Value of Demand Information in Autonomous Mobility-on-Demand Systems

Jian Wen[a], Neema Nassir[a], Jinhua Zhao[a,*]

[a] *Massachusetts Institute of Technology, Cambridge, MA 02139*

Abstract

Effective management of demand information is a critical factor in the successful operation of autonomous mobility-on-demand (AMoD) systems. This paper classifies, measures and evaluates the demand information for an AMoD system. First, the paper studies demand information at both individual and aggregate levels and measures two critical attributes: dynamism and granularity. We identify the trade-offs between both attributes during the data collection and information inference processes and discuss the compatibility of the AMoD dispatching algorithms with different types of information. Second, the paper assesses the value of demand information through agent-based simulation experiments with the actual road network and travel demand in a major European city, where we assume a single operator monopolizes the AMoD service in the case study area but competes with other transportation modes. The performance of the AMoD system is evaluated from the perspectives of travelers, AMoD operators, and transportation authority in terms of the overall system performance. The paper tests multiple scenarios, combining different information levels, information dynamism, and information granularity, as well as various fleet sizes. Results show that aggregate demand information leads to more served requests, shorter wait time and higher profit through effective rebalancing, especially when supply is high and demand information is spatially granular. Individual demand information from in-advance requests also improves the system performance, the degree of which depends on the spatial disparity of requests and their coupled service priority. By designing hailing policies accordingly, the operator is able to maximize the potential benefits. The paper concludes that the strategic trade-offs of demand information need to be made regarding the information level, information dynamism, and information granularity. It also offers a broader discussion on the benefits and costs of demand information for key stakeholders including the users, the operator, and the society.

Keywords: value of information, autonomous vehicles, mobility-on-demand systems, fleet management, agent-based simulation

1. Introduction

The online-enabled mobility-on-demand (MoD) services have been launched in hundreds of cities around the world in the past few years. Powered by the autonomous vehicles (AV), the

* Corresponding author
Email address: jinhua@mit.edu (Jinhua Zhao)

concept of shared autonomous mobility-on-demand (AMoD) systems is poised to be one of the most disruptive innovations in the transportation industry. Although AMoD applications remain niche products, research on better design and operation of such systems have led to system design evaluation [1, 2, 3, 4, 5, 6], advanced dispatching algorithms [7, 8, 9], demand prediction tools [10, 11] and dynamic pricing [12, 13].

Managing information and leveraging the value of it are key factors for successful AMoD applications. A well-designed system will allow the smooth exchange of information between travelers and the operator, enabling both parties to be better informed and make more coordinated use of the resources. Current MoD applications often rely on online-enabled platforms for communication, and the conceptual autonomous systems are likely to follow the same technology choice. As the amount of data is multiplying and the desire for real-time solutions continues to grow, the effective management of information becomes even critical and requires more investigation in such AMoD systems.

However, most of the existing research has focused on the supply information provided to the consumers. In contrast, this paper focuses on the demand information provided to the operators in order to improve the efficiency of the AMoD system. We classify, measure and evaluate the demand information in an AMoD system. The demand information includes all details the operator has about the future trips, including basic information (such as origins and destinations of on-demand requests), and aggregate or individual information about the future trips for better operational decisions with regard to fleet sizing, dynamic pricing, and real-time dispatching. Potentially, the demand information will help achieve better service quality and deliver higher profitability. The value of information in this paper represents the overall benefits that demand information will bring to the operators. The costs for the operators to collect and manage information are not discussed in this paper.

This paper aims to make two contributions. First, the paper studies demand information at both individual and aggregate levels and measures two key attributes: dynamism and granularity. We identify the important trade-offs between both attributes during data collection and information inference processes. The compatibility of dispatching algorithms with different types of information is also discussed. Second, the paper examines the relationship between information and its value through agent-based simulation experiments. The simulation platform is capable of incorporating critical system design decisions including fleet sizes, sharing policies, and hailing strategies such as in-advance and on-demand requests. It also includes an insertion heuristic which assigns both on-demand and in-advance requests to vehicles and an optimal rebalancing policy to reposition idle vehicles based on the demand information about the future.

The case study uses the actual road network and travel demand in a suburban neighborhood of a major European city, where we assume a single operator monopolizes the AMoD service in the study area but competes with other transportation modes. Based on the case study, we test a series of scenarios with different settings of information in terms of levels, dynamism and granularity. We also compare different fleet sizes to reflect the operational choice of supply, which also depends on information availability. The performance of the simulated AMoD system is then evaluated from the perspectives of travelers, AMoD operator, and transportation authority regarding the overall system performance.

The remainder of the paper is organized as follows. Section 2 reviews the existing works and identifies the research gaps. Section 3 describes dynamism, granularity among others attributes for both individual and aggregate demand information. It also proposes the value of information as a measurement of the performance gain. Section 4 presents the case study, simulates the selected scenarios, and examines the simulation results. Section 5 discusses the limitations and points out the future. It also offers a broader discussion on the benefits and costs of demand information for key stakeholders including the users, the operator and the society.

2. Literature Review

With the advancement of information technology over the past decades, the value of traveler information has attracted a significant amount of research. However, most of the existing research has been focused on the supply information provided to the consumers. This includes the studies on real-time traffic information, traffic congestion updates, incidents alerts, and their impacts on drivers path choices [14, 15], users compliance with information [16], perceived user benefits [17, 18] and on mode choice and accessibility [19].

In public transit systems, it has been hypothesized and evidenced that providing real-time service information to passengers reduces indeterminacy of the service and helps passengers optimize their travel decisions [20, 21, 22, 23, 24, 25]. Tang and Thakuriah [26] observe that the weekday ridership for a bus route in Chicago with Bus Tracker service is 126 passengers more than that without such information, controlling for all other factors including population, socioeconomic characteristics, and transit service attributes. The ridership increase in the presence of information is also confirmed in the New York City Transit network [27], thanks to the improved passenger experience as the result of their ability to optimize their path choice decisions. In a behavioral experiment, Brakewood et al. [28] conclude that users of real-time information can on average save about 2 minutes of waiting time compared to non-users and have lower levels of anxiety and frustration while waiting. They also have higher levels of satisfaction with time spent waiting.

In contrast, this paper focuses on the demand information provided to the operators in order to improve the system efficiency of the AMoD system. The demand information, either real-time or historical, is critical for the mobility-on-demand systems. On one side, on-demand services are desirable to travelers because they provide flexibility in both time and space. On the other side, the dispatching of vehicles, including both request-vehicle assignment and anticipatory rebalancing, requires the presence of real-time demand information. Moreover, the dispatching can be even more efficient if (some of) the demand requests are known in advance so that the operators can plan ahead. As such, the operators are motivated to understand future demand and appreciate the value of the demand information.

The request-vehicle assignment problem has been intensively studied in the literature as the vehicle routing problem (VRP) [29], and many algorithmic studies of VRP have been focused on the performance gap between dynamic (online) systems, in which requests are registered during operation, and static (offline) ones, in which requests are known in advance. Borodin and El-Yaniv [30] are among the first to present the competitive ratio criteria as a metric to evaluate the performance of online algorithms. The competitive ratio is defined as the worst-case ratio of the cost of the online algorithm to the cost of an optimal offline algorithm. An algorithm is c-competitive if the competitive ratio of the algorithm is at most c. Jaillet and Wagner [31] then demonstrate that, the best possible competitive ratio of an online algorithm is 2 for VRP with m vehicles, infinite capacity and no precedence constraints.

Pillac et al. [32] argue that the competitive ratio criteria have drawbacks when applied to real-world applications, since it requires the proof of the c-competitiveness. The value of information presented by Mitrović-Minić et al. [33] constitutes a more flexible and practical metric. This metric indicates the performance of an online algorithm based on empirical results and captures the impact of the dynamism on the solution yielded by the algorithm. The larger the value is, the bigger the performance gap of the algorithm will be when a static system becomes completely dynamic. Gendreau et al. [34] report a value of information between 2.5% and 4.1% using their tabu search algorithm for the vehicle routing problem without capacity constraints. Tagmouti et al. [35] report a value larger than 10% for a neighborhood search descent heuristic when arcs in the graph are capacitated.

In contrast to using the value of information as a metric to assess dynamic VRP algorithms, understanding the value of information may also potentially improve AMoD operations by making use of it in system design, fleet management and real-time dispatching. In the literature, the most studied attribute of information is the dynamism, which describes to what extent information is known a priori. Larsen et al. [36] define the degree of dynamism as the proportion of requests that are sent on demand. The simulation results show that increasing the degree of dynamism results in a linear increase in operational cost. The in-advance information has therefore positive value for the operators. Diana [37] studies the impact of fleet size and assignment interval (cycle time) on dynamic and partially dynamic systems under different degrees of dynamism. The paper concludes that, when fleet size is small, the performance is more susceptible to the lack of a priori information. The cycle time does not significantly affect the solution. Wong et al. [38] investigated the system performance of demand responsive transport using simulation and found that a system incurs a higher operational cost and accommodates fewer requests when the request arrivals are partially dynamic, as compared to static or fully dynamic scenarios.

Rebalancing also represents a critical part of effective AMoD dispatching. Aggregate information, taking the form of demand volumes (outward flows) at station or zone levels, is often used to support rebalancing decisions. Based on the fluid model, Pavone et al. [39] propose an optimal rebalancing problem and simulate it on a 12-station AMoD system. In this system, every station will reach equilibrium so that there are excess vehicles and no waiting customers. Zhang and Pavone

[40] then extend the idea of the fluid model and present a queueing-theoretical approach within the framework of Jackson networks. Assuming that the demand volumes during the entire period of study are known at all stations, the approach gives a solution to an offline optimal rebalancing problem. They continue that, if taking only current demand information at a specific time point, the problem could be adopted to online applications.

Marczuk et al. [7] test both offline and online systems using Singapore travel data. The results show that about 28% and 23% fewer vehicles are required to guarantee the same service rate when offline and online rebalancing are in use respectively. Moreover, online policy outperforms the offline one by reducing the average wait time from 11 minutes to 9. Spieser et al. [41] tackle the rebalancing issues from the perspective of the operators. They evaluate the operational cost as a function of fleet size, service rate and vehicle utilization and demonstrate that rebalancing can reduce the cost significantly. However, both of the works are limited to systems with fixed demand pattern. The different attributes in the demand information and their impacts on cost have been omitted in their research.

The value of information as a performance metric of VRP algorithms has been intensively examined in the literature, yet the scale and scope of the studies are still limited to algorithmic progress and theoretical analysis. Recently, agent-based simulation has become popular in AMoD research for its advantages in capturing individual behaviors, enabling dynamic operations and accounting for stochasticity. However, none of the existing agent-based simulation applications have addressed the operational needs of determining the value that different kinds of demand information bring. This paper is an attempt to bridge the research gap between the two sides. Specifically, the paper extends the degree of dynamism with the motivation of developing a set of attributes to measure demand information. We are then able to study the value of demand information with different attributes using `amod-abm`, an agent-based modeling platform for AMoD simulation [42]. The simulation experiments are based on a case study area in a major European city under realistic operational settings.

3. Measuring Demand Information

3.1. Individual and Aggregate Information

The demand information consists of the origins, destinations, time constraints, and other necessary details of the upcoming trips. Typically, the time constraints are the latest pick-up time or the latest drop-off time. If a request is sent in advance for a future trip, the earliest pick-up time other than the request sending time is also included. The supply side is subject to constraints such as the vehicle capacity keeps within bounds the total number of passengers on board at a time; the maximum wait time controls the wait between request sending time (or earliest pick-up time if in advance) and actual pick-up time; and the maximum detour factor sets a limit on the actual detour of a shared ride, defined as the ratio of actual in-vehicle travel time to the shortest in-vehicle travel time.

Based on the assignment strategies, availability of vehicles as well as the relevant constraints, the operator responds to each request by either accepting or rejecting it. The accepted request is then assigned to a selected vehicle in the efforts to minimize the operational cost or maximize the level of service. A new route is then generated and the traveler gets notified. The rejection of a request often results from the constraint dissatisfaction, while a profit-driven operator may also refuse service due to other reasons, such as poor profitability associated with the trip and low user rating of the traveler.

Requests fed into the request-vehicle assignment are demand information at the individual level. If they are averaged/summed over a group of relevant entries, the information is considered aggregate. Aggregate demand information may also come from a combination of data sources, as well as the direct outputs of inference models. Although this type of demand information cannot provide the level of details which individual information does, it has the advantage of: (1) reducing the amount of presented information for fast queries and statistical analysis; (2) requiring less efforts in data collection and information inference, and (3) eliminating personally identifiable information so as to be free of privacy implications.

In practice, the demand information is often aggregated spatially. Earlier literature assumes simplified station-based AMoD systems and groups trips by origin stations. In the latest works, free-floating systems have prevailed in order to provide door-to-door service, in which aggregate demand information takes the form of zonal demand volumes. This type of information, even in a static format, plays a critical role in system-level decisions such as determining the best fleet size, designing the operational modes, and configuring the fare policies. When aggregate demand information becomes available in real time, it would also improve the performance of stochastic request-vehicle assignment algorithms, as future demand could be taken into consideration with probability. Similarly, the optimal rebalancing algorithm relies on the predicted zonal demand volumes in the near future to estimate the potential benefit of repositioning vehicles to each traffic analysis zone.

3.2. Sources of Information

Demand information usually comes from two sources: direct information exchange with the travelers and indirect information inference based on historical data [43, 44, 45, 46, 47].

The information exchange involves the mutual communication between the travelers and the operators. In most of the current MoD applications, the operators mainly serve on-demand requests. The on-demand services are attractive to both travelers and operators because of the flexibility: travelers can start their trips with little spatial and temporal limitations; operators reserve the right of rejecting requests when the system is saturated, thus avoid the uneconomical efforts to maintain an oversized fleet for all-time service availability. However, such systems have limitations due to the limited knowledge of future trips and the high computational requirements for real-time solutions.

Obtaining individual demand information a priori can be beneficial to the operators in many aspects. First, extra information of trips enlarges the search space by providing more combinations and options among which the optimal solution may exist. This makes better request-vehicle assignments possible. Extra information is especially helpful when most travelers are willing to share, as such a traveler is more likely to be paired with co-riders with close origin and destination, reducing both travel time and vehicle distance. Second, predictive demand information contributes to many marketplace decisions, such as fleet sizing, dynamic pricing and vehicle rebalancing. For these reasons, the AMoD operators want to look forward into the future horizon, ranging from dozens of minutes to a few hours, and estimate the demand ahead of time by collecting data and making predictions.

One way of doing so is by incentivizing travelers to send requests in advance. From the perspective of product design, by sending his/her request in advance, the traveler naturally expects being assigned a vehicle at the time of sending; and once assigned, s/he supposes that the service is secured and no further changes should be anticipated. It implies that by design, travelers requesting in advance can avoid competing with other on-demand requests traveling at the same time and have a higher probability of getting served. Such mechanism reduces the uncertainty for travelers in trip planning, including the potentially long wait, late arrival and even denied service. Therefore, the policy of in-advance requests is inevitably coupled with service priority, which will incentivize some travelers despite the possible loss of flexibility in trip-making. For operators, allowing in-advance requests is therefore double-edged. On the one hand, the operators benefit from the a priori information collected from the requests to improve dispatching efficiency; on the other hand, this information comes at the cost of promising service priority, even though some of the prioritized trips may put strong constraints on the operation and turn out not to be profitable.

The operators can also infer information about future trips using the historical data, and this approach becomes more attractive and feasible when the operators accumulate more historical data of their customers. Technically, the inference of individual demand information involves predicting the spatio-temporal behavior of frequent travelers based on Bayesian models or Markov chain-based models. For example Zhao et al. [47] predicts daily individual mobility represented as a chain of trips based on the Bayesian n-gram model. In terms of individual-level prediction, stochastic assignment algorithms are often preferred to their deterministic counterparts in order to deal with the uncertainty. For aggregate demand information, the mathematical models are abundant in the literature [48, 49, 50]. These models are usually descriptive, generalized and perform well at the system level. The demand is usually generated spatially and aggregated by socioeconomic characteristics to represent different market segments.

3.3. Attributes of Information

3.3.1. Dynamism and Granularity

Demand information, individual or aggregate, can be described and measured from different perspectives. In this paper, we refer to the different features of demand information as “attributes”. For example, the data collection process requires that source information should be reliable and valid, inference models require accuracy and consistency, and real-time systems talk about timeliness and integrity. In fact, the comprehensive set of attributes for demand information is vast. This paper will focus on two attributes: dynamism and granularity.

Dynamism measures the average reaction time the demand information permits during a given period of operation. In this paper, the reaction time is defined as the duration between the time point a piece of information is available (i.e. when a request is sent to the operator), and the time point it expires (i.e. when the traveler is supposed to be picked up). The longer the average reaction time is, the less dynamic the AMoD system will be.

Granularity in both spatial and temporal dimensions indicates the level of details presented in a set of demand information. For example, the tolerance of a prediction location or the size of a traffic analysis zone implies the spatial granularity. If the period of study is divided into a series of time segments, the time segment defines the temporal granularity of information. This paper only examines the spatial dimension of the information granularity.

As discussed in Section 3.2, the primary source of individual demand information is the requests from travelers, both on-demand and in-advance. In contrast, the aggregate demand information, if it's about the future, often comes from inference models. The measurement of dynamism and granularity depends on how information is processed and presented. Consequently, it's important to separate individual information from aggregate information.

Table 1 presents the measures for dynamism and granularity at both the individual and aggregate levels. This section will discuss each of the cells with special attention to the degree of dynamism for individual information (Cell 1) and the spatial granularity for aggregate information (Cell 4).

Table 1: Measuring the Attributes of Demand Information

	Dynamism	Granularity
Individual	Cell 1 (degree of dynamism)	Cell 2 (-)
Aggregate	Cell 3 (look-ahead time)	Cell 4 (spatial: zone size)

Cell 1: For individual requests, the reaction time is defined as the period between the time a request is known to the operator, and the time it is supposed to be picked up [36]. Assuming that an AMoD system serves a mix of on-demand and in-advance requests, we propose the "degree of dynamism" as a normalized measurement of the dynamism as follows.

The information that a request r contains can be described as a combination of origin o_r, destination d_r, earliest pick-up time $t_{e,r}$ and latest pick-up time $t_{l,r}$. $t_{s,r}$ is the sending time of request r. Given a set of requests $\mathcal{R}$, the degree of dynamism is defined as the average reaction time divided by the total time of study T_s.

$$\text{degree of dynamism} = \frac{1}{R} \sum_{r \in \mathcal{R}} \left(1 - \frac{t_{e,r} - t_{s,r}}{T_s}\right) \tag{1}$$

$\mathcal{R}$ includes on-demand requests $\mathcal{R}_{ond}$ and in-advance requests $\mathcal{R}_{adv}$. The sizes of the sets are R, R_{ond} and R_{adv} respectively. $\mathcal{R} = \mathcal{R}_{ond} \cup \mathcal{R}_{adv}$. $R = R_{ond} + R_{adv}$. The difference between $t_{e,r}$ and $t_{s,r}$ is therefore the reaction time for system to make dispatching decisions. For any on-demand request $r_{ond} \in \mathcal{R}_{ond}$, the system should respond in real time and we have $t_{e,r_{ond}}$ very close to $t_{s,r_{ond}}$. In an ideal dynamic system, the reaction time is zero. The degree of dynamism is 100% when the system is fully dynamic and all requests are sent on demand. The degree of dynamism decreases as the proportion of in-advance requests increases and as the reaction time gets longer.

Cell 2: Trip requests often contain specific locations and times which makes the discussion on the granularity trivial. However, some MoD applications do propose selected spots for picking up and dropping off, as an alternative to the door-to-door service. In this case, the high spatial granularity is unnecessary and travelers may be invited to walk for a short distance. The stop-to-stop option facilitates the operation, as multiple riders are more likely to be grouped together where rides are shared. By deliberately avoiding locations that require long detour and choosing selected spots instead, the efficiency of vehicle routing can be improved. But the granular spatial information is indeed useful if the system also tries to guide the travelers from their doors to the specified pick-up locations.

Cell 3: Aggregate information often “looks ahead” for some time to understand the demand in the future. This period is called the “look-ahead horizon”. For problems like fleet sizing, dynamic pricing, the look-ahead period can be a couple of hours or an entire day. To rebalance the idle vehicles dynamically, the look-ahead period is relatively short, from several minutes to an hour or so.

Cell 4: When conducting demand analysis, the operators often divide the service area into traffic analysis zones and choose the appropriate size of zones for the desired level of details. The zone size here represents the spatial granularity of the aggregate demand information, and fine division results in high granularity in space. In practice, each zone typically covers roughly the same area or population. High granularity of aggregate demand information often leads to good performance in rebalancing, however additional zones also add to the computational expenses.

3.3.2. Trade-offs between Attributes

Considering the financial, technological and regulatory limitations, the demand information to which the AMoD operator has access will never be perfect. An operator should make decisions on how to best spend efforts on data collection and management and find the balance between information dynamism and information granularity–the two important attributes that determine the quality of demand information.

In this section, we examine the information accuracy and discuss its two different aspects: request accuracy and prediction accuracy. A traveler may send wrong information, not follow original trip plans, and even cancel trips after being assigned services. The request accuracy describes the likelihood that a traveler actually follows up with his/her trip request. When an operator tries to predict the future demand based on the historical data, the prediction inaccuracy represents the systematic uncertainties in this process driven by the limitation of the prediction models and insufficiency of the input data.

An important trade-off during the information inference process is the balance between accuracy and granularity. When uncertainty in inference is significant, the point estimate that gives a single estimated value is unlikely to be true. Instead, the operator may use the confidence interval to specify a value range and its probability. The confidence interval is tightly coupled with the confidence level: the narrower the range is, the less probable that the true value will fall within the estimated range. Granularity and accuracy are therefore negatively correlated.

A variety of dispatching algorithms have been devised to handle the accuracy-granularity trade-off. If we assume accurate individual demand, deterministic request-vehicle assignment algorithms in the literature can apply almost without modifications. However, if individual demand information is inaccurate and the probability distribution gets involved, stochastic assignment algorithms that require much more computational resources should be used instead. As for rebalancing, the granularity of aggregate demand information is critical. The finer the zoning is, the better the spatial disparity of demand can be represented, and the more beneficial the rebalancing will be. The larger sample size will normally lead to a better estimate. However, in practice, the cost of collecting and processing a large amount of data should be taken into account.

The privacy issues are also closely tied to the accuracy-granularity trade-off. The travel information often contains sensitive personal information that may be used to locate a single person's home and workplace, to follow his/her activities, even to distinguish the individual identity. In this sense, travelers may find it unwanted to share highly granular data.

Another trade-off can be seen between dynamism and reliability. On-demand requests directly communicated to the operators often correctly describe the actual trips the travelers are willing to make and are reliable. For this reason, when making dynamic dispatching and routing decisions, current MoD applications often ignore the errors and unexpected changes in requests caused by travelers. However, the travelers' trip plans are more subject to changes in a longer time span and this unreliability represents one of the major challenges of implementing in-advance requests. If monetary leverage is allowed, the operators may impose a cancellation fee for those who have requested but decide to change. This paper does not focus on the monetary leverage for information, but points out the trade-off between the potential value that in-advance requests bring and the cost that the unreliable information generates.

3.4. Value of Information

Inspired by Mitrović-Minić et al. [33], the value of information in this paper is defined as:

$$voi = \frac{F - F'}{F'} \tag{2}$$

The metric measures the percentage gain in the objective. F is the value of the objective function when additional information is available; F' is the objective of the base scenario. Generally, the value of information of an AMoD system is positive since information enlarges the search space and has the potential of improving the performance. However, if the information comes with coupled service requirements that introduce new constraints to the service design, it is still possible that in some scenarios information may lead to a negative value of information. Example of this will be demonstrated in the case study in Section 4.

The objective of AMoD operations should reflect the interests of the stakeholders which may correlate or contradict with each other. For example, profit-driven operators often seek for maximum profit while travelers care about the level of service and travel cost. The governmental organizations, in contrast, have more considerations regarding system-level performance and overall social welfare. In this paper, we focused on three key indicators: ridership, waiting time and profit. We use the average number of served requests as the indicator of ridership. For travelers, we use the adjusted wait time as the main measure of service quality, which includes the average wait time for those served and a penalty for those rejected. The shorter the adjusted wait time is, and the more attractive the AMoD service is. For the operator, we use the profit as an indicator of its overall performance. In our simulation, the profit made during the period of study is derived from its fare revenue and operational cost as in Wen et al. [51]. The fare scheme takes the form of the base fare plus per-unit-distance and per-unit-time. The operational cost includes both fixed cost and variable cost. The former is based on fleet size, and the latter is proportional to the total vehicle distance traveled.

4. Simulation Experiments

4.1. Case Study Area

To demonstrate the framework in a real urban setting, we select a case study area (CSA) in a major European city. This city has an extensive and developed transportation network in which public transportation has a high mode share (45% in 2015). Commuter rail travel has also shown strong growth over the past decade, providing good service from the outskirt of the city to the downtown area.

CSA is based on a spread-out residential area located about 25 kilometers outside the downtown. It is centered around a commuter rail station with frequent and high-speed train service to downtown. However, bus service in this area is infrequent and not economically efficient as a result of the low residential density. Consequently, local trips are particularly car-dependent. The area is chosen as CSA because (a) it possesses a significant first-mile demand to the train station, (b) the inefficient bus service requires improvement, and (c) it has an appropriate density for initial AV trials. We choose the area of 15km$\times$10km and about 159 thousand people, which is a bit larger than the local administrative boundary to include all important bus routes originating from the rail station.

We use the official annual household travel diary surveys from 2005 to 2014 to analyze the current travel demand. This proportion contains 2709 respondents (1.7% of the population) and accounts for around 74,000 trips made by all residents after the sample is expanded to match the population. The all-mode demand is around 28,000 trips/h. According to an established mode choice model, the demand for the AMoD service is around 600 to 700 trips per hour, depending on the level of service and the fare structure.

4.2. Simulation Platform and System Settings

The value of demand information is studied through agent-based simulation experiments. The agent-based modeling platform used in this paper, `amod-abm`, has been specifically designed for simulating large-scale shared AMoD applications. Notably, it models individual travelers and vehicles as agents and analyzes system performance through various metrics of indicators. It also addresses the special dispatching need that is critical for effective AMoD operation. This includes insertion heuristic for assigning vehicles and pairing rides on the fly, and optimal rebalancing problem to offset the imbalance between vehicle supply and travel demand. The objective of the implemented dispatching algorithms is to maximize the number of served requests, and for those served, to minimize their total travel time including wait time and in-vehicle travel time.

Insertion heuristic finds an approximate solution to the assignment problem by considering each new request individually and independently from others and serving them on a first-come-first-serve basis. For each request, it searches for the best available vehicle with the objective of minimizing the incremental travel time that the inserted request imposes on all the travelers including itself. If none of the vehicles satisfies the constraints associated with the incoming request, the system rejects it. According to this algorithm, the AMoD operator prioritizes high ridership despite the profitability, and the travel time remains minimal whenever a new request comes in, although the global optimality that minimizes the total travel time is traded in exchange of a series of local optima for computational speed.

Rebalancing helps further reduce the travel time, or more specifically, the wait time. In this paper, this is done by applying an optimal rebalancing problem (ORP) to the system [39, 52]. ORP formulates a Mixed Integer Nonlinear Programming (MINLP) problem, the decision variables of which are the numbers of rebalancing vehicles sent from each zone to another at the beginning of a rebalancing period. The problem takes into account both the supply availability and the aggregate demand information, and the objective is to minimize the average wait time for all travelers, which is then translated to maximizing the total expected number of requests that can be served by vehicles from the same zone at the end of the rebalancing period. The simulation platform in this paper uses a combination of incremental-optimal and branch-and-bound methods to give a close approximation to the optimum [52].

The simulated AMoD system has a fixed-size fleet of taxi-like vehicles to provide service, and it serves a mix of on-demand and in-advance requests. Each vehicle may have up to 4 travelers on board at a time. Every trip has a chance to be shared. The maximum detour factor for a shared ride is 1.5. The maximum wait time is assumed to be 10 minutes. If a traveler is estimated to wait longer before any vehicle becomes available, the system refuses service and the notified traveler will walk away. The period of study (T_s) is 3600 seconds with 1800 seconds before as a warm-up period and 1800 seconds after as a cool-down period. The request-vehicle assignment algorithm - insertion heuristic - runs every 30 seconds. The optimal rebalancing executes every 150 seconds.

We distinguish three operational cases: (1) the low-supply case, in which fleet size is set to 110 and we observe a significant demand over supply; (2) the mid-supply case, in which the fleet size is 130 and demand and supply are comparable, and (3) the high-supply case, in which the fleet size is 150 and supply is over demand. In the low-supply case, the service rate is below 95% and travelers are likely to be rejected service due to the unavailability of vehicles. In the mid-supply case, the service rate is increased to 98%. In the high-supply case, the service rate reaches above 99% and travelers are almost sure to be served. As the fleet size becomes larger, not only will more travelers get served during the period of study, but also they will have a shorter adjusted wait time on average. However, from the operator's perspective, the vehicle idleness will become more significant as well, degrading the profitability of the operation.

4.3. Selected Scenarios

In this paper, we pay special attention to two attributes defined in Section 3.3.1 that determine the demand information: spatial granularity for aggregate information, and degree of dynamism for individual information. For illustration purposes, we select only a subset of scenarios as shown in Table 2. This table is designed to demonstrate the impact of the two attributes on the value of information using current system settings and deterministic dispatching algorithms. Table 3 shows how the comparison among the scenarios are made.

4.3.1. Spatial Granularity for Aggregate Information

We first compare scenarios with different levels of spatial granularity with regard to the aggregate demand information. The base scenario (Scenario 1) treats the entire CSA as a whole and the operator is only informed of the estimated total demand. This scenario has the coarsest granularity. In Scenario 2a and Scenario 2b, CSA is divided into selected numbers of gridded zones ($5 \times 5 = 25$ zones and $10 \times 10 = 100$ zones for 2a and 2b), the zone size of which are approximately 2km$\times$2km and 1km$\times$1km respectively. For each specific zone, the predicted demand volume is made available to the operator.

In practice, the aggregate demand information is usually available for most of the current MoD operators. However, current systems barely rebalance idle vehicles due to the difficulty in negotiating with human drivers. The conceptual autonomous systems, in contrast, comply perfectly with the dispatching decisions and make the execution of optimal rebalancing algorithms almost effortless. We also assume that the predicted demand volumes are known 150 seconds ahead of time (look-ahead period), exact (no errors) and constant throughout the simulation period (zero temporal granularity). The AMoD operator is then able to rebalance the idle vehicles dynamically following the optimal rebalancing strategy presented above.

4.3.2. Degree of Dynamism for Individual Information

In this part of the simulation, we assume that the individual demand information comes directly from the requests of travelers and is perfectly reliable. A portion of the travelers are allowed to send

Table 2: Selected Simulation Scenarios

Simulation Scenario	Level of Demand Information			
	Aggregate		Individual	
	Spatial Granularity	Other Attributes	Dynamism	Other Attributes
Scenario 1	total demand (coarsest)	as in 4.3.1	on-demand requests only	as in 4.3.2
Scenario 2a	zonal demand, 5×5 griding	as in 4.3.1	on-demand requests only	as in 4.3.2
Scenario 2b	zonal demand, 10×10 griding	as in 4.3.1	on-demand requests only	as in 4.3.2
Scenario 3a	zonal demand, 10×10 griding	as in 4.3.1	5% in-advance requests from CSA	as in 4.3.2
Scenario 3b	zonal demand, 10×10 griding	as in 4.3.1	10% in-advance requests from CSA	as in 4.3.2
Scenario 4	zonal demand, 10×10 griding	as in 4.3.1	5% in-advance requests, all in UST	as in 4.3.2

* The number of trips from UST accounts for approximately 10% of the total number of trips. In this case, 5% of the total trips being in-advance and concentrated in UST means that approximately 50% of the UST trips are in-advance.

** In simulation, a request is in-advance request with the aforementioned probability. The choice of a request is independent to others.

Table 3: The Comparison of Scenarios

Comparison	Base Scenario	New Scenario(s)	Related Section	Purpose of Comparision
Comparison 1	Scenario 1	Scenario 2a and Scenario 2b	Section 4.4.1	Spatial Granularity for Aggregate Information
Comparison 2	Scenario 2b	Scenario 3a and Scenario 3b	Section 4.4.2	Degree of Dynamism for Individual Information
Comparison 3	Scenario 3a	Scenario 4	Section 4.4.3	Spatial Disparity of Individual Information

Figure 1: Under-served satellite town is an agglomeration in the southwest corner of the case study area.

their requests in advance. The partially dynamic AMoD system then serves a mix of on-demand and in-advance requests in real time.

We first compare the fully dynamic system (Scenario 2b) with two less dynamic systems: (1) Scenario 3a: 5% of the travelers randomly selected from the entire CSA are in-advance users; (2) Scenario 3b: 10% of the travelers randomly selected from the entire CSA are in-advance users. The results from above are also compared with Scenario 4. In Scenario 4, we also have 5% of the travelers sending requests in advance, but all of them are from an underserved satellite town (UST, see Figure 1). This scenario is especially helpful to understand the spatial disparity of service. The satellite town in this case study is an agglomeration in the southwest corner of the CSA, which accounts for 10% of the total trips. Travelers from UST are particularly underserved because the town is distant from the main concentration of demand and vehicles are not always available. The simulated experiments indicate that the service rate in UST is below 50% if both optimal rebalancing and in-advance requesting are disabled.

We also assume that the reaction time is 1800 seconds for all in-advance requests. That being said, the in-advance travelers will be notified of the assignment decisions 1800 seconds before their planned departure time. Under these assumptions, the degree of dynamism of a system with 100% in-advance requests is 0.5. For Scenarios 3a, 3b and 3c, the degrees of dynamism are 0.975, 0.950 and 0.975 respectively. We notice that the degree of dynamism reflects the percentage of in-advance requests in the system; however, the spatial distribution of the in-advance requests is omitted in this metric, and 3a and 3c have the exactly same degree of dynamism. This observation requires further discussion together with the simulation results.

The adoption of in-advance requests is not widely seen in real-world applications. However, as the share of the MoD service continues to grow, we expect that differentiated products such as in-advance requests and prioritized services will emerge in the market. The data collected during the service will also contribute to individual trip prediction. This in return pushes forward the intelligent operation and product differentiation. In this case, travelers will also have the freedom to choose from a variety of products. When both on-demand and in-advance requests are available, Scenario 4 appears to be more realistic, since travelers from the under-served areas are more likely to take advantage of the in-advance requests in order to secure service, while others may prefer flexibility [51].

4.4. Results and Analysis

4.4.1. Aggregate Information

Figure 2 illustrates the impact of aggregate demand information on the performance of the simulated AMoD system. Three indicators proposed in Section 3.4 are used in the following analysis: the average number of served requests as the system-level objective, adjusted wait time for travelers, and profit for the AMoD operator.

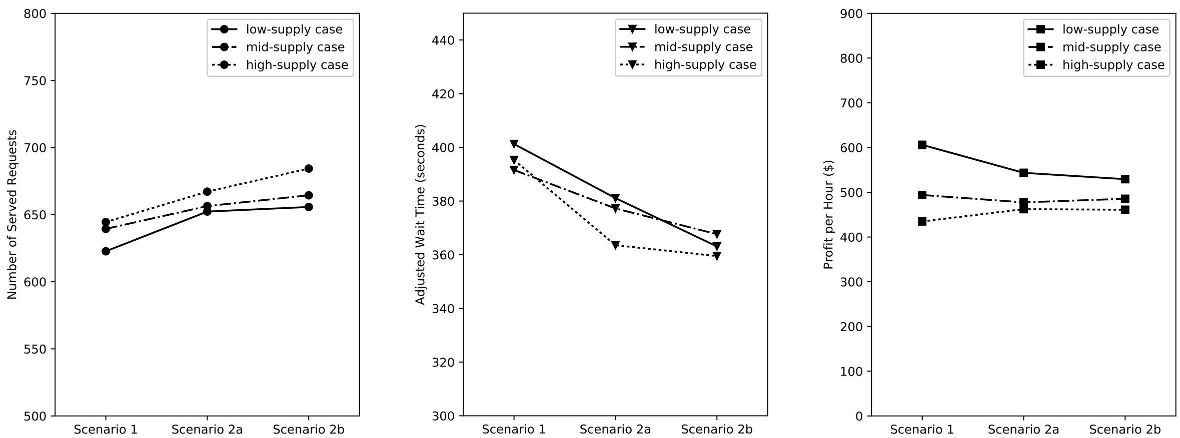

(a) On number of served requests. (b) On adjusted wait time. (c) On operator's profit.

Figure 2: Impact of aggregate demand information on system performance.

Figure 2a shows that, in spite of the different operational cases, rebalancing has a generally positive impact on the system, in the sense that it increases the number of requests served during the period of study. The more granular the aggregate demand information is, the more effective the rebalancing is, and the higher the ridership will be. Effective rebalancing also leads to shorter adjusted wait times as shown in Figure 2b. This could be explained by two reasons: (a) rebalancing moves idle vehicles closer to the potential demand and reduces the average wait time; (b) higher ridership imposes less penalty on the adjusted value.

The fleet size is also one of the determining factors for system performance. As the fleet size becomes larger, the supply gradually outgrows the demand. Consequently, there is a general trend that travelers will have a higher chance of getting served; the adjusted wait time will also be reduced in the meantime.

From the operator's perspective, applying rebalancing strategies does not always imply high profitability as in Figure 2c. On the one hand, rebalancing contributes to the increased revenue as the system serves more travelers during the period of study. On the other hand, rebalancing inevitably increases the operational cost, since it induces empty rebalancing distance. In the high-supply case, the fleet experiences considerably high idleness. Rebalancing will help make good use of the idle vehicles and turns out to be profitable. However, in the low-supply case, the revenue growth is little and could not compensate for all the induced cost of rebalancing. The operator's profit decreases as a result.

If the AMoD operator was profit-driven, it would decide to deploy a small fleet and operate without rebalancing (as Scenario 1 in the low-supply case) in order to maximize its return. This decision would be unfortunately against the interests of travelers and the system as a whole. In this paper, we assume a non-profit-driven operator, who may choose a larger fleet and enable rebalancing for the public good (as Scenario 2b in the high-supply case, which maximizes the number of served requests and minimizes the wait time). The profit generated in this case, however, is the lowest. Consequently, the governmental organizations may subsidize the operator in recognition of its loss.

4.4.2. Individual Information

We also compare the different degrees of dynamism and present the results in Figure 3.

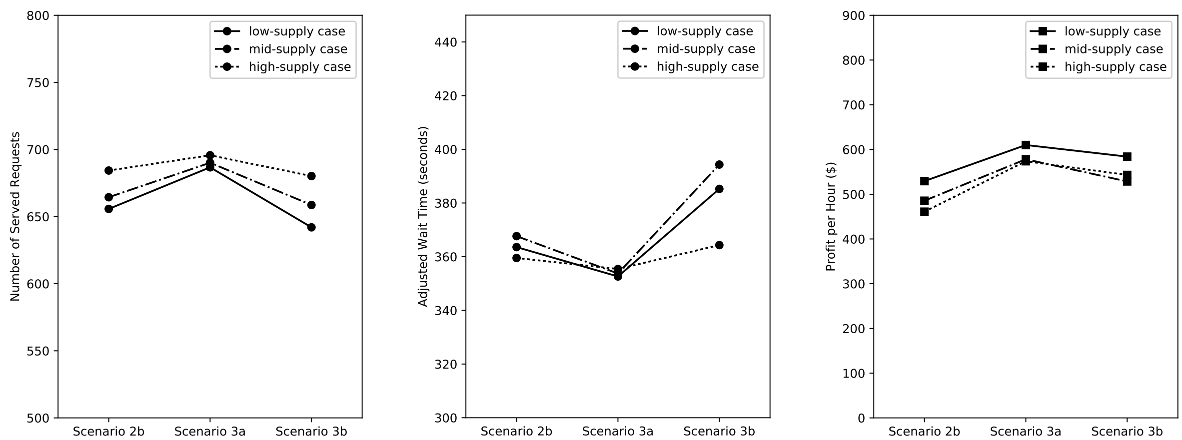

(a) On number of served requests. (b) On adjusted wait time. (c) On operator's profit.

Figure 3: Impact of the degree of dynamism of individual demand information on system performance.

Figure 3a and Figure 3b show the changes in the number of served requests and adjusted wait time. As the percentage of in-advance requests grows from 0% to 5% (from Scenario 2b to Scenario 3a), the number of served requests increases, which meets our expectation that information brings value. However, as the percentage of in-advance requests continues to grow from 5% to 10% (from Scenario 3a to Scenario 3b), its impact on the system becomes surprisingly negative: the number of served requests turns downward, and the adjusted wait time increases significantly. This negative value of information can be explained by the associated service priority. In Section 3.2, the assumption has been made that a traveler sending a request in advance expects no changes after it has been assigned a vehicle. In this case, the in-advance requests are naturally prioritized in the assignment and guaranteed service afterwards. When in-advance requests are many, this imposes binding constraints on the dispatching and may do harm to the system-level performance.

A solution to this problem is to increase the supply level. As we can see from Figure 3a and Figure 3b, the negative impact on the system is less significant when it's running with high supply, in which case the vehicles are abundant and could handle in-advance requests with more ease. Applying metaheuristic assignment algorithms that reoptimize the pairing solutions may also be of help. In fact, the implemented insertion heuristic processes the incoming requests one after another and may be caught in local optima. Reoptimization provides a scheme to search with larger search space, improving the existing solutions and making the dispatching less constrained by trips already assigned service. However, due to the limited computational capacity, we do not use any reoptimization methods in this paper.

The decrease in the total number of served travelers leads to the decrease in profit, as shown in Figure 3c. However, this decrease is also less important and the second segments of the curves are almost horizontal. This might be due to the fact that having more demand information a priori makes ride-sharing more efficient. If the AMoD operator was profit-driven, allowing in-advance requests could be even more profitable. In fact, the in-advance travelers benefit from the prioritized service and create negative externalities to the rest of the system. It is justifiable that they should be charged differently for this. In the long run, an operator may also consider extending its product line, which helps differentiate travelers that have higher willingness-to-pay and generate more added value in the service.

4.4.3. Spatial Disparity of Individual Information

Lastly, we compare Scenario 3a with Scenario 4. In both scenarios, 5% of the total requests are assumed to be in advance. Consequently, their degrees of dynamism are identically 0.975 according to Equation 1. The only factor that makes a difference in this comparison is the spatial distribution of the in-advance requests: the former scenario distributes the in-advance requests uniformly to the entire CSA, while the latter concentrates them into the under-served satellite town.

Table 4: Impact of Request Distribution on System Performance and Value of Information

Operational Case	Scenario 2b	Scenario 3a	Scenario 4
low-supply	657	687 (4.6%)	694 (5.6%)
mid-supply	664	690 (3.9%)	705 (6.2%)
high-supply	684	696 (1.8%)	722 (5.6%)

* The upper value in each cell indicates the number of served requests; the lower (bracketed) value in each cell indicates the value of information using the upper one as the objective.

Table 4 presents the performance of Scenario 3a and Scenario 4 and gives the value of information using Scenario 2b as the benchmark. It shows that, by limiting in-advance requests within the boundary of UST instead of having them randomly distributed in CSA, the number of served requests per hour increases significantly for each of the three operational cases. If we use the number of served requests as the objective F in Equation 2 and Scenario 2b as the base scenario, the value that the in-advance requests in UST bring to the system is more than 5.8% on average. As a comparison, the value of information in Scenario 3a is merely 3.4%. The profit that the operator makes also increases comparatively.

The distribution of information matters. However, the general metric 'degree of dynamism' failed to reflect the spatial disparity in the AMoD operation. One potential way of improving this is by adding a weight to requests to describe their difference in importance, priority and associated service rate. However, this requires field knowledge and would be feasible only when operational data are available. In practice, the operator should also deliberately design the hailing policies to match the demand patterns. This includes allowing in-advance requests in specific areas, imposing change fee and cancellation fee, as well as pricing for differentiated service.

5. Conclusion and Discussion

5.1. Summary of Findings

This paper presents a simulation-based approach that classifies, measures and evaluates the demand information at both individual and aggregate levels. The approach is tested in a major European city using scenarios with different information settings in terms of level, granularity and dynamism. It also compares various fleet sizes to study the supply-demand interaction. Results show that aggregate demand information leads to more served requests, shorter wait time and higher profit through effective rebalancing, especially when supply is high and demand information is spatially granular. Individual demand information from in-advance requests also improves the system performance, the degree of which depends on the spatial disparity of requests and their coupled service priority. By designing hailing policies accordingly, the operator is able to maximize the potential benefits.

5.2. Policy Discussion

The demand information for AMoD services, similar to all business information sources in the age of big data and information technologies, can be treated as a commodity that has inherent value and is produced, transformed, stored, commercially traded and used to generate added value for consumers, operators, and society at large. Regulatory considerations with respect to proprietary rights, quality standards, transaction and storage security measures commonly practiced for commercial goods and services should be discussed for the demand information as well. The authorities should be responsible for regulating the AMoD demand information and should work with the industry and consumer groups to design and implement standards for privacy protocols, security measures for storing and processing private, sensitive, or personally identifiable information, and guidelines for transferring and sharing data and developing legitimate use cases of data by entities that have access to the collected and processed user data.

This section offers a qualitative discussion on the stakeholders involved in the AMoD services, namely the users, the operators and the society at large, and for each of them, the benefits and costs associated with extracting, transforming, storing and using the demand information, as illustrated in Table 5.

Table 5: Benefits and Costs of Collecting and Utilizing Demand Information for Key Stakeholders

Stakeholder	Benefits	Costs
Users	short waiting times reliable service high service rates possible discounts	compromised flexibility and spontaneity risk of data privacy
Operator	operational efficiency increased revenue increased demand/economies of scale	data collection/management efforts service guarantees
Society	decreased private mode shares sustainability/operational efficiency high accessibility	risk of data privacy increased vehicle miles and congestion

As demonstrated in the simulation experiments in this paper, granular and in-advance demand information can help improve the level of service by reducing the waiting times, increasing the service rate, and improving the utilization of the AMoD vehicles. As a result, users would benefit from more timely and reliable AMoD service, which is important since the out-of-vehicle times are typically perceived as more onerous than in-vehicle travel times. Regarding the potentially negative aspects of collecting and utilizing demand information, the users may perceive a risk of compromised privacy as the result of their trip itineraries being collected, accumulated and processed. A different aspect is a trade-off between service priority and service flexibility: requesting trips in-advance increases the service priority for the user at the price of the flexibility and spontaneity. Requesting trips in advance will require planning and more rigid travel schedules. This is more of an individuals choice rather than a negative aspect.

From the operator's point of view, the better level of service and the higher vehicle utilization thanks to the demand information will improve the attractiveness and competitiveness of AMoD over alternative transportation modes and the overall transportation accessibility of the service area, both resulting in higher demand for AMoD service. In addition to the direct increase in the demand and revenue, the operator also benefits from the economies of scale that will improve operational efficiency and profitability even further [51]. Depending on the market share, competitive landscape, and business model, the operators may consider offering discounts to the users who share their travel information or opt into advance request schemes as leverage for more in-advance information, which stimulates a cycle of more information, better service, increased demand, and profitability. Collecting, processing, analyzing, and protecting the data of demand information may require significant investment in human resources, and data and computational infrastructure. To utilizing the data effectively also requires advanced analytic capacity and operational agility. As likely the main investor of the required technical, intellectual and infrastructural resources for the demand information system, the operator need to recognize and strike a balance between value creation, cost, and risk exposure.

The society as a whole is an important stakeholder in the equation. A well designed and operated AMoD service, strengthened by an advanced demand information system, may bring higher vehicle occupancy and lower empty vehicle kilometers traveled, encourage multimodality and shared mobility, and help transform the SOV dominated transportation system. If well designed and regulated, the improvement in the AMoD operation efficiency can also be channeled to increase accessibility to jobs, education, and health care, and promote economic prosperity, mobility equity, and social inclusion. Shifting demand from private single occupancy vehicles (SOV) is applaudable but we should also pay attention to the possible demand shifts from active and sustainable modes, such as walk, bike, and transit, and the negative impacts of the increased congestion and emissions. [51] demonstrate that there are ways in which the authorities can mitigate the negative modal shifts such as introducing high base fare and minimum distance for AMoD service to discourage short walk and bike trips from shifting to AMoD.

In certain scenarios, it can be argued that the transportation data is a public good, and therefore the authorities may take action to establish the data and modeling infrastructure to improve the transportation system as a whole. For example, transportation authorities in many state and city governments often develop and maintain travel demand models for strategic planning purposes. They may provide the baseline and forecast population, employment, and land use data, aggregate travel demand data, and transportation network data to the operators to improve their business intelligence and operation efficiency. Other data sources such as transit smart card data, automatic plate number recognition systems, automatic traffic flow data can also be useful [53]. In return, the operator may share data with the authorities to help them understand the sustainability, accessibility and equity consequences of the AMoD services. Of course, it is naive to simply ask all parties to share data with each other. Examining today's interaction between transportation network companies such as Uber and Lyft, public transit operators, and transportation authority, one witnesses a huge amount of confusion, ad-hoc negotiations and transactions, and even distrust and blaming between them. How to design an incentive structure so that all parties are willing to participate is a critical technical, policy, and legislative question. This paper only intends to illustrate the value of demand information, recognize that the property right of the transportation data needs to be defined and regulated, and call for innovation in both policy and technology to design and implement the mechanism for the smooth and fair exchange of transportation information between the private and public sectors.

5.3. Future Research

We identify four main areas of future research as below.

First, this paper focuses on three attributes of demand information: level, dynamism, and granularity, but has not discussed in-depth other important attributes such as accuracy and reliability. It is important to establish a more comprehensive framework of demand information, including the full set of attributes as well as the relationship and trade-offs between them. Conceptually it will be valuable to position this framework in the formal information theory.

The second limitation is regarding the dispatching algorithms in their relations with different types of demand information. Some algorithms are more capable than others when dealing with in-advance information with uncertainty. For example, metaheuristics that re-optimize the solutions may improve the system performance and alleviate the drawbacks of the prioritized requests, fitting with a much wider range of degrees of information dynamism. In addition, stochastic algorithms should be carefully investigated because they have the advantages of being compatible with stochastic features in the AMoD service as well as the probabilistic representation of the demand information.

Third, the paper makes the strong assumptions about the AMoD market structure, fare schemes, hail policies, and sharing policies. The findings of the value of information depend on these assumptions. It is important to examine a wide range of business models, as well as different fare, hailing and sharing policies, under which the value of demand information could vary. The diverse business and regulatory settings will also entail more complex interest structures of different stakeholders, and the value of the information needs to be carefully defined to reflect these complexities.

The last is to understand the cost of obtaining the demand information. This paper focuses on the value of demand information for key stakeholders such as the users, the operator and the society. By combining the cost and the value, we can then optimize the extent to which AMoD operators should spend effort collecting data, inferring information, managing the communication flow with the travelers.

Acknowledgment

The authors thank our research sponsor for providing financial support and travel data. We also thank Leo Yu Xin Chen, Nate Bailey, and members from MIT Mobility Lab (mobility.mit.edu) and MIT Transit Lab (transitlab.mit.edu) for their comments.

References

[1] Y. Shen, H. Zhang, J. Zhao, Integrating shared autonomous vehicle in public transportation system: A supply-side simulation of the first-mile service in singapore, Transportation Research Part A: Policy and Practice 113 (2018) 125–136.

[2] D. J. Fagnant, K. M. Kockelman, Dynamic ride-sharing and fleet sizing for a system of shared autonomous vehicles in austin, texas, Transportation (2016) 1–16.

[3] K. A. Marczuk, H. S. S. Hong, C. M. L. Azevedo, M. Adnan, S. D. Pendleton, E. Frazzoli, et al., Autonomous mobility on demand in simmobility: Case study of the central business district in singapore, in: Cybernetics and Intelligent Systems (CIS) and IEEE Conference on Robotics, Automation and Mechatronics (RAM), 2015 IEEE 7th International Conference on, IEEE, 2015, pp. 167–172.

[4] L. M. Martinez, G. H. Correia, J. M. Viegas, An agent-based simulation model to assess the impacts of introducing a shared-taxi system: an application to lisbon (portugal), Journal of Advanced Transportation 49 (2015) 475–495.

[5] L. M. Martinez, J. M. Viegas, Assessing the impacts of deploying a shared self-driving urban mobility system: An agent-based model applied to the city of lisbon, portugal, International Journal of Transportation Science and Technology 6 (2017) 13–27.

[6] J. Yang, J. Dong, W. Wang, A Simulation Model for Performance Analysis of Electric Autonomous Taxi Systems, Technical Report, 2017.

[7] K. A. Marczuk, H. S. Soh, C. M. Azevedo, D.-H. Lee, E. Frazzoli, Simulation framework for rebalancing of autonomous mobility on demand systems, in: MATEC Web of Conferences, volume 81, EDP Sciences, 2016, p. 01005.

[8] J. Jung, R. Jayakrishnan, J. Y. Park, Dynamic shared-taxi dispatch algorithm with hybrid-simulated annealing, Computer-Aided Civil and Infrastructure Engineering 31 (2016) 275–291.

[9] J. Alonso-Mora, S. Samaranayake, A. Wallar, E. Frazzoli, D. Rus, On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment, Proceedings of the National Academy of Sciences (2017) 201611675.

[10] J. Chen, K. H. Low, Y. Yao, P. Jaillet, Gaussian process decentralized data fusion and active sensing for spatiotemporal traffic modeling and prediction in mobility-on-demand systems, IEEE Transactions on Automation Science and Engineering 12 (2015) 901–921.

[11] J. Miller, A. Hasfura, S.-Y. Liu, J. P. How, Dynamic arrival rate estimation for campus mobility on demand network graphs, in: Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on, IEEE, 2016, pp. 2285–2292.

[12] M. K. Chen, M. Sheldon, Dynamic pricing in a labor market: Surge pricing and flexible work on the uber platform., in: EC, 2016, p. 455.

[13] H. Qiu, R. Li, J. Zhao, Daily level dynamic pricing in shared mobility-on-demand service, Transportation Research Board 97th Annual Meeting Compendium of Papers (2018).

[14] D. Levinson, The value of advanced traveler information systems for route choice, Transportation Research Part C: Emerging Technologies 11 (2003) 75–87.

[15] S. Gao, H. Huang, Real-time traveler information for optimal adaptive routing in stochastic time-dependent networks, Transportation Research Part C: Emerging Technologies 21 (2012) 196–213.

[16] C. G. Chorus, J. L. Walker, M. Ben-Akiva, A joint model of travel information acquisition and response to received messages, Transportation Research Part C: Emerging Technologies 26 (2013) 61–77.

[17] L. Zhang, D. Levinson, Determinants of route choice and value of traveler information: a field experiment, Transportation Research Record: Journal of the Transportation Research Board (2008) 81–92.

[18] C. G. Chorus, J. L. Walker, M. E. Ben-Akiva, The value of travel information: A search-theoretic approach, Journal of intelligent transportation systems 14 (2010) 154–165.

[19] B. Van Wee, K. Geurs, C. Chorus, Information, communication, travel behavior and accessibility, Journal of Transport and Land Use 6 (2013) 1–16.

[20] M. D. Hickman, N. H. Wilson, Passenger travel time and path choice implications of real-time transit information, Transportation Research Part C: Emerging Technologies 3 (1995) 211–226.

[21] G. Gentile, S. Nguyen, S. Pallottino, Route choice on transit networks with online information at stops, Transportation science 39 (2005) 289–297.

[22] K. Nökel, S. Wekeck, Boarding and alighting in frequency-based transit assignment, Transportation Research Record: Journal of the Transportation Research Board (2009) 60–67.

[23] O. Cats, H. N. Koutsopoulos, W. Burghout, T. Toledo, Effect of real-time transit information on dynamic path choice of passengers, Transportation Research Record 2217 (2011) 46–54.

[24] A. Fonzone, J.-D. Schmöcker, Effects of transit real-time information usage strategies, Transportation Research Record: Journal of the Transportation Research Board (2014) 121–129.

[25] N. Nassir, J. Zhao, J. Attanucci, F. Salvucci, N. Wilson, Bayesian inference of passenger boarding strategies at express stops with real-time bus arrival information, Transportation Research Part C, under review (2017).

[26] L. Tang, P. V. Thakuriah, Ridership effects of real-time bus information system: A case study in the city of chicago, Transportation Research Part C: Emerging Technologies 22 (2012) 146–161.

[27] C. Brakewood, G. S. Macfarlane, K. Watkins, The impact of real-time information on bus ridership in new york city, Transportation Research Part C: Emerging Technologies 53 (2015) 59–75.

[28] C. Brakewood, S. Barbeau, K. Watkins, An experiment evaluating the impacts of real-time transit information on bus riders in tampa, florida, Transportation Research Part A: Policy and Practice 69 (2014) 409–422.

[29] P. Toth, D. Vigo, The vehicle routing problem, Society for Industrial and Applied Mathematics, 2002.

[30] A. Borodin, R. El-Yaniv, Online computation and competitive analysis, cambridge university press, 2005.

[31] P. Jaillet, M. R. Wagner, Generalized online routing: New competitive ratios, resource augmentation, and asymptotic analyses, Operations research 56 (2008) 745–757.

[32] V. Pillac, M. Gendreau, C. Guéret, A. L. Medaglia, A review of dynamic vehicle routing problems, European Journal of Operational Research 225 (2013) 1–11.

[33] S. Mitrović-Minić, R. Krishnamurti, G. Laporte, Double-horizon based heuristics for the dynamic pickup and delivery problem with time windows, Transportation Research Part B: Methodological 38 (2004) 669–685.

[34] M. Gendreau, F. Guertin, J.-Y. Potvin, E. Taillard, Parallel tabu search for real-time vehicle routing and dispatching, Transportation science 33 (1999) 381–390.

[35] M. Tagmouti, M. Gendreau, J.-Y. Potvin, A dynamic capacitated arc routing problem with

time-dependent service costs, Transportation Research Part C: Emerging Technologies 19 (2011) 20–28.

[36] A. Larsen, O. Madsen, M. Solomon, Partially dynamic vehicle routing-models and algorithms, Journal of the operational research society (2002) 637–646.

[37] M. Diana, The importance of information flows temporal attributes for the efficient scheduling of dynamic demand responsive transport services, Journal of advanced Transportation 40 (2006) 23–46.

[38] K. Wong, A. Han, C. Yuen, On dynamic demand responsive transport services with degree of dynamism, Transportmetrica A: Transport Science 10 (2014) 55–73.

[39] M. Pavone, S. L. Smith, E. Frazzoli, D. Rus, Robotic load balancing for mobility-on-demand systems, The International Journal of Robotics Research 31 (2012) 839–854.

[40] R. Zhang, M. Pavone, Control of robotic mobility-on-demand systems: a queueing-theoretical perspective, The International Journal of Robotics Research 35 (2016) 186–203.

[41] K. Spieser, S. Samaranayake, W. Gruel, E. Frazzoli, Shared-vehicle mobility-on-demand systems: a fleet operators guide to rebalancing empty vehicles, in: Transportation Research Board 95th Annual Meeting, 16-5987, 2016.

[42] J. Wen, amod-abm, `https://github.com/wenjian0202/amod-abm`, 2017.

[43] C. Chen, J. Ma, Y. Susilo, Y. Liu, M. Wang, The promises of big data and small data for travel behavior (aka human mobility) analysis, Transportation research part C: emerging technologies 68 (2016) 285–299.

[44] N. Van Oort, T. Brands, E. de Romph, Short term ridership prediction in public transport by processing smart card data, Transportation Research Record (2015).

[45] Y. Wang, N. J. Yuan, D. Lian, L. Xu, X. Xie, E. Chen, Y. Rui, Regularity and conformity: Location prediction using heterogeneous mobility data, in: Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, 2015, pp. 1275–1284.

[46] C. Zhong, M. Batty, E. Manley, J. Wang, Z. Wang, F. Chen, G. Schmitt, Variability in regularity: Mining temporal mobility patterns in london, singapore and beijing using smart-card data, PloS one 11 (2016) e0149222.

[47] Z. Zhao, H. N. Koutsopoulos, J. Zhao, Individual mobility prediction using transit smart card data, Transportation Research Part C: Emerging Technologies 89 (2018) 19–34.

[48] D. McFadden, The measurement of urban travel demand, Journal of public economics 3 (1974) 303–328.

[49] M. E. Ben-Akiva, S. R. Lerman, S. R. Lerman, Discrete choice analysis: theory and application to travel demand, volume 9, MIT press, 1985.

[50] C. Daganzo, Multinomial probit: the theory and its application to demand forecasting, Elsevier, 2014.

[51] J. Wen, Y. X. Chen, N. Nassir, J. Zhao, Transit-oriented autonomous vehicle operation with integrated demand-supply interaction framework, 2017. Under review at Transportation Re-

search Part C.

[52] J. Wen, J. Zhao, P. Jaillet, Rebalancing shared mobility-on-demand systems: A reinforcement learning approach, in: Intelligent Transportation Systems (ITSC), 2017 IEEE 20th International Conference on, IEEE, 2017, pp. 220–225.

[53] R. Kitchin, The real-time city? big data and smart urbanism, GeoJournal 79 (2014) 1–14.