning algorithm then the estimates may not be accurate. Here, it may be desirable for drivers to bid so that the provider can obtain accurate knowledge of their profit preferences. This is revealed in an auction, but not in posted-price mechanisms. Second, auctions often yield higher revenue than posted-price mechanisms when the pdf of the driver preferences has large weight in the tails [29]. Thus, when there is large variation in driver preferences, an auction may lead to higher revenue for the provider. Third, an auction to price and allocate drivers has been implemented by Liftago³. Their success suggests that this approach can be practical. It is thus important to understand the mechanism's properties.

Although provider profit is not the only mechanism feature of interest—average passenger prices and driver profits also matter—it suggests that hybrid mechanisms where passengers are priced via a posted-price mechanism and drivers via an auction may play a useful role in ODT systems. To this end, in the following sections, we introduce an online hybrid mechanism (Section IV) and study its properties from the perspective

of the allocation of a single passenger request (Section V), followed by a system level evaluation (Section VII).

IV. THE HYBRID MECHANISM

In this section, we develop the pricing and allocation mechanism that is run by the provider. We call our approach a hybrid mechanism since passengers are priced by a posted-price mechanism and driver payments are determined by a single-sided auction. As such, it is neither a two-sided posted-price mechanism nor a classical double auction [47]. Note that since a posted-price mechanism is used for the passengers, only one passenger at a time can be priced and allocated to one of a subset of all available drivers.

In practice, when there are multiple simultaneous requests, a subset of the available drivers must be selected for each passenger, a process known as *market formation*. As we focus on the pricing and allocation mechanism, we assume that only one passenger request arrives at a time and that all available drivers (without a passenger) can potentially be allocated.

In our hybrid mechanism, there are two phases for each passenger request. In the first, passengers are priced by an online posted-price mechanism. In the second phase, a driver is selected via an auction. We detail each phase corresponding to passenger $i \leq n$ (i.e., the i -th request) as follows.

A. Phase 1: Passenger Pricing

Phase 1 of the mechanism is run immediately after a passenger requests a journey. Let $K \in \mathbb{N}$ and define the price set $\mathcal{P} = \{\frac{\rho_{\max}}{K}, \frac{2\rho_{\max}}{K}, \dots, \rho_{\max}\}$. The price set consists of the possible price-rates that the provider can offer each passenger. Note that using a finite set of prices rates is consistent with current practice in Hackney carriage or dispatcher mechanisms.

Passenger $i \leq n$ is then priced according to the following online posted-price mechanism:

Algorithm 1 Passenger pricing.

Offer: The provider offers passenger i a price p_i computed as follows.

- If $1 < i \leq K$, then set $p_i = R_i \frac{j^* \rho_{\max}}{K}$.
- Otherwise, set the price to be

$$p_i = R_i \frac{j^* \rho_{\max}}{K} \quad (12)$$

where j^* maximizes $\bar{w}_j + \sqrt{\frac{2 \log i}{n_j}}$, $w_j = \frac{j \rho_{\max}}{K}$, n_j is the number of times a passenger has been offered the price w_j , \bar{n}_j is the number of times a passenger offered price w_j has accepted, and $\bar{w}_j = w_j \bar{n}_j / n_j$.

Response: The passenger either accepts or rejects the offer.

Observe that the mechanism in **Phase 1** is online, which means that prices are adapted as the number of transactions grows. In particular, **Phase 1** is based on the online posted-price mechanism proposed in [28], which exploits the UCB1 learning algorithm [48]. As such, the mechanism inherits desirable properties such as truthfulness and known bounds for the profit loss due to imperfect estimates of the pdf for the maximum price-rate of each passenger. We explore these properties further in Section V.

³For more details on Liftago's hybrid mechanism see [46].

B. Phase 2: Driver Allocation and Payment

Phase 2 of our hybrid mechanism allocates a driver to the passenger and determines the payment they receive. Algorithm 2 describes this process for a request from passenger i .

Algorithm 2 Driver allocation and payment.

Driver Identification: The provider identifies those drivers able to reach passenger i within Δ minutes.

Offer: Feasible drivers are informed how much passenger i paid for the service in **Phase 1** and the distance R_i , allowing them to compute their potential profit for transporting i .

Driver Bidding: Each driver makes a sealed bid for the minimum payment they are willing to receive. The bid for driver j corresponds to η_j and is only considered if $\eta_j \geq 0$.

Allocation: The provider selects driver j with the highest bid η_j , with ties broken at random. This driver receives profit $(1 - \eta^*)r_i R_i - c_{j,i}$, where η^* is the value of the second highest bid. The provider receives a profit $\eta^* r_i R_i$.

Observe that the mechanism in **Phase 2** has low communication requirements. In particular, the driver identification step reduces the number of drivers that are allowed to bid. Moreover, the driver bidding step only requires each driver to make a single bid per passenger request, which is equivalent (in terms of communication requirements) to the two-sided posted-price mechanism of Section III. **Phase 2**'s duration is short and the bandwidth required is low. Given the simple strategy in this phase, automating driver bids is also possible.

We also note that there is no guarantee that a driver will be found in **Phase 2** who is willing to transport each passenger and is able to arrive within Δ minutes of the request. This problem can be resolved by allowing the provider to retract the journey offer and refund any payment made by the passenger. However, this situation is undesirable and it is therefore important for the provider to give incentives to drivers to ensure that at least one driver is willing to transport each passenger. We explore variations of the mechanism to improve the proportion of passengers transported in Section VI.

C. Numerical Example

We provide an illustrative numerical example for our hybrid mechanism. Let p_i be the price obtained for passenger i from the **Offer** step in Algorithm 1. This price depends on the distance, R_i , of the journey and previous offers to other passengers. For concreteness, we assume that $p_i = 10$ euros. Next, suppose that the passenger is willing to pay 11 euros for their journey. By the **Response** step in Algorithm 1, the passenger accepts. The provider then seeks to find a driver.

Suppose that two drivers are able to reach the passenger within a time, Δ , which in Section VII we take to be 10 minutes. The provider informs each driver that the passenger is willing to pay 10 euros for the journey. Each driver bids η_j , $j = 1, 2$. For concreteness, we assume that $\eta_1 = 0.6$, $\eta_2 = 0.4$. Since $\eta_1 > \eta_2$, driver 1 is allocated the journey and receives a profit of $6 - c_{1,i}$ euros, where $c_{1,i}$ is the cost for driver 1 to reach the pickup location of the passenger. The provider then receives 4 euros profit.

V. PROPERTIES OF THE HYBRID MECHANISM

In this section, we explore the properties of the hybrid mechanism. In Section V-A, we show that each phase our mechanism ensures that passengers and drivers have incentives to report truthfully. We then consider the effect of historical data on the profit the mechanism in Section V-B, which leads to a data-profit tradeoff. In Section V-C we derive bounds to relate the profit to the passenger preferences and imperfections in the provider's estimates.

Throughout this section, we compare the profits obtained via our hybrid mechanism to a mechanism which has perfect knowledge of the passenger preference distribution, which we denote by \mathcal{M}^* . In the mechanism \mathcal{M}^* , passenger i is priced via the optimization problem

$$r_{\mathcal{M},i} = \arg \max_{r \in [0, \rho_{\max}]} r \int_r^{\rho_{\max}} f_{r_{\max}}(x) dx \quad (13)$$

To simplify notation, we define the *expected gross profit* based on a price-rate r as

$$J(r) = r \int_r^{\rho_{\max}} f_{r_{\max}}(x) dx \quad (14)$$

For passenger i , the expected gross profit from our hybrid mechanism is denoted by $J_{H,i}^*$, while the expected gross profit from mechanism \mathcal{M}^* is denoted by $J_{\mathcal{M},i}^* = J(r_{\mathcal{M},i})$.

A. Truthful Reporting

As **Phase 1** is based on an online posted-price mechanism and **Phase 2** is based on a sealed-bid second-price auction, our hybrid mechanism inherits the truthful reporting properties of these mechanisms under no collusion. This ensures that passengers and drivers do not have any incentive to use reporting strategies that are not straightforward to characterize by the provider. In particular, the provider can use simulations to obtain more reliable estimates of its long-term profits, the profits of drivers and the prices that passengers will pay. These are important to verify the suitability of the mechanism.

We first consider passenger reports in **Phase 1**, which consist of the desired time of pick-up and the passenger's accept or reject response. Observe that in **Phase 1**, the price for passenger i does not depend on the time passenger i makes her request nor on whether they accept or reject the journey. As such, each passenger cannot influence its price by reporting an earlier or later pick-up time. Moreover, if a passenger untruthfully reports her accept or reject response, she will either receive a unwanted journey or miss out on an opportunity to receive a desirable journey. This means that passengers will report truthfully.

In **Phase 2**, only one passenger request is active in each hybrid mechanism run. Thus, the truthfulness of driver bids is inherited as a property of sealed-bid second-price auctions, which are themselves Vickrey auctions [25]. It follows that the hybrid mechanism is truthful for both passengers and drivers.

B. Data-Profit Tradeoff

We now turn to the tradeoff between historical data available to the provider and the profit of the mechanism. In particular,

we consider the sum expected gross profit after n passenger requests. Our first proposition provides an asymptotic scaling law for the expected loss of profit by the provider when using the hybrid mechanism rather than another mechanism:

$$\mathbb{E}[L] = \mathbb{E} \left[\sum_{i=1}^n J_{\mathcal{M},i}^* - \sum_{i=1}^n J_{H,i}^* \right] \quad (15)$$

Proposition 1. *Let $K = \lceil (n/\log n)^{\frac{1}{4}} \rceil$ and suppose that the parameters of the passenger price preferences satisfy $\alpha = 1$ and $\beta > 1$. Then, $\mathbb{E}[L]$ satisfies*

$$\mathbb{E}[L] = O(\sqrt{n \log n}) \quad (16)$$

The proof of Proposition 1 relies on the following lemma.

Lemma 1. *Suppose $\alpha = 1$ and $\beta > 1$. Then the unique maxima of $J(r)$ in (14) is*

$$r^* = \frac{\rho_{\max}}{1 + \beta} \quad (17)$$

Proof. The condition for a critical point is

$$\frac{\partial}{\partial r} \left\{ r \int_r^{\rho_{\max}} f(x) dx \right\} = 0 \quad (18)$$

Applying the Leibniz rule and evaluating the integral yields

$$\frac{\rho_{\max}}{\beta} \left(1 - \frac{u}{\rho_{\max}} \right)^{\beta} - u \left(1 - \frac{u}{\rho_{\max}} \right)^{\beta-1} = 0 \quad (19)$$

which has the solution

$$u = \frac{\rho_{\max}}{\beta + 1}. \quad (20)$$

We now show that the point u is the unique maximum of $J(r)$. Without loss of generality, we focus on the case $\rho_{\max} = 1$ since this only affects the support. Observe that

$$\int_r^1 \frac{1}{B(1, \beta)} (1-x)^{\beta-1} dx - \frac{1}{B(1, \beta)} r (1-r)^{\beta-1} < 0$$

$$\Leftrightarrow \frac{1}{\beta} (1-r) - r < 0 \Leftrightarrow r > \frac{1}{1+\beta}$$

which implies that $J'(r) < 0$ for $r > \frac{1}{1+\beta}$. Using a similar argument, it follows that $J'(r) > 0$ for $r < \frac{1}{1+\beta}$. This implies that u is in fact the unique maximum for $J(r)$. \square

We now complete the proof of Proposition 1.

Proof. Since $J(r)$ has a unique maximum on $[0, 1]$ it follows that we can apply [28, Theorem 3.14], which states that $\mathbb{E}[L] = O(\sqrt{n \log n})$ when $K = \lceil (n/\log n)^{\frac{1}{4}} \rceil$, from which Proposition 1 then follows. We note that the statement of Theorem 3.14 in [28] corresponds to the case $\rho_{\max} = 1$; however, the result can be extended to $\rho_{\max} \in \mathbb{R}_+$. \square

We remark that Proposition 1 does not apply to every choice of α_r, β_r for the distribution of the maximum price-rate. However, it holds for a wide range of parameter choices where the tail of the density as $r \rightarrow \rho_{\max}$ is decaying. Moreover, a numerical study suggests that the result holds for all α, β . As passengers are typically less likely to pay for higher priced journeys, the choices of parameters in Proposition 1 are the most applicable to ODT systems.

A drawback of Proposition 1 is that it does not provide any information on the effect of passenger preferences on the expected loss. However, intuition suggests that the profit loss $\mathbb{E}[L]$ should be smaller for lower values of the variance for the maximum price-rate $r_{i, \max}$, which we verify next.

C. Influence of Preferences

While Proposition 1 provides insight into how the expected gross profit varies with the number of requests (corresponding to the available data), it does not indicate how passenger and driver preferences affect the performance of the mechanism.

Suppose passenger i is offered a price-rate r_i by the hybrid mechanism in **Phase 1**, which is typically suboptimal. This offer can be interpreted in terms of the mechanism \mathcal{M}^* , where the price-rate is optimized using a pdf $\hat{f}_{r_{\max}}$ with parameters $\hat{\alpha}, \hat{\beta}, \rho_{\max}$. The expected gross profit loss due to the imperfect estimate $\hat{f}_{r_{\max}}$ is then given by

$$\mathbb{E}[L_i] = J_{\mathcal{M},i}^* - r_i \int_{r_i}^{\rho_{\max}} \hat{f}_{r_{\max}}(x) dx \quad (21)$$

Due to the dependence on r_i arising from the mechanism in **Phase 1**, it is not possible to obtain a closed-form expression for $\mathbb{E}[L_i]$. As such, we focus on obtaining a bound to examine the effect of parameters for passenger preferences on the profit loss. We begin by considering the upper bound for the expected gross profit loss.

Proposition 2. *Suppose that $\alpha = 1$ and $\beta > 1$, then $\mathbb{E}[L_i]$ is given by (22), where ψ is the digamma function.*

Proof. Observe that

$$\begin{aligned} \mathbb{E}[L_i] &= r^* \int_{r^*}^{\rho_{\max}} f(x) dx - r_i \int_{r_i}^{\rho_{\max}} \hat{f}(x) dx \\ &\leq r^* \int_{r^*}^{\rho_{\max}} f(x) dx - r^* \int_{r^*}^{\rho_{\max}} \hat{f}(x) dx \\ &\leq r^* \int_0^{\rho_{\max}} |f(x) - \hat{f}(x)| dx \end{aligned} \quad (23)$$

Consider the change of variables $y = x/\rho_{\max}$, which yields

$$\mathbb{E}[L_i] \leq r^* \int_0^1 |\tilde{f}(y) - \tilde{\hat{f}}(y)| dy \quad (24)$$

where \tilde{f} and $\tilde{\hat{f}}$ are beta pdfs with support $[0, 1]$ and parameters α, β and $\hat{\alpha}, \hat{\beta}$ respectively.

Note that the integral in (24) is twice the total variation. Applying Pinsker's inequality [49] to bound the total variation in terms of the Kullback-Leibler divergence, D_{KL} , we obtain

$$\mathbb{E}[L_i] \leq \sqrt{2D_{KL}(\tilde{f}, \tilde{\hat{f}})} \quad (25)$$

which yields the desired result after evaluating the Kullback-Leibler divergence of beta distributed random variables. \square

In the special case that $\alpha = \hat{\alpha} = 1$ and $\beta, \hat{\beta} \approx 1$, we can obtain further insight into the effect of imperfect knowledge

$$\mathbb{E}[L_i] \leq \sqrt{2 \left[\log \left(\frac{B(\hat{\alpha}, \hat{\beta})}{B(\alpha, \beta)} \right) + (\alpha - \hat{\alpha})\psi(\alpha) + (\beta - \hat{\beta})\psi(\beta) + (\hat{\alpha} - \alpha + \hat{\beta} - \beta)\psi(\alpha + \beta) \right]} \quad (22)$$

of β . In particular, $B(1, \beta) = \frac{1}{\beta}$ and $\psi(1 + \beta) = \psi(\beta) + \frac{1}{\beta}$. As such, the bound in Proposition 2 simplifies to

$$\mathbb{E}[L_i] \leq \sqrt{2 \left[\log \left(\frac{\beta}{\hat{\beta}} \right) + \frac{\hat{\beta} - \beta}{\beta} \right]} \quad (26)$$

Since $\beta, \hat{\beta} \approx 1$, we obtain the approximation

$$\mathbb{E}[L_i] \lesssim \sqrt{2 \left[\beta - \hat{\beta} + \frac{\hat{\beta} - \beta}{\beta} \right]} \quad (27)$$

from which it follows that

$$\mathbb{E}[L_i] \lesssim \sqrt{K|\beta - \hat{\beta}|} \quad (28)$$

where K is a constant depending only on β and not $\hat{\beta}$. As such, the bound on $\mathbb{E}[L_i]$ approximately scales with the square root of the estimation error.

For providers, it is also of interest to understand how the mechanism performs in the presence of different passenger preferences using the mechanism \mathcal{M}^* . Suppose that the mechanism has perfect knowledge of passenger preferences. Let $r_{\alpha, \beta}^*$ be the optimal price-rate with parameters α, β . We then have the following result.

Proposition 3. *Consider two sets of passenger parameters α, β and α', β' such that $\alpha = \alpha' = 1$ and $\beta, \beta' > 1$. Then, the expected gross profit difference is given by*

$$\mathbb{E}[L_D] = \rho_{\max} \left| \frac{\beta^\beta}{(1 + \beta)^{\beta+1}} - \frac{\beta'^{\beta'}}{(1 + \beta')^{1+\beta'}} \right| \quad (29)$$

Proof. Since $\alpha = \alpha' = 1$ and $\beta, \beta' > 1$, we can apply Lemma 1. In particular, the expected gross profit difference is — as required — given by

$$\begin{aligned} \mathbb{E}[L_D] &= \left| r_{1, \beta}^* \int_{r_{1, \beta}^*}^{\rho_{\max}} f(x; 1, \beta) dx - r_{1, \beta'}^* \int_{r_{1, \beta'}^*}^{\rho_{\max}} f(x; 1, \beta') dx \right| \\ &= \left| \frac{\rho_{\max}}{\beta(1 + \beta)B(1, \beta)} \left(1 - \frac{1}{1 + \beta} \right)^\beta \right. \\ &\quad \left. - \frac{\rho_{\max}}{\beta'(1 + \beta')B(1, \beta')} \left(1 - \frac{1}{1 + \beta'} \right)^{\beta'} \right| \\ &= \rho_{\max} \left| \frac{\beta^\beta}{(1 + \beta)^{\beta+1}} - \frac{\beta'^{\beta'}}{(1 + \beta')^{1+\beta'}} \right| \end{aligned} \quad (30)$$

Observe from Proposition 3 that for large β, β' , the expected gross profit difference satisfies $\mathbb{E}[L_D] \approx 0$.

VI. VARIATIONS OF THE HYBRID MECHANISM

The hybrid mechanism aims to maximize the provider's profit. However, this is not necessarily desirable, e.g., if the provider seeks to serve more passengers. Similarly, the provider may wish to modify the mechanism if there are insufficient drivers willing to bid above the reservation price in the auction, such as when driver costs are too high.

A key motivation for providers to modify the hybrid mechanism is to improve the efficiency, which is a general notion in mechanism design. In the context of ODT it is natural to interpret the efficiency as the proportion of passengers that accept their offer and have a driver to transport them. That is, efficiency is the proportion of passengers able to successfully complete their desired journey and is denoted by \mathcal{E} .

There are two modifications the provider can make to the hybrid mechanism, namely to either adjust the price-rate, or to introduce driver subsidies. In the remainder of this section, we detail these variations on the mechanism developed in Section IV and explore their impact on efficiency.

A. Adjusting the Price-Rate

In the ideal scenario where the provider has perfect knowledge of passenger preferences and has elicited the minimum profit each driver requires, the profit maximization problem that the provider aims to solve for each passenger i is

$$P^* = \max_r \min_{\{S_{j, \min}\}} R_i r \left(1 - \frac{S_{j, \min} + \kappa_j(R_i + R_{j, i})}{r R_i} \right) \times \mathbb{P}(r \leq r_{i, \max}) \quad (31)$$

This problem can also be written as

$$P^* = \max_r r R_i \mathbb{P}(r \leq r_{i, \max}) - (S^* + C^*) \mathbb{P}(r \leq r_{i, \max}) \quad (32)$$

where S^* is the minimum profit of the selected driver and C^* is the corresponding travel cost. As such, the driver profit requirements regularize the optimization problem. In particular, the subtracted term results in a price-rate higher than the case where only passenger preferences are considered.

We note that the problem in (32) cannot be solved directly as it is necessary to elicit the driver profit requirements, which in turn depend on the price-rate that the passenger is charged. For this reason, a mechanism such as our hybrid mechanism (as described in Section IV) is required. However, the problem in (32) illustrates the important dependence of the price-rate on the efficiency (or equivalently the likelihood that both passengers and drivers are willing to accept the journey).

To gain further insight into the dependence of efficiency on the price-rate, consider the probability that at least one driver

is willing to transport passenger at price-rate r . That is, for a given passenger i there exists at least one driver j such that

$$S_{j,\min}\tau_{j,i} + c(R_i + R_{j,i}) \leq rR_i \quad (33)$$

The probability this holds is equivalent to the efficiency \mathcal{E} .

To explore the influence of journey parameters on the efficiency, consider driver j and define the random variable

$$X = S_{j,\min}\tau_{j,i} + \kappa_j(R_i + R_{j,i}) \quad (34)$$

Since X has bounded support determined by $S_{j,\min}, \kappa_j, R_i, R_{j,i}$, and $\tau_{j,i}$, we model X as a beta distributed random variable with parameters α_X, β_X . To justify this model, observe that for known $R_i, R_{j,i}, \tau_{j,i}, \kappa_j$

$$S_{j,\min} = \frac{X - \kappa_j(R_i + R_{j,i})}{\tau_{j,i}} \quad (35)$$

which is the beta distributed minimum driver profit (c.f., Sec. II).

With the definition of X , the optimization problem in (32) can be approximated as

$$P_{approx}^* = \max_r rR_i \mathbb{P}(r \leq r_{i,\max}) \Pr(X \leq rR_i) \quad (36)$$

which corresponds to the optimal profit when the provider allows only one driver to transport passenger i . This optimization problem provides useful insights into the tradeoff between the expected provider profit per journey and the efficiency of the mechanism. Observe that as the price-rate increases, both the approximate profit and the efficiency lower bound increase before decreasing again for large price-rate r . Intuitively, no drivers will accept a journey at a low price-rate while no passengers will accept a journey at a very high price-rate.

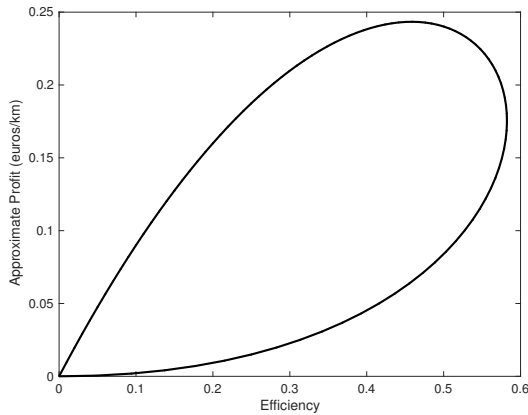


Fig. 1. Plot of the approximate profit (per km) in (36) versus the efficiency. Parameters: $\alpha_p = 1, \beta_p = 1, \alpha_d = 1, \beta_d = 5, \rho_{\max} = \sigma_{\max} = 1, R_i = 1$.

Fig. 1 illustrates the tradeoff between profit and efficiency. A key observation is that the optimal price-rate to maximize the approximate profit differs from the optimal price-rate to maximize the efficiency lower bound. Providers seeking to maximize efficiency will not in general choose the same price-rate as in the case of profit maximization. This is consistent with the differences between profit maximization and social planning objectives considered in auction design [29].

In practice, the provider does not have perfect knowledge of the efficiency-profit tradeoff curve. However, the information obtained from early passenger requests can provide clues as to this curve. In particular, as the online mechanism in **Phase 1** explores a range of prices, this data can also be used to construct an estimate of the efficiency-profit tradeoff curve.

B. Driver Subsidies

The efficiency of the hybrid mechanism is influenced by the requirement that each driver receive a minimum profit for each journey, which induces a loop in the profit-efficiency curve (as illustrated in Fig. 1). In particular, reducing the price-rate does not necessarily lead to improvements in the efficiency of the mechanism. This is due to the fact that reducing the price-rate can also reduce the probability that there exists a driver willing to transport each passenger.

An alternative method to improve the efficiency of the mechanism is to introduce subsidies. In particular, when no driver is able to serve a passenger profitably, the provider can compensate a driver that would otherwise not meet her minimum profit criterion. In this case, the provider pays the driver more than the price it charges the passenger.

While an unbounded subsidy can make every journey profitable for at least one driver, in general the subsidy must be bounded. This is to ensure that highly unprofitable passengers are not serviced so that the provider makes a long-term profit.

In the hybrid mechanism, subsidies are implemented by modifying the reservation price. Instead of requiring that $\eta > 0$ for each bid, the broker also accepts negative bids subject to the restriction that $\eta > \eta_{res}$ where $\eta_{res} < 0$ is the new reservation price. This increases the efficiency of the mechanism as drivers with bids

$$\eta = 1 - \frac{S_{j,\min} + \kappa_j(R_i + R_{j,i})}{rR_i} < 0 \quad (37)$$

can still bid for a journey.

To ensure that the provider does not suffer a significant loss, the reservation price $\eta_{res} \in \mathbb{R}$ is bounded. To evaluate the effect of the subsidy it is necessary to consider a system level evaluation which we describe next.

VII. SYSTEM LEVEL EVALUATION

Due to the complexity and spatial-temporal non-homogeneity of ODT systems, to properly understand the behavior of any mechanism it must be evaluated at the system level. In particular, the mechanism must be evaluated with realistic demand profiles of passenger (i.e., pick-up and drop-off locations and times), time-varying locations of available drivers, and city road networks. In this section, we carry out a simulation-based system level evaluation of our hybrid mechanism and the variations detailed in Section VI. In Section VII-A, we provide a detailed description of our simulation setup. At a high-level, we consider an ODT provider based in the Hague and build our simulator by modifying the Mobility Testbed [50], which provides realistic demand profiles for passengers. In Section VII-B, we detail the findings from our system level evaluation.

A. Simulation Setup

To evaluate our hybrid mechanism at the system level, we simulate the dynamics of the provider and passenger and driver agents. Our simulator is based on the multi-agent Mobility Testbed simulation tool for ODT systems. In contrast with other simulation tools⁴ for transportation systems built from scratch (e.g., [51]–[55]) and those based on AnyLogic [56], MATSIM or SUMO [57], the Mobility Testbed allows for tracking and modifications of driver routes facilitating the analysis of routing, scheduling and pricing algorithms.

The Mobility Testbed was originally designed to evaluate dial-a-ride allocation algorithms [50], [58]. As such, we make a number of modifications to incorporate taxi mechanisms and our hybrid mechanism, as well as more realistic driver behavior. We now describe the setup of our simulator.

We consider an ODT provider based in the Hague over a period of one day, with a total of 1000 passenger requests and 100 drivers (not all of whom are always available at a given time). The road network in the simulator is obtained from the OSM map of the Hague⁵, from which journey distances are computed. The region of the city considered is shown in Fig. 2. Journey durations and passenger requests (pick-up and drop-off locations) are based on the realistic traffic and demand models instantiated in the Mobility Testbed [58].

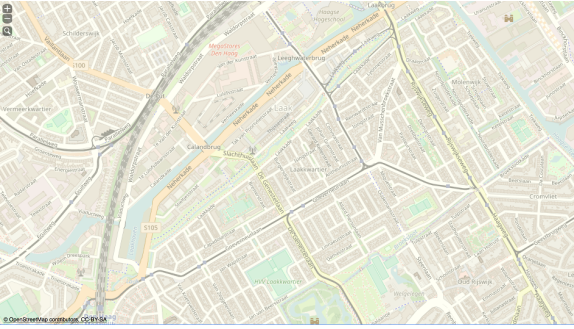


Fig. 2. OSM map of the Hague considered in the simulation study.

We set the maximum waiting time of each passenger, Δ , to 10 minutes, thereby imposing restrictions on the drivers that can feasibly transport any passenger. As such, the driver locations when they are not transporting a passenger affect the proportion of passengers that are transported (the efficiency) and also the revenue of the ODT provider.

To account for how drivers choose their location when they are not transporting a passenger, we adopt a hill-climbing approach. In particular, each node in the road network is assigned a probability that a passenger will make a pick-up request at the node. To determine where each available driver will travel in the next time period, each node k is assigned a score which is equal to the sum of the probabilities associated to each node that can be reached from node k within Δ minutes. Each available driver then travels to a neighboring node with the highest score and in the case of a tie, the

driver selects nodes with equal score uniformly. To obtain the score for each node, we used a training set 1000 passenger requests before the simulation began (in a real system, such pre-training could be performed using existing data available to the ODT provider). The scores were then updated as the simulation progressed.

To evaluate our hybrid mechanism, we compared it with the dispatcher and two-sided posted-price mechanisms. In the dispatcher mechanism, passengers are offered a price-rate of 2 euros/km and the selected driver is the nearest feasible one. This driver is paid a 10% commission. If no feasible driver is found, then the passenger is not transported. The price-rate was selected as representative of real-world rates. In the online two-sided posted-price mechanism, passengers and drivers are priced using **Phase 1** of our hybrid mechanism. For the drivers, the mechanism offers the nearest feasible driver the minimum acceptable profit. The full details for the online two-sided posted-price mechanism are provided in Algorithm 3, where similar to **Phase 1** the parameter $\bar{w}_{p,j}$ is the average passenger revenue from offering a price-rate $\rho_{\max}j/K$ and $n_{p,j}$ corresponds to the number of times that price-rate $\rho_{\max}j/K$ has been offered to passengers. For the drivers, we use a similar pricing scheme also based on UCB1, where $\bar{w}_{d,l}$ is the average amount paid to drivers from offering a payment of $\sigma_{\max}l/K$ and $n_{d,l}$ corresponds to the number of times that the payment $\sigma_{\max}l/K$ has been offered to drivers.

Algorithm 3 Two-sided online posted-price mechanism.

Initialization: Passenger i requests a journey.

if $1 \leq i \leq K$ **then**

Set $p_i = R_i i \rho_{\max} / K$

else

Set $p_i = R_i (\bar{w}_{p,j^*} + \sqrt{2 \log i / n_{p,j^*}})$, where j^* is chosen to maximize $\bar{w}_{p,j} + \sqrt{2 \log i / n_{p,j}}$.

Offer passenger i the journey at price p_i .

if $p_i > p_{i,\max}$ **then**

Passenger rejects offer, end algorithm.

else

Select closest feasible driver k .

if $1 \leq i \leq K$ **then**

$S_{k,i} = \tau_{k,i} i \sigma_{\max} / K$.

else

$S_{k,i} = -\tau_{k,i} (\bar{w}_{d,l^*} + \sqrt{2 \log i / n_{d,l^*}})$, where l^* is chosen to maximize $\bar{w}_{d,l} + \sqrt{2 \log i / n_{d,l}}$.

Offer driver k the journey with minimum profit $S_{k,i}$.

if $S_{k,i} < \tau_{k,i} S_{k,\min}$ **then**

Driver rejects offer, end algorithm.

else

Successful transaction: passenger pays p_i ; driver receives $S_{k,i}$; provider receives $p_i - \kappa_k (R_{k,i} + R_i) - S_{k,i}$.

Next, we detail our findings for the performance of our hybrid mechanism and its variations, as well as these two baseline mechanisms used for comparison. A full summary of simulation parameters is provided in Table I.

⁴For a detailed survey of simulation tools see [50].

⁵The Hague OSM map can be found in the Hague benchmark for the Mobility Testbed <https://github.com/agents4its/mobilitytestbed/wiki/Benchmarks>.

TABLE I
SIMULATION PARAMETERS.

Parameter	Parameter Value
Maximum Waiting Time Δ	10 minutes
$\alpha_d = \beta_d$	1
σ_{\max}	0.2 euros/km/minute
α_r	1
ρ_{\max}	10 euros/km
Number of Drivers	100
Number of Passengers	1000
Duration	24 hours
Vehicle Speed	15 km/hour
Vehicle Cost	1 euro/km

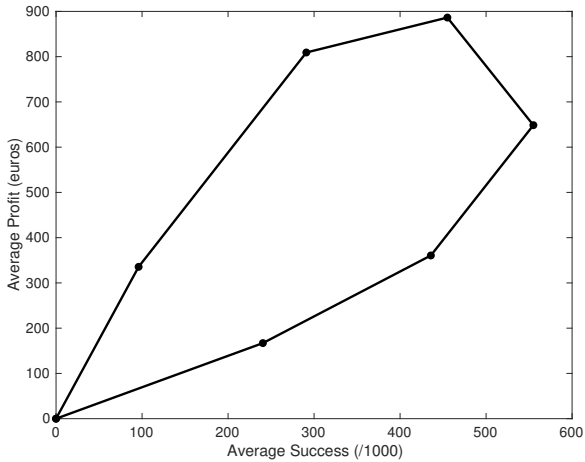


Fig. 3. Empirical efficiency-profit curve obtained by varying the offered price-rate in the dispatcher mechanism.

B. Simulation Results

Before evaluating the performance of our proposed mechanism, we first consider the behavior of the dispatcher model with varying passenger prices. Fig. 3 plots the tradeoff between the average number of successful journeys and the average profit received by the provider. The tradeoff is obtained by varying the price-rate offered to each passenger. Observe that the shape of the curve is consistent with the efficiency-profit curve in Fig. 1, obtained via theoretical analysis. As such under the dispatcher model, it is not possible to simultaneously maximize the profit and the number of successful journeys.

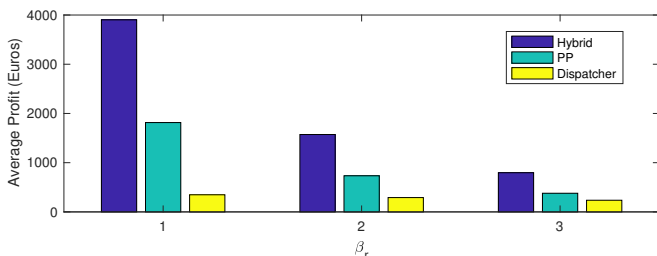


Fig. 4. Effect of passenger preferences on the average profit for the proposed, posted-price (PP) and dispatcher mechanisms.

While it is not possible to simultaneously maximize profit and the number of successful journeys, it is feasible to sig-

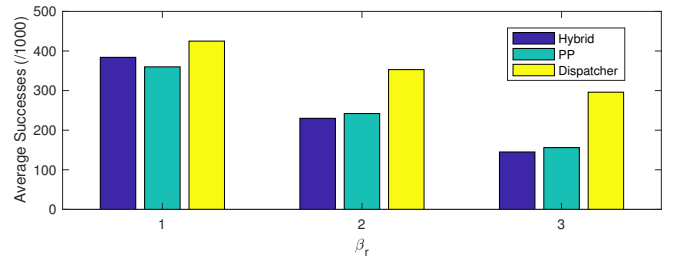


Fig. 5. Effect of passenger preferences on the success rate for the proposed, posted-price (PP) and dispatcher mechanisms.

nificantly improve the profit using our proposed mechanism. Figures 4 and 5 compare the average profit and success rate for each mechanism for a system with passenger and driver demand parameters $\alpha_r = \alpha_d = 1$, $\beta_d = 5$. The parameter β_r governs the shape of the beta distribution from which each passenger's maximum price-rate is drawn. A larger β_r means that passengers are less likely to have a high maximum price-rate and can be interpreted as a demand parameter. In particular, when there is low demand for a journey each passenger is prepared to pay less.

Observe that the proposed mechanism yields a significantly higher average profit compared with the posted-price and dispatcher models. Moreover, the profit and success values depend on the passenger demand. In particular, observe that as the passenger demand parameter β_r increases the profit reduces. This occurs as passengers are less likely to have a large maximum price-rate when β_r is large. As already noted, the average profit for the hybrid mechanism is significantly higher than for the posted-price mechanism. On the other hand, the success rate for the hybrid mechanism and the posted-price mechanism are comparable. The success rate for the dispatcher mechanism is higher than for the hybrid and PP mechanisms, at the cost of a significantly lower average profit.

We now study the impact of subsidies introduced in Section VI-B for a choice of $\beta_r = 1$. Table II shows the effect of introducing the subsidies. As expected, increasing the subsidy improves the success rate at the cost of the average profit. In particular, the profit loss corresponding to an improvement of approximately 80 passengers in the success rate is approximately 640 euros. However, a gain in 20 passengers corresponds to a profit loss of only 70 euros. This captures the intuition that there are diminishing gains in how many passengers can be reliably served.

TABLE II
EFFECT OF SUBSIDIES ON THE AVERAGE PROFIT AND SUCCESS RATE.

Subsidy Level (euros)	-1	-5	-10
Hybrid Ave. Profit (euros)	3766	3394	3069
Hybrid Ave. Success (/1000)	417	477	487

VIII. CONCLUSIONS

The wide availability of data means that it is possible for on-demand transport providers to improve pricing and allocation of passengers and drivers. In this paper, we proposed a hybrid

mechanism and studied its performance and sensitivity to data available to the provider from previous passenger requests. We show that the provider’s revenue loss scales with $\sqrt{n \log n}$, where n is the number of passenger requests the provider has received. We also compared the performance of the hybrid mechanism—for realistic models of passengers and drivers—with a two-sided posted-price mechanism and a mechanism based on the dispatcher taxi model. Our simulation results show that the success rate of our mechanism is comparable with an improvement in the average profit.

In future work, we intend to examine how interactions between multiple providers and passengers and drivers affect the mechanism, and what (if any) changes are required to function in such settings. We also intend to examine whether more sophisticated mechanisms for routing drivers without passengers can affect the latter’s profits.

ACKNOWLEDGEMENTS

This work was partly funded by the Technology Agency of the Czech Republic (grant no. TE01020155). The authors acknowledge the support of the OP VVV MEYS funded project CZ.02.1.01/0.0/0.0/16_019/0000765 “Research Center for Informatics”.

REFERENCES

- [1] J.-F. Cordeau, “The dial-a-ride problem: models and algorithms,” *Annals of Operations Research*, vol. 153, pp. 29–46, 2007.
- [2] D. Barbucha, “A multi-agent approach to the dynamic vehicle routing problem with time windows,” in *Proc. of the International Conference on Computational Collective Intelligence*, 2013.
- [3] G. Berbeglia, J.-F. Cordeau, and G. Laporte, “Dynamic pickup and delivery problems,” *European Journal of Operational Research*, vol. 202, no. 1, pp. 8–15, 2010.
- [4] C. Cubillos, F. Guido-Polanco, and C. Demartini, “Madarp: multi-agent architecture for passenger transportation systems,” in *Proc. IEEE International Conference on Intelligent Transport Systems*, 2005.
- [5] A. Glaschenko, A. Ivaschenko, G. Rzevski, and P. Skobolev, “Multi-agent real time scheduling system for taxi companies,” in *Proc. International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2009.
- [6] R. Bai, J. Li, J. Atkin, and G. Kendell, “A novel approach to independent taxi scheduling problem based on stable matching,” *Journal of the Operational Research Society*, vol. 65, 2014.
- [7] K. Seow, N. Dang, and D.-H. Lee, “A collaborative multiagent taxi-dispatch system,” *IEEE Transactions on Automation Science and Engineering*, vol. 7, no. 3, 2010.
- [8] Y. Dumas, J. Desrosiers, and F. Soumis, “The pickup and delivery problem with time windows,” *European Journal of Operational Research*, vol. 54, pp. 7–22, 1991.
- [9] M. Kümmel, F. Busch, and D. Wang, “Framework for automated taxi operation: the family model,” in *Transport Research Procedia*, 2017.
- [10] H. Yang, M. Ye, W. Tang, and S. Wong, “A multiperiod dynamic model of taxi services with endogenous service intensity,” *Operations Research*, vol. 53, no. 3, pp. 501–515, 2005.
- [11] H. Yang and T. Yang, “Equilibrium properties of taxi markets with search frictions,” *Transportation Research Part B: Methodological*, vol. 45, no. 4, pp. 696–713, 2011.
- [12] H. Yang, C. Fung, K. Wong, and S. Wong, “Nonlinear pricing of taxi services,” *Transportation Research Part A: Policy and Practice*, vol. 44, no. 5, pp. 337–348, 2010.
- [13] J. Gan, B. An, and C. Miao, “Optimizing efficiency in taxi systems: scaling-up and handling arbitrary constraints,” in *Proc. of the International Conference on Autonomous Agents and Multiagent Systems*, 2015.
- [14] J. Gan, B. An, H. Wang, X. Sun, and Z. Shi, “Optimal pricing for improving efficiency of taxi systems,” in *Proc. of the International Joint Conference on Artificial Intelligence*, 2013.
- [15] R. Balan, N. Khoa, and L. Jiang, “Real-time trip information service for a large taxi fleet,” in *Proc. International Conference on Mobile Systems, Applications, and Services*, 2011.
- [16] C. Borgs, O. Candogan, J. Chayes, I. Lobel, and Nazerzadeh, “Optimal multiperiod pricing with service guarantees,” *Management Science*, vol. 60, no. 7, pp. 1792–1811, 2014.
- [17] B. Cici, A. Markopoulou, and N. Laoutaris, “Designing an on-line ride-sharing system,” in *Proc. ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, 2015.
- [18] N. Ta, G. Li, T. Zhao, J. Feng, H. Ma, and Z. Gong, “An efficient ride-sharing framework for maximizing shared route,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 30, no. 2, pp. 219–233, 2018.
- [19] M. Egan, M. Schaefer, M. Jakob, and N. Oren, “A double auction mechanism for on-demand transport networks,” in *In Proc. PRIMA 2015: Principles and Practice of Multi-Agent Systems*, 2015.
- [20] J. Zhang, D. Wen, and S. Zeng, “A discounted trade reduction mechanism for dynamic ridesharing pricing,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 17, no. 6, pp. 1586–1595, 2015.
- [21] C. Zeng and N. Oren, “Dynamic taxi pricing,” in *In Proc. European Conference on Artificial Intelligence*, 2014.
- [22] J. Gan and B. An, “Game theoretic considerations for optimizing taxi system efficiency,” *IEEE Intelligent Systems*, vol. 32, no. 3, pp. 46–52, 2017.
- [23] M. Asghari, D. Deng, C. Shahabi, U. Demiyurek, and Y. Li, “Price-aware real-time ride-sharing at scale: an auction-based approach,” in *Proc. ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, 2016.
- [24] A. Kleiner, B. Nebel, and V. Ziparo, “A mechanism for dynamic ride sharing based on parallel auctions,” in *Proc. International Joint Conference on Artificial Intelligence (IJCAI)*, 2011.
- [25] Y. Shoham and K. Leyton-Brown, *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, 2008.
- [26] N. Nisan and A. Ronen, “Algorithmic mechanism design,” *Games and Economic Behavior*, vol. 35, no. 1-2, pp. 166–196, 2001.
- [27] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [28] R. Kleinberg and T. Leighton, “The value of knowing a demand curve: bounds on regret for online posted-price auctions,” in *IEEE Symposium on Foundations of Computer Science*, 2003.
- [29] R. Wang, “Auctions versus posted-price selling,” *The American Economic Review*, vol. 83, no. 4, pp. 838–851, 1993.
- [30] S. Chawla, J. Candogan, J. Chayes, I. Lobel, and H. Nazerzadeh, “Multi-parameter mechanism design and sequential posted pricing,” in *Proc. ACM Conference on Theory of Computing (STOC)*, 2010.
- [31] O. Besbes and A. Zeevi, “Dynamic pricing without knowing the demand function: risk bounds and near-optimal algorithms,” *Operations Research*, vol. 57, no. 6, pp. 1407–1420, 2009.
- [32] X. Wang, F. He, H. Yang, and H. Gao, “Pricing strategies for a taxi-hailing platform,” *Transportation Research Part E: Logistics and Transportation Review*, vol. 93, pp. 212–231, 2016.
- [33] L. Zha, Y. Yin, and H. Yang, “Economic analysis of ride-sourcing markets,” *Transportation Research Part C: Emerging Technologies*, vol. 71, pp. 249–266, 2016.
- [34] S. Djavadian and J. Chow, “An agent-based day-to-day adjustment process for modeling ‘Mobility as a Service’ with a two-sided flexible transport market,” *Transportation Research Part B: Methodological*, vol. 104, pp. 36–57, 2017.
- [35] M. Egan and M. Jakob, “A profit-aware negotiation mechanism for on-demand transport services,” in *In Proc. European Conference on Artificial Intelligence*, 2014.
- [36] —, “Market mechanism design for profitable on-demand transport services,” *Transportation Research Part B: Methodological*.
- [37] L. Zhou and H. Xu, “An efficient double auction mechanism for on-demand transport services in cloud-based mobile commerce,” in *IEEE International Conference on Mobile Cloud Computing, Services and Engineering*, 2017.
- [38] S. Guo, Y. Liu, K. Xu, and D. Chiu, “Understanding ride-on-demand service: demand and dynamic pricing,” in *Proc. IEEE International Conference on Pervasive Computing and Communications*, 2017.
- [39] —, “Understanding passenger reaction to dynamic prices in ride-on-demand service,” in *Proc. IEEE International Conference on Pervasive Computing and Communications Workshop*, 2017.
- [40] L. Moreira, J. Gama, M. Ferreira, J. Mendes-Moreira, and J. Damas, “Predicting taxi-passenger demand using streaming data,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 3, pp. 1393–1402, 2013.

- [41] C. Brooks and E. Durfee, "Congregating and market formation," in *Proc. International Joint Conference on Autonomous Agents and Multiagent Systems*, 2002.
- [42] B. Verplanken and W. Wood, "Interventions to break and create customer habits," *Journal of Public Policy and Marketing*, vol. 25, no. 1, pp. 90–103, 2006.
- [43] S. Parker, "The generalised beta as a model for the distribution of earnings," *Economics Letters*, vol. 62, no. 2, pp. 197–200, 1999.
- [44] K. Bury, *Statistical Distributions in Engineering*. Cambridge University Press, 1999.
- [45] H. Yang, S. Wong, and K. Wong, "Demand-supply equilibrium of taxi services in a network under competition and regulation," *Transportation Research Part B: Methodological*, vol. 36, pp. 799–819, 2002.
- [46] J. Mrkos, J. Drchal, M. Egan, and M. Jakob, "Liftago on-demand dataset and market formation algorithm based on machine learning," *arXiv:1608.02858*, 2016.
- [47] R. McAfee, "A dominant strategy double auction," *Journal of Economic Theory*, vol. 56, no. 2, pp. 434–450, 1992.
- [48] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Machine Learning*, vol. 47, no. 2-3, pp. 235–256, 2002.
- [49] T. Cover and J. Thomas, *Elements of Information Theory, Second Edition*. John Wiley and Sons, Inc., 2006.
- [50] M. Čertický, M. Jakob, R. Přebil, and Z. Moler, "Agent-based simulation testbed for on-demand transport services (demonstration)," in *Proc. International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2014.
- [51] L. Fu, "A simulation model for evaluating advanced dial-a-ride para-transit systems," *Transportation Research Part A: Policy and Practice*, vol. 36, no. 4, pp. 291–307, 2002.
- [52] M. Horn, "Multi-modal and demand-responsive passenger transport systems," *Transportation Research Part A: Policy and Practice*, vol. 36, no. 2, pp. 167–188, 2002.
- [53] M. Jakob, M. Pechoucek, M. Cap, P. Novák, and O. Vanek, "Mixed-reality testbeds for incremental development of HART applications," *IEEE Intelligent Systems*, vol. 27, no. 2, pp. 19–25, 2012.
- [54] J. Jlassi, J. Euchí, and H. Chabchoub, "Dial-a-ride and emergency transportation problems in ambulance services," *Computer Science and Engineering*, vol. 2, no. 3, pp. 17–23, 2012.
- [55] L. Quadrifoglio, M. Dessouky, and F. Ordóñez, "A simulation study of demand responsive transit system design," *Transportation Research Part A: Policy and Practice*, vol. 42, no. 4, pp. 718–737, 2008.
- [56] J. Banks *et al.*, *Discrete-event system simulation*. Pearson, 2005.
- [57] P. d'Orey, R. Fernandes, and M. Ferreira, "Empirical evaluation of a dynamic and distributed taxi-sharing system," in *Proc. of ITSC*, 2012.
- [58] M. Čertický, M. Jakob, R. Přebil, and Z. Moler, "Agent-based simulation testbed for on-demand mobility services," in *Proc. International Workshop on Agent-based Mobility, Traffic and Transportation Models, Methodologies and Applications (ABMTRANS)*.



Nir Oren is a Reader in Computing Science at the University of Aberdeen. His primary research area is in the field of multi-agent systems, and in particular reasoning, explanation and optimisation within such systems. To this end, Dr Oren has worked in areas such as formal argumentation theory, computational trust and policy based reasoning, as well as virtual organisation creation and on-demand transport.



Michal Jakob is an associate professor of computer science at Czech Technical University in Prague, Czechia. His primary research area is the application of artificial intelligence in transport and mobility, in particular related to mobility as a service and large-scale electric mobility. He was the principal investigator of several large European research projects in the area of sustainable urban mobility. Michal has published over 50 scientific papers and is the principal inventor on two international patents.



Malcolm Egan received the Ph.D. in Electrical Engineering in 2014 from the University of Sydney, Australia. He is currently an Assistant Professor in CITI, a joint laboratory between INRIA, INSA Lyon and Université de Lyon. His research interests include information theory, statistical signal processing and mechanism design with applications in wireless and molecular communications, and intelligent transportation.