Using Bilateral Trading to Increase Ridership and User Permanence in Ridesharing Systems

April 15, 2017

Neda Masoud
Assistant Professor
Department of Civil and Environmental Engineering
University of Michigan, Ann Arbor, MI, USA, 48109
nmasoud@umich.edu
Corresponding author

Roger Lloret-Batlle
Ph.D. Candidate
Department of Civil and Environmental Engineering
University of California, Irvine, CA, USA, 92697
rlloretb@uci.edu

R. Jayakrishnan
Professor
Department of Civil and Environmental Engineering
University of California, Irvine , CA, USA, 92697
rjayakri@uci.edu

Abstract

One of the main obstacles that has challenged peer-to-peer (P2P) ridesharing systems in operating as stand-alone systems is reaching a critical mass of participants. Toward this goal, we propose what we call the *P2P ride exchange* mechanism to increase matching rate and customer retention in a ridesharing system. This mechanism gives riders the opportunity to purchase other riders' itineraries while it provides suitable alternative rides to the sellers, thus increasing the service rate in a ridesharing system. The proposed mechanism aims to maximize expected user surplus, is robust towards selfish user manipulation, and has very low information requirements. Using numerical experiments, we demonstrate what type of ridesharing systems can benefit the most from P2P ride exchange. Furthermore, we study the impact of customer flexibility on the rate of exchange. If implemented properly, P2P ride exchange can effectively increase the number of served riders and enhance customer loyalty by engaging customers in the ride-matching process.

Key words: Bilateral trading, Ridesharing, Ride-Matching, Mechanism Design, Ride Exchange, Posted-price mechanism

1 Introduction

Peer-to-peer (P2P) dynamic ridesharing is a shared mobility alternative in which peer drivers and riders (passengers) share the space in the drivers' personal vehicles. The term "P2P" implies that drivers are not hired by companies to transport passengers, but are rather using their personal vehicles to carry out their personal tasks, which makes them peers to riders. The term 'dynamic' highlights the fact that customers can join the system at any point in time and do not have to book their trips in advance.

P2P ridesharing manages to eliminate vehicles from roads by getting people who are traveling in the same direction in the same vehicle. P2P ridesharing benefits drivers, riders, non-users, the transportation infrastructure, and the environment. Drivers receive monetary compensation for the service they provide while following their own daily schedules, and riders are charged less than other transportation alternatives, such as taxis. By reducing the number of traveling vehicles and hence congestion levels, the benefits of P2P ridesharing are extended to the entire community as well as the environment.

Contrary to traditional service businesses where servers belong to the business and their number is proportional to the demand for service, in P2P ridesharing servers are also customers. Therefore, it is important for the system operator to attract the right proportion of riders and drivers. Another feature of P2P ridesharing systems is that drivers typically have specific locations where they start and end their trips, and tight travel time windows to carry out their trips. This limits the level of spatiotemporal coverage of the network by each driver. Therefore, in order to serve a higher number of riders, a ridesharing system needs to increase the spatiotemporal coverage of the network by increasing the number of drivers. To motivate, attract and retain a high number of drivers, a high number of riders is necessary. Therefore, the number of customers in a P2P ridesharing system should pass a certain critical mass with a specific proportion of drivers to riders, in order for the system to be able to operate independently and without a need for outsourcing supply.

A ride-matching algorithm is the engine of a P2P ridesharing system, determining how drivers and riders should be paired. Except for very simple and non-efficient ridesharing systems (which we will discuss later), the ride-matching problems are computationally hard to solve. A good ride-matching method is one that can provide the highest number of matches in an attempt to engage the highest number of customers and bypass the critical mass of participants.

Customer experience is another factor that plays a role in the success of a P2P ridesharing system, especially during the initial phases of implementing the system. A customer (rider or driver) may give the system a chance by attempting to use the system a few times, but if he/she is not matched, there is a possibility that such a customer would never return to the system. Therefore, it is essential for a P2P ridesharing system to involve and retain as many customers as possible.

Customers in many transportation systems are served on a first-come, first-served (FCFS) or a similarly pre-ordered basis. For P2P ridesharing, in which customer retention is especially important, considering riders on an FCFS basis is an inefficient use of the very limited available resources (drivers). The FCFS rule, however, is the natural order of serving riders in a dynamic system, where riders announce their trips not long before departure. In addition, dropping the FCFS principle may lead to high solution times for the resulting matching problem, and is therefore not an appropriate implementation strategy for a dynamic real-time system.

In this paper, we introduce what we call *P2P ride exchange*, a mechanism to improve the number of matches in an FCFS-based system. In a system where P2P ride exchange is implemented, riders will still be considered for service on an FCFS basis. Upon joining the system, a rider will be offered the best available itinerary, according to certain criteria which we will discuss later. However, if no match exists, the rider will be given the chance to buy a previously-matched rider's itinerary under specific circumstances. Purchasing an itinerary from a previously-matched rider is in fact reversing the FCFS rule. This exchange of rides is accompanied with an exchange of money through the system. Since the objective of the system from implementing the exchange mechanism is to increase the total number of matched riders, only riders for whom an alternative itinerary is available will receive a proposal to sell their current itineraries.

There are, admittedly, considerable regulatory obstacles to overcome for such P2P exchange or trade schemes to be used in transportation systems. The legal battles faced by ridesourcing firms are now well-known. Transportation supply being considered a public good, any breaking of the traditional FCFS operational paradigms also could face objections based on socio-political arguments of inequity across users. While important, such topics are considered beyond the scope of this paper that focuses only on showing the performance potential of the proposed scheme.

2 Related Work

P2P ridesharing systems are a member of the family of shared-use mobility alternatives. There is an abundance of work in the literature on the benefits ridesharing systems offer in terms of reduced direct and indirect cost to the environment and the society (Chan and Shaheen, 2012; Morency, 2007; Heinrich, 2010; Kelly, 2007). Despite these benefits, ridesharing operators have been facing multiple challenges in running ridesharing systems as stand-alone businesses. Furuhata et al. (2013) conduct a thorough survey of different types of ridesharing systems, and discuss some of the challenges that have prevented these systems from reaching their potential, despite the improvements in communication technology, prevalence of GPS-enabled cell-phones, and ease of developing cellphone applications that greatly facilitate participating in ridesharing systems.

Ultimately, for a ridesharing system to operate successfully, it has to attract and maintain a critical mass of customers. An essential challenge practitioners face is finding the most effective way to build this critical mass (Cervero and Griesenbeck, 1988; Brereton and Ghelawat, 2010; Raney, 2010). James Shield of Carma Technology Corporation which develops carpooling applications is quoted in an article by Gaynor (2015) to believe that although there is no definitive answer to this question, attracting a higher number of drivers, increasing marketing efforts, improving the technology, and attempting to use a societal/behavioral approach to engage people and make habits are all valid approaches.

Research in the field of marketing has found customer satisfaction, among other factors, to be a great predictor of customer retention rate (Gustafsson et al., 2005; Ranaweera and Prabhu, 2003; Rust and Zahorik, 1993). A satisfied customer not only has a higher probability of returning to the system, but also generates positive word of mouth (WOM) that helps in attracting new customers (Söderlund, 1998). Research has shown that WOM is a more important factor when it comes to deciding on services, rather than goods (Buttle, 1998). In addition, the cost of customer acquisition is about five times the cost of customer retention (Pfeifer, 2005), suggesting that a customer's first few experiences with the system play a central role in its long-term success. In light of these research studies, it is very important for ridesharing systems, especially in their initial stages, to serve as many ride requests at possible. In addition to a high matching rate, the responsiveness of the system to dynamic ride requests could play a role in customer satisfaction. Dynamic systems which try to address requests in real-time score high in this respect.

Encouraging a high number of drivers to participate in the system is another goal of a ridesharing system, albeit not as important as the first one. The reason is that firstly, drivers who participate in ridesharing usually receive a base fare regardless of the extent of their contribution. For a given level of demand (ride requests), there is an optimal amount of supply (drivers) above which the contribution of additional supply is only marginal, and therefore attracting drivers with only marginal contributions is not financially wise. Secondly, drivers in a ridesharing system are traveling to perform their personal activities. Even if they are not matched on a regular basis, entering their fixed daily schedules in the system only once could earn them extra revenue.

The ride-matching algorithm used by a ridesharing system plays an important role in the number of riders the system can serve. The simplest form of ride-matching algorithm matches each driver with a single rider (Agatz et al., 2011; Wang, 2013). This problem can be formulated as a maximum cardinality matching problem on a bipartite graph and solved quickly using efficient algorithms (Alt et al., 1991). The more sophisticated ride-matching problems are capable of allocating more than one rider to a single driver (Wolfler Calvo et al., 2004; Baldacci et al., 2004; Teodorović and Dell'Orco, 2005; Herbawi and Weber., 2012; Di Febbraro et al., 2013), proposing multi-hop itineraries to a single rider, where the rider can transfer between multiple drivers (Agatz et al., 2009; Masoud and Jayakrishnan, 2017; Ghoseiri, 2013; Herbawi and Weber., 2011a,b), and finally considering multiple riders and drivers in the same problem, and proposing multi-hop itineraries at the same time (Masoud and Jayakrishnan, 2017; Ghoseiri, 2013; Regue et al., 2016).

A many-to-many ride-matching problem, where a rider can transfer between multiple drivers and a driver can carry multiple passengers at any moment in time, is the most comprehensive form of ride-matching and can yield the highest number of matches. Not surprisingly, a many-to-many problem is also the hardest matching problem to solve. For a ridesharing system to enjoy the benefits of many-to-many ride-matching, it should have access to information on future trips of riders and drivers. This property of a many-to-many problem coupled with its higher solution time prohibit such a system from being used in a dynamic setting where matches need to be made in real-time and information on future trips is not typically available. Many-to-many ride-matching is, however, very effective in static implementations of ridesharing, where system participants are required to announce their trips by a certain deadline.

In a many-to-one matching problem, a rider can transfer between drivers, but the matching problem is solved for one rider at a time. By definition, a many-to-one problem provides the best (multi-hop) solutions for a dynamic system where riders need to be informed of the status of their requests as soon as they input their register their trips. There are two ways, however, to shift the solution of a many-to-one problem towards that of a many-to-many problem in a dynamic system. After a rider is matched in a real-time system, the itinerary of the rider is fixed, and the drivers constructing the itinerary will have to remain committed to their assignments. It is possible, however, to include previously matched drivers with fixed itineraries (that respect their previous assignments) in the matching problem for the current rider, increasing the level of supply available to the current rider and hence enhancing our chances of satisfying their request while at the same time increasing vehicle occupancies. Introducing this small variation will transform a many-to-one system in which each driver carries one rider at a time, to a many-to-many system, in which each driver can have multiple passengers on board. Note that the type of matching method used in this case is still a many-to-one method.

The second way of shifting the solution of a many-to-one problem to that of a many-to-many problem is by implementing the P2P ride exchange mechanism introduced in this paper. Contrary to a many-to-many matching problem that assumes no particular order in serving requests and hence achieves a high matching rate as a result of introducing this additional level of flexibility, the nature of a dynamic system, which requires attending to requests as soon as they arrive, calls for an FCFS-based implementation. P2P ride exchange attempts to reverse the impact of the FCFS rule by proposing to a rider who has been offered an itinerary to switch to a less attractive itinerary, liberating the drivers contributing to the rider's current itinerary from their commitments. This exchange is motivated by a monetary compensation from a second rider who finds the liberated drivers more valuable. The P2P exchange mechanism proposes the amount of this compensation ensuring that the system will remain budget-balanced.

In this paper in order to match riders and drivers, we use the many-to-one ride-matching algorithm proposed by Masoud and Jayakrishnan (2015). We choose this algorithm because it can solve matching problems in real-time and can be easily modified for use in one-to-one and many-to-many ridesharing systems as well, giving us the ability to study the impact of the P2P ride exchange mechanism on a wide range of ridesharing systems. Furthermore, using this algorithm, not only can we find the optimal itinerary for a rider, but we can also identify and store other feasible itineraries that can be used later by the exchange mechanism.

There have been a few attempts in the literature to design mechanisms for para-transit and ridesharing systems. Furuhata et al. (2015) propose the Proportional Online Cost Sharing Mechanism for demand-responsive transport Systems. This mechanism is capable of proposing an upper bound on the fare a potential user has to pay. The mechanism relies heavily on having the passenger requests in advance of the start time of the trips. The focus of the work is on proving the online fairness, budget balancedness, individual rationality and ex-post incentive-compatible properties of the mechanism under certain conditions. The mechanism, however, unlike the P2P ride exchange mechanism proposed in this paper, is not designed to increase operational efficiency of the system. Wang (2013) proposes a stable matching game between riders and drivers in a one-to-one system, where no rider/driver can be better of by unilaterally switching to other drivers/riders. Although such a system can lead to an equilibrium, for it to yield operationally efficient results, it requires access to the participants' trip information in advance. Kleiner et al. (2011) proposes an auction-based allocation mechanism that incorporates users' valuations on the ride assignment. While this mechanism violates the FCFS rule, it is not real-time since it uses a rolling horizon in which decisions are delayed.

P2P ride exchange is the first real-time mechanism designed to address the inherent trade-off between the two factors that influence customer satisfaction in a ridesharing system, namely rider matching rate and system responsiveness. To the best of our knowledge, the proposed P2P ride exchange mechanism is the first trading mechanism to increase ridership in a dynamic P2P ridesharing system. The designed mechanism is limited to bilateral trades, where there is a single buyer and a single seller. This mechanism, therefore, is optimal for a one-to-one matching system, and provides a lower-bound on the increase in ridership in one-to-many and many-to-many systems. In the rest of this paper, we first provide a brief summary of the ride-matching algorithm used.

We then officially introduce the mechanism, and elaborate on some of its properties. Finally, we conduct extensive numerical experiments to quantify the performance of the P2P ride exchange mechanism under different parameter values for the system.

3 One-to-many Ride-Matching Algorithm

We solve the ridesharing problem using the dynamic programming (DP) algorithm proposed in Masoud and Jayakrishnan (2015). This algorithm is suitable for the purpose of P2P ride exchange because firstly, using this algorithm, real-life size problems can be solved in a very short period of time (a fraction of a second in most settings), and secondly, all feasible solutions to the problem are retrievable using the set of trees that are generated while solving the problem. In this section, we provide a brief review of the algorithm.

Let us define graph $G = (N, L)$. Each vertex $n \in N$ in this graph is a tuple $(s, t) \in S \times T$, where S is a pre-defined set of stations in the network where participants can start or end their trips and/or transfer between vehicles, and T is the set of time intervals during the study time horizon (set to be one minute in this study). An edge $\ell = (n_1, n_2) = (s_1, t_1, s_2, t_2) \in L$ in this graph corresponds to trip between stations s_1 and s_2 that begins at interval t_1 and ends at interval t_2. Each participant (rider or driver) upon registering in the system provides information on their origin and destination stations and their travel time window, which is bounded from below by the earliest departure time and from above by the latest arrival time of their trip.

For each rider r, a graph G_r can be constructed based on these parameters. A link on this graph exists if it is spatiotemporally accessible by at least one driver (i.e., both the rider and at least one driver can travel on that link). The algorithm searches on this graph for a minimum cost path that starts from the origin station of the rider and ends at his/her destination station. We define the cost of a path as a weighted linear combination of the in-vehicle travel time, waiting travel time, and number of transfers. Note that a rider can use multiple vehicles/modes of transportation to accomplish his/her trip.

The algorithm tries to find the best itinerary for rider r, by first topologically sorting the graph G_r, and then searching on this graph using a DP algorithm for an optimal path. The Bellman equation for the DP algorithm is presented in equation (1). $V(n_j, d)$ in this equation is the value of the minimum cost path from node 1 in the topologically ordered graph, to node n_j, with d being the last driver in the set of drivers that form the itinerary of the rider. The cost of each path is a linear combination of the travel time between nodes n_i and n_j, $C(n_i, n_j)$, which is itself a linear function of the in-vehicle and waiting travel times, and a fixed penalty, C_T, for each transfer. Set D_i^{in} denotes the set of drivers who enter node i, and set DN_j denotes a set of tuples (n_i, d), such that there is a link for driver d from node n_i to node n_j. Finally, set $ED(n_i, d')$ contains a list of drivers on the optimal path to node n_i, excluding the driver on the last link. The purpose of using this set is to prevent an itinerary to switch between the same set of drivers (e.g., traveling with driver 1, then transferring to driver 2, and then transferring to driver 1 again.) For more details on this algorithm refer to Masoud and Jayakrishnan (2015).

$$V(n_j, d) = min_{n_i:(n_i,d) \in DN_j} \left(min_{d' \in D_i^{in} \setminus ED(n_i,d')} (v(n_i, d') + C_T 1_{\{d \neq d'\}}) + C(n_i, n_j) \right) \tag{1}$$

The above algorithm has multiple practical advantages in the context of P2P ride exchange. First, the algorithm can be easily modified to make it suitable for one-to-one matching, either by setting C_T to a large value, or by modifying the input sets. Second, the algorithm can be used to run a many-to-many ridesharing system, by taking into consideration the previously matched drivers when building the graph G_r for rider r. Third, the trees generated during the iterations of

DP can be stored and used later to retrieve additional feasible itineraries for a rider in case he/she is a candidate to trade his/her itinerary for a sub-optimal one in P2P ride exchange.

4 Peer-to-Peer Ride Exchange

Dynamic ridesharing systems should have the capability of matching riders and drivers in real-time. Since participants in a dynamic ridesharing system announce their trips not long before they are ready for departure, the attempt to find a match for them should start as soon as the trip announcement is received by running the DP algorithm described in section 3. If all the itineraries generated by the algorithm are infeasible due to their conflicts with itineraries of the previously assigned riders (i.e., if the itineraries use the same drivers, but through different paths), then the system evaluates the possibility of a trade. In this section, we show through an example the benefits of a P2P exchange program, discuss the conditions under which trade can happen, and devise a mechanism that ensures a fair trade.

Let P^r denote the set of itineraries for rider r. Each itinerary has a value that is determined by a pre-specified objective function (the DP objective function), based on which the itineraries within P^r are ranked. Let p_i^r denote the i^{th} itinerary of rider r, and $d(p_i^r)$ denote the set of drivers who contribute to itinerary p_i^r. Note that there is no need to know all members of set P^r in advance, but we will generate them as (and if) needed. Furthermore, let p_k denote the itinerary of the assigned driver k.

Once rider r joins the system, the system uses the DP algorithm to generate a set of trees from which members of set P_r can be retrieved. The system starts by evaluating members of set P^r in order of their ranking. If an itinerary with no conflicts with the itineraries of previously matched drivers is found, this itinerary will be assigned to rider r. If the system exhausts all members of set P_r, and is not successful in finding a non-conflicting itinerary for rider r, then it considers the possibility of a trade.

Assume that rider 1 enters the system, and has two itineraries: $P_1 = \{p_1^1, p_2^1\}$, where $d(p_1^1) = \{d_1\}$ and $d(p_2^1) = \{d_2\}$. The left hand side picture in Figure 1 shows the rider and his itinerary set. Assuming that the minimum cost itinerary for this rider is the first one, this itinerary will be announced to both rider 1 and driver 1. Next, rider 2 joins the system. Because rider 1's itinerary has been fixed, there are no feasible itineraries for rider 2. However, rider 2 has a chance to buy rider 1's itinerary if rider 1 has not started his trip yet. The right hand side picture in Figure 1 shows this scenario after the trade. In this trade, rider 2 buys rider 1's assigned itinerary, and by doing so liberates driver 1, who in turn forms a feasible itinerary for rider 2. Rider 1 switches to a less convenient itinerary (with driver 2) in exchange for a monetary compensation. This trade's contribution to customer retention is double-folded. Not only are both riders served, but now both drivers are participating in the system as well.

Note that it is possible to obtain the same optimal solution by solving a many-to-many ride-matching problem that is capable of considering both riders at the same time. There are, however, two issues with such an approach: (1) An optimal matching algorithm that could consider both riders at the same time is computationally harder to solve (specially for real-world size problems), and therefore might not be able to yield solutions in real-time. (2) Even if the system is equipped with a many-to-many ride-matching algorithm that can yield solutions in a reasonable period of time, the information on the two drivers and the two riders need to be available in advance for the many-to-many matching problem to generate the solution that can serve both riders.

The system studies the possibility of a trade if the following three conditions hold. First, the buyer does not have any feasible itineraries; second, the seller has an alternative feasible itinerary

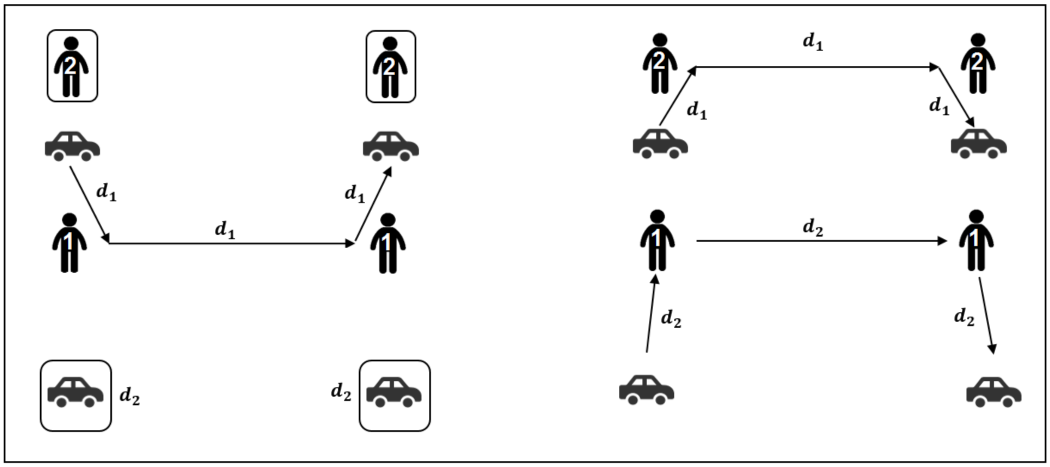

Figure 1: An example of a successful trade. Rider 1 arrives at the system first, followed by rider 2. Both drivers 1 and 2 are available at the time of rider 1's arrival. The figure on the left shows the state of the system without trade, where rider 1 is matched with driver 1 (with their itinirary demonstrated by solid lines), and rider 2 and driver 2 remain unmatched (in rectangles). The figure on the right shows the state of the system after trade, where all four participants are matched.

to his current one; and third, both parties will be better off with the trade than without it.

The monetary transfer from the buyer to the seller covers the additional cost the seller has to endure due to itinerary-switching. This cost includes the additional monetary cost due to a potentially increased travel distance, and a compensation to the seller for a potentially increased travel time. A proportion of this money will be used by the system operator to cover the cost of the seller's new itinerary that is now more expensive, and the rest will be transferred to the seller himself.

4.1 The Scope of the Trade

Assume a set of itineraries P^r for rider r. Drivers contributing to itinerary i are stored in set $d(p_i^r)$. Let us divide members of set $d(p_i^r)$ into two mutually exclusive sets, $d_a(p_i^r)$ and $d_f(p_i^r)$. Drivers in set d_a have been previously assigned to other riders, but their corresponding riders' trips have not started yet. Drivers in set d_f are free, and have not been assigned to any riders. The necessary condition for rider r to have a feasible itinerary is for at least one of the driver sets $d_a(p_i^r)$ and $d_f(p_i^r)$ to be non-empty. As the sufficient condition for p_i^r to be a feasible itinerary for rider r, one of the following conditions should hold: (1) $d_a(p_i^r) = \emptyset$, i.e., none of the drivers that contribute to the itinerary are assigned to other riders, and (2) $\forall k \in d_a(p_i^r)$, $p_i^r(k) \in p_k$, i.e., drivers in set $d_a(p_i^r)$ can still follow their previously assigned itineraries. If none of these two conditions hold, then the system tries to find a good candidate for a trade between the assigned riders.

To find the candidates for a trade, the system has to first identify the itinerary that rider r is interested in. It starts from the best itinerary, i.e., p_1^r, and moves to the next itinerary if the trade on the current itinerary is not possible. In order for the system to offer itinerary p_i^r to rider r, it has to liberate all the drivers in set $d_a(p_i^r)$ from their previous assignments. Therefore, the system has to find all the riders who are using these drivers, and find alternative itineraries for them as well. These riders form the sellers in the first level of trade (Figure 2a). In order for the system to propose an exchange to a rider r' in the first level of trade, it should find an alternative itinerary for this rider first. This task can be accomplished by identifying the set of assigned drivers for the

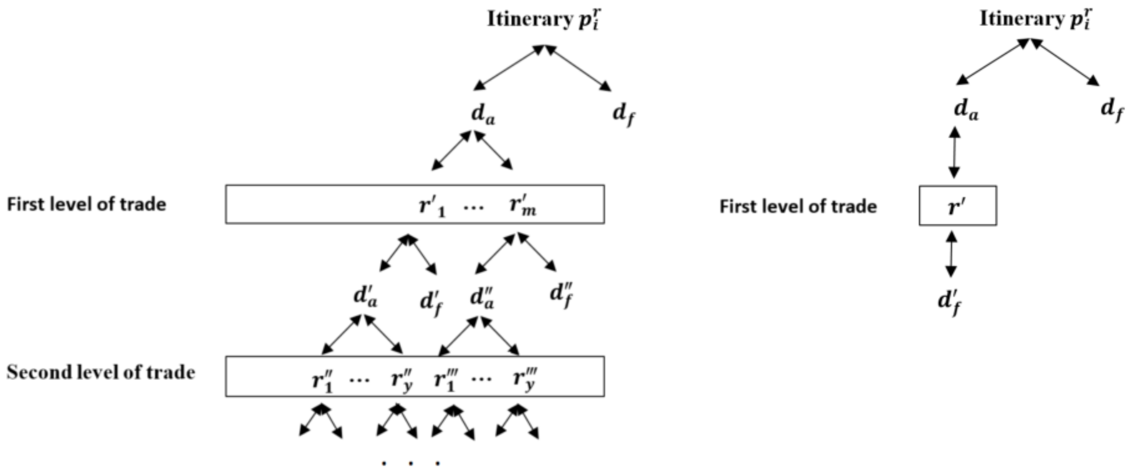

(a) General levels of trade

(b) Level of trade covered in this paper

Figure 2: Levels of trade

rider (d'_a), finding the rest of the riders whose itineraries are affected by these drivers, and finally finding alternative itineraries for them as well. This procedure continues until the system reaches a level of trade where all riders have itineraries with free drivers (or previously assigned drivers with non-conflicting assignments).

The system will then start proposing trades to riders, starting from those in the last level of trade. In order for a trade to be approved at any level, all the riders at that level should approve the trades proposed to them. For the n^{th} level of trade to take place, the trade at level $n+1$ should have been approved. Once all riders in a given level approve the proposed trades, the system can move to the next (higher) level of trade (moving upwards in Figure 2a). Therefore, it is clear that the more the levels of trade there are, the less likely it is for rider r to obtain itinerary p_i^r.

Another complication is that even if only one rider does not agree to the trade at a certain level, the trade cannot happen. In this case all the riders in the same and lower levels who have agreed to the trades proposed to them have to go back to their previous itineraries. Therefore, in order to simplify this procedure and make it easy to implement in practice, this paper only considers trades in settings where the level of trade is limited to 1, and the number of riders in the first level of trade is limited to 1 as well, i.e., the set of assigned drivers affect only the itinerary of a single previously assigned rider (Figure 2b). These simplifications limit the trade between two individuals only: the buyer, r, and the seller, r'.

Figure 3 displays two examples involving multilateral trade. In the first example (Figure 3a), in order for the system to serve rider 2 by liberating driver 1, it must find alternative itineraries for both riders currently served by driver 1. For this to happen, rider 2 should negotiate with both riders 1 and 3, which is beyond the scope of the bilateral trade covered in this paper. Notice that this example is still limited to the first level of trade (Figure 2a). Furthermore, even if a multilateral trade mechanism was available, in order for the trade to happen, both riders 1 and 3 should have agreed to the trade.

Figure 3b demonstrates an example of a simple scenario involving riders beyond the first level of trade. In this example, two alternative itineraries are available for rider 1, just as in Figure 1. When rider 2 joins the system, he is interested in purchasing the itinerary assigned to rider 1. In this example, however, driver 2 has been matched with rider 3, who belongs to the second level of

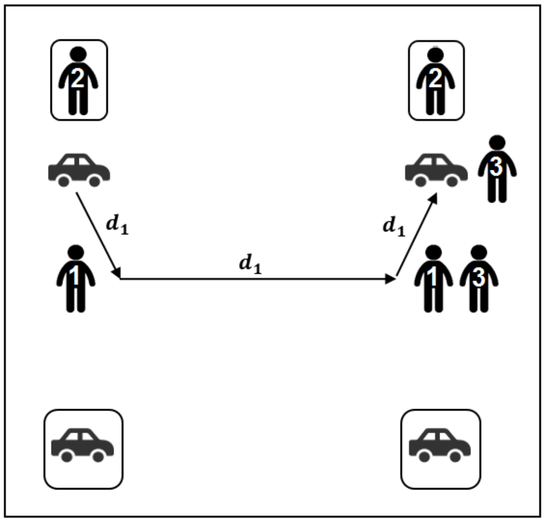

(a) A more complex trade at the first level: To release driver 1 from its previous assignment, the system has to find alternative itiniraries for both riders 1 and 3.

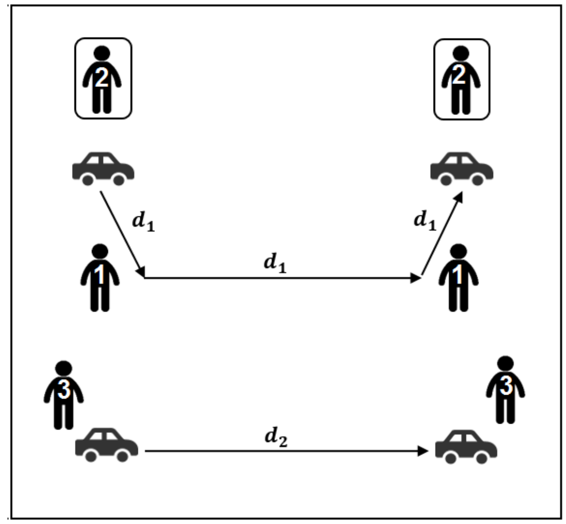

(b) A trade at the second level: To serve rider 2, driver 1 should be released from its previous assignment and matched with ridr 2, and driver 2 should be released from its previous assignment and matched with rider 1, extending the trade to the second level.

Figure 3: Examples of multilateral trade. Lines mark the itiniraries of participants, and rectangles identify un-matched participants.

trade. Therefore, for the trade to happen, the system should find an alternative itinerary for rider 3 first.

4.2 P2P Ride Exchange Mechanism

Besides ensuring that the trade makes both parties better off, the designer (operator) should also ensure that the trading parties cannot manipulate the outcome of the trade. Since both the buyer and the seller hold private information not known to the operator (e.g., their value of time (VOT)), this could lead to an inefficient outcome. This issue is addressed by modeling the trade from a mechanism design perspective. Informally, a mechanism is a method that defines rules for a game with incomplete information (Bayesian game) to influence agents' behavior and reach a particular goal, which in this case is efficiency maximization. Excellent introductions to mechanism design can be found in (Mas-Colell et al., 1995; Nisan and Ronen, 1999). The basic definitions are provided next, but a complete understanding of mechanism design may require reading the above introductory references.

Let $I = 1, \dots, n$ be the set of agents. Each agent has a type (value of time), $\theta_i \in \Theta_i$ which is private. $\Theta = \times_{(i \in I)} \Theta_i$ is the type profile set. Agent i has the (quasi-linear) utility function $u_i(\theta_i, \theta_{-i}; \theta_i) = v_i\big(k(\theta_i, \theta_{-i}); \theta_i\big) - p_i(\theta_i, \theta_{-i})$. Where $v_i(\boldsymbol{.})$ is his valuation and $p_i(\boldsymbol{.})$ is the price charged to him. The types to the left of the semicolon are the types announced to the designer, while the type at its right is the agent i's actual type. Making a distinction between the values of time reported to the designer and the actual value of time of an agent is important, since it allows us to take measures to ensure that agents cannot manipulate the system by falsely reporting their values of time. A (direct revelation) mechanism is composed of two interrelated functions. The first is an allocation function $k : \Theta \to K$ that maps the type space to an outcome set K. That is, for every announced type profile, an allocation $k(\theta) \in K$ is given. In the case of a bilateral trade, K is composed of the two allocations (trading states): either there is trade or there is not. The allocation rule that maximizes the sum of agents' valuations is the efficient allocation rule, $k^*(\theta) \in K$. Secondly, there is a payment function $p : \Theta \to R^N$. This function assigns a transfer amount to every agent i in accordance with its announced typeθ_i.

Mechanism design defines concepts that address how the strategic interests of agents are satisfied. The main one is truthfulness, or incentive compatibility, which states that truthful bidding forms an equilibrium. In other words, any participating agent is always better off by truthfully eliciting its type rather than lying, subject to others telling the truth. A mechanism (k, p) is (Dominant-Strategy) Incentive Compatible (DSIC) if:

$$v_i(k(\theta_i, \theta_{-i}); \theta_i) - p_i(\theta_i, \theta_{-i}) \geq v_i\big(k(\hat{\theta}_i, \theta_{-i}); \theta_i\big) - p_i(\theta_i, \theta_{-i}), \forall i{\in}I, \forall \theta_i{\in}\Theta_i, \forall \theta_{-i}{\in}\Theta_{-i}, \forall \hat{\theta}_i{\in}\Theta_i$$

Besides truthfulness, a designer is interested in the users' willingness to participate in the mechanism, called individual rationality. A mechanism (k, p) is Ex-Post Individual Rational (EPIR) if:

$$v_i\big(k(\theta_i, \theta_{-i}); \theta_i\big) - p_i(\theta_i, \theta_{-i}) \geq \bar{u_i}(\theta_i) \forall i{\in}I, \forall \theta_i{\in}\Theta_i, \forall \theta_{-i}{\in}\Theta_{-i}$$

where $\bar{u_i}(.)$ is agent i's utility from not participating in the mechanism. If EPIR is satisfied, an agent would be willing to truthfully participate in the mechanism rather than stay out. Finally, the designer may be interested in the mechanism being self-sufficient from a budget point of view. That is, the mechanism should not require an external subsidy to achieve the desired outcome. This condition is known as balanced budget. A mechanism (k, p) is strictly budget balanced (SBB) if:

$$\sum\nolimits_{(i \in I)} p_i(\theta_i, \theta_{-i}){=}0$$

The trade is modeled as a bilateral trade with private information (Hagerty and Rogerson, 1987; Myerson and Satterthwaite, 1983). We follow a "robust" implementation approach in which the designer attempts to maximize the expected surplus from the trade. It is assumed that the designer has a prior on the private information from agents, but the agents themselves do not have a prior of the other agents' type, unlike in common truthfulness concepts such as Bayesian Incentive Compatibility (Myerson and Satterthwaite, 1983). The designer proceeds to find the optimal posted price that maximizes expected surplus based on that information. This framework is very convenient for our purposes, since the mechanism has to be designed far in advance, with no previous experience or learning on trading outcomes from either the users' or the designer's part, while, at the same time, it aims an increase in the number of served riders to achieve user permanence in the system. Hagerty and Rogerson (1987) has shown that in this kind of a bilateral trading setting, any DSIC mechanism is a posted-price mechanism.

Let $I = 1, 2$ be the set of agents, $i = 1$ being the seller and $i = 2$ being the buyer. Each agent i has type $v_i = [\underline{v_i}, \overline{v_i}]$. These types are drawn from an empirical VOT distribution estimated from a survey on households conducted in Stockholm, Sweden in 2005 (Abou-Zeid et al., 2010). In that research, the Stated Preferences (SP) choice scenarios are composed of car alternatives that differ on attributes such as travel times and travel costs. Since only the main statistics are available in the publication, the distribution is recalibrated as a lognormal distribution given these statistics. Its parameters are location $\mu = 2.16$ and scale $\sigma = 0.40$.

The mechanism lies in the space $(q, p) \in [0, 1] \times R$, where q is the probability of trade and p is the payment from the buyer to the seller. For clarity in the exposition, we use the following change in notation $c_1 \overset{def}{=} -v_1$. c_1 is the opportunity cost of the seller. By definition, the bilateral trading setting satisfies the strict budget balance property, thus the seller has utility $u_1 = p - c_1 q$ and the buyer $u_2 = v_2 q - p$. Both agents are proposed the price p and if both agree, the trade takes place. This occurs when $v_2 > p > c_1$. The surplus of such as trade is $w\big((q,p);\theta\big) = (v_2 - c_1)q$.

Instead of valuing an object by a scalar as in the original bilateral trading environment, riders value their allocation (assigned ride) by its generalized cost, which is an affine transformation of their private type. For a rider i, this cost is the product of the travel time t_{ri} and the sum of the value of time θ_i, plus the fare per unit of time c_{ri}. These valuations are normalized with regard to the initial situation (no trade) to fit the bilateral trading original setting: c_1 and v_2 are in fact the valuation difference between states "trade" and "no trade". When there is a trade, $c_1(\theta_1)= \theta_1(t'_{r1} - t_{r1}) + c'_{r1}t'_{r1} - c_{r1}t_{r1}$ and $v_2(\theta_2) = \theta_2(t_{out} - t_{r2}) + c_{out}t_{out} - c_{r2}t_{r2}$. Here, t_{r1}, t'_{r1}, t_{r2}, and t_{out} refer to the travel time of the seller's current and new itineraries, travel time of the buyer's itinerary, and the travel time of the buyer's outside alternative, respectively. c_{r1}, c'_{r1}, c_{r2} and c_{out} are the costs per unit time of seller initial ride, seller alternative proposed ride, buyer proposed ride and the outside option cost to the buyer.

These time and cost variables have bounds and relative magnitudes. $t_{out} \leq t_{r2}$ since we consider the outside option to use the shortest path between buyer's origin and destination. We assume that cost of the rideshare option to the buyer is less than that of the outside option, i.e., $c_{out}t_{out} - c_{r2}t_{r2} \geq 0$; Otherwise the buyer would not have selected to use the rideshare option. This assumption leads $v_2 = [\underline{v_2}, \overline{v_2}]$ to be positive in our analysis. Note that as θ_2 increases, v_2 decreases, so $\underline{v_2} = max\big(0, v_2(\underline{\theta_2})\big)$ and $\overline{v_2} = v_2(\overline{\theta_2})$. Since the seller is offered a longer ride than the one he holds, $t'_{r1} \geq t_{r1}$ and $c'_{r1}t'_{r1} - c_{r1}t_{r1} \geq 0$, with a high probability. If due to higher number of riders involved in the new itinerary c'_{r1} becomes smaller than c_{r1}, the system will set $c'_{r1}t'_{r1} - c_{r1}t_{r1} = 0$. In this way, $c_1 \in [\underline{c_1}, \overline{c_1}] \geq 0$ and it is increasing with θ_1. Its bounds are $\underline{c_1} = c_1(\underline{\theta_1})$ and $\overline{c_1} = c_1(\overline{\theta_1})$. Without loss of generality, we assume $\underline{\theta_1} = \underline{\theta_2} = \underline{\theta}$ and $\overline{\theta_1} = \overline{\theta_2} = \overline{\theta}$. Trade is only possible when $\underline{v_2} \geq \overline{c_2}$. The next proposition formally presents the bilateral trade mechanism for a given price p and its properties:

Proposition 1. *The posted-price mechanism with price p for the bilateral trade setting is a revelation mechanism (q,p) such that:*

$$(q,p) = \begin{cases} (1,p) & v_2 > p > c_1 \\ (0,0) & otherwise \end{cases}$$

This mechanism is DSIC, EPIR, SBB and guarantees the following expected surplus $W(p)$:

$$W(p) = \iint\limits_{(c_1,v_2)=\left(c_1(\underline{\theta_1}),p\right)}^{(c_1,v_2)=\left(p,v_2(\underline{\theta_2})\right)} (v_2 - c_1)\phi(c_1,v_2)dv_2dc_1$$

Proof. Let the strategy of the seller s_{seller} be: sell if $p > c_1$ and do not sell if $p \leq c_1$. This is the only DSIC strategy for the seller (we assume that the seller prefers to keep the object when the gain is zero). Suppose that the seller follows a different strategy s'_{seller}: sell only when $p - c_1 > \epsilon, \epsilon > 0$. Then the seller would lose the opportunity to make profit $p - c_1$ when $c_1 + \epsilon > p > c_1$. Equivalently, when ϵ is negative, the seller would incur a loss of $c_1 - p$ when $c_1 + \epsilon < p < c_1$.

Moreover, s_{seller} is the only strategy that is EPIR. Seller's payoff from not participating in the mechanism is zero. Following the same reasoning as above on incentive compatibility, s_{seller} is the only strategy that provides a non-negative profit for every buyer and seller type. The same reasoning applies for the buyer, with strategy s_{buyer}: buy if $v_2 > p$ and do not buy if $v_2 \leq p$.

The mechanism can be easily seen to be strictly budget-balanced (SBB). When there is no trade, the monetary transfer is zero. When there is a trade, the positive price paid by the buyer goes to the seller. There is no waste in the numeraire in either case, which ensures the SBB property.

The expected welfare $W(p)$ is the integral of welfare function weighted by the joint type probability distribution function. The bounds determine the entire valuation range over which the trade happens and therefore captures the cases when there is positive surplus from a trade. When a trade does not happen, the surplus is zero. Since agents' VOT distributions are assumed to be independent, the VOT joint distribution ϕ is the product of the two marginal distributions. $\square$

The problem the designer faces is to set an optimal price p^* which maximizes the probability of trade, and thus maximizes the total surplus. In the present study, the optimal price p^* is found by maximizing $W(p)$ with linear search in price intervals of 5 cents.

4.3 Pricing

There are many factors that should be taken into account in determining the fare for ridesharing services. Setting the right fare is essential to the success of a ridesharing system, and deserves design of a separate mechanism which ensures that no incentive exists for drivers and riders to falsely report their preferences in order to affect the amount of transaction.

The fare a rider is charged in our system is made of two components. The first component is a variable, distance/time dependent fee. Assume that rider r's itinerary involves traveling on link set L^*, and that at the time of matching the rider, on each link $\ell \in L^*$, n_ℓ number of individuals (including the driver) share the same vehicle with the rider. We assume that the cost of travel on each link is equally shared by the individuals who travel on the link. Therefore, the variable fee of rider r will be $\sum_{\ell \in L^*} \frac{d_\ell}{n_\ell}$. In this equation, d_ℓ is the general cost of traveling on link ℓ, and $\sum_{\ell \in L^*} \frac{d_\ell}{n_\ell}$ is the total share of the rider from the cost. Note that the general cost of a link can be time-based, distance-based or a combination of both.

In addition to a variable component, the fare also has a fixed component. Since drivers may have to divert from their shortest/preferred paths in order to accommodate riders, they need to be compensated for the extra travel. We calculate the base fare based on the average extra travel time drivers have to spend in the network, assuming an average speed of 40 mph, and a payment of 60 cents per mile. These fares could vary for different times the day, and days of the week, based on the composition of the ridesharing system, i.e., number of participants, the driver to rider ratio, and the degree of flexibility of riders and drivers. Although a pricing scheme that can distribute fares among drivers based on their contribution to the system may be fairer, in the interest of simplicity we use the more preliminary pricing scheme introduced in this section.

5 Numerical Study

In order to study the impact of the P2P ride exchange mechanism on the performance of a ridesharing system, we generate and solve multiple random instances of the ridesharing problem. All results reported here are averaged over 30 runs for each problem instance.

In each problem instance, we generate a number of participants with varying ratio of riders. The origins and destinations of participants are selected based on a uniform random distribution from a pre-specified set of stations in a grid network. The earliest departure time of each participant is selected uniformly by a random distribution within a certain "departure period". The travel time budget of each participant is determined as a factor of his/her shortest path travel time. The latest arrival time of a participants at his/her destination station is then computed as the sum of the participant's earliest departure time and travel time budget. All these parameters impact the level of spatiotemporal proximity between trips.

For each participant, a VOT is drawn from the lognormal distribution described earlier. Each individual is assumed to have a separate transportation alternative outside of the system, with a travel time equal to the shortest path travel time between the individual's origin and destination. The unit distance-based cost of the outside alternative is assumed to be equal to that of the ridesharing system.

We solve each problem instance using three different ridesharing implementation strategies. The first strategy referred to as "one-to-one" matches a single rider with a single driver. The second strategy referred to as "one-to-many" allows a driver to carry multiple riders. Riders complete their trips in a single vehicle. The last implementation strategy referred to as "many-to-many" allows each driver to carry multiple riders, and each rider to transfer between multiple drivers. Note that the P2P exchange mechanism is optimal only for the one-to-one matching method. The number of additional riders served due to exchange in one-to-many and many-to-many systems is only a lower bound and may increase if we use a more sophisticated mechanism that can include higher levels of trade.

In this section, we perform sensitivity analysis over system parameters, namely the number of participants, ratio of riders, travel time budget factor, and number of stations. Through this analysis, we study the impact of different parameter values on the percentage increase in the number of matched riders (to which we refer as the *exchange rate*). Finally, we generate different ridesharing scenarios with different levels of spatiotemporal proximity between trips, and use statistical tests to confirm whether the observed difference in the exchange rates in these scenarios is statistically significant.

5.1 Base Fares

In order for a rider to decide whether to participate in a trade or not, he/she should have information on the cost of the proposed itinerary. As discussed in section 4.3, the cost of an itinerary entails a variable, route-dependent cost, and a fixed cost. In this section, we demonstrate for certain parameter values how this fixed cost is calculated, and how it may differ from hour to hour or day to day.

Figure 4 displays the base fares charged to riders and payed to drivers in a ridesharing system at different levels of trip spatiotemporal proximity. As mentioned in the previous section, this study uses the same base fares for all drivers and a different but equal base fare for all riders in a given time period, for example, during weekday morning peak hours. These fares may vary from location to location, and depend on the number of participants and system composition (number of participants, and ratio of riders to total number of participants). In this section, we show sample base fares for a system with 200 participants with departure period of 60 minutes and travel time budget factor of 1.5, under different ratios of riders and number of stations.

As figure 4 suggests, for a one-to-one system, the fare paid by riders is similar to the fare received by drivers, since the number of served riders and matched drivers in a one-to-one system are equal. In a one-to-many system, in which each driver carries multiple riders, the fixed fare paid by riders decreases as the ratio of served riders to matched participants increases, since now multiple riders are being served by a single driver and hence they each pay a portion of the fixed fare the driver receives. Interestingly enough, in a many-to-many system, when the ratio of riders in the problem is small, each rider pays more than what each driver receives, suggesting that riders are receiving multi-hop itineraries (i.e., transferring between drivers), especially when the number of stations is high and therefore the spatial proximity between trips is too low. After a certain point, however, the fare received by each driver surpasses the fare paid by each rider, suggesting a high level of sharing.

5.2 Number of Participants

In this section, we study the performance of the P2P ride exchange mechanism under different participation rates, and different numbers of stations. For a given number of participants, number of riders and drivers are assumed to be equal, since a rider to driver ratio of close to 1 is where a ridesharing system yields the highest matching rate (Masoud and Jayakrishnan, 2017). Figure 5 shows the initial matching rate of riders, and the exchange rate under different implementations of the system.

The results suggest that under all implementation strategies, both the initial matching rate and the exchange rate are positively correlated with the participation rate, i.e., the higher the number of participants, the higher the performance of the system and the exchange rate.

Another general observation is that with all matching methods, lower numbers of stations result in higher initial matching rates, and higher exchange rates. This result is intuitive, since participants have to choose their origins and destinations from the set of stations. Therefore, a lower number of stations results in more participants sharing the same origin and/or destination stations, and therefore a higher spatial proximity between trips. Higher spatial proximity also suggests a higher probability of finding a driver that can serve the seller in the ride exchange, and therefore higher success rate in the exchange.

Despite this general trend, the impact of spatial proximity depends on the number of participants as well. For a one-to-one system with 300 participants, the exchange rate increases as we increase the spatial proximity by moving from from 49 to 25, and finally 9 stations. With 500 participants, however, it takes lower levels of spatial proximity than 49 stations to impact the exchange rate since a higher number of participants by itself increases the proximity of trips, to some degree.

The one-to-one matching method is the only method for which the result of the exchange mechanism is optimal; therefore, the matching rate for such a system is the highest. Another reason why the P2P exchange scheme may serve one-to-one ridesharing systems well is that in such systems each driver serves a single rider (if matched), which increases the likelihood of liberating the desired driver, and requires negotiations only with a single seller.

In one-to-many and many-to-many systems there are trades that expand to higher levels than the first level of trade, and therefore are not successfully completed. This is one reason behind the lower exchange rates for these systems, compared to the one-to-one system. Another reason for lower exchange rates in one-to-many and many-to-many systems is that the exchange rate measures the percentage increase in the matching rate. Even if the same additional number of riders are served due to exchanges in all implementation strategies, the many-to-many and one-to-many systems will show lower exchange rates because their initial matching rates are higher.

Figure 6 shows the driver matching rates and the impact of the exchange mechanism on the increase in the percentage of matched drivers. In general, this figure follows similar trends to Figure 5. In a many-to-many system the driver exchange rate is higher than the rider exchange rate, which implies a high level of sharing before the exchange.

5.3 System Composition

In this section, we study the impact of system composition on the performance of the exchange mechanism. Figure 7 summarizes the results. For all ridesharing implementations, as the ratio of

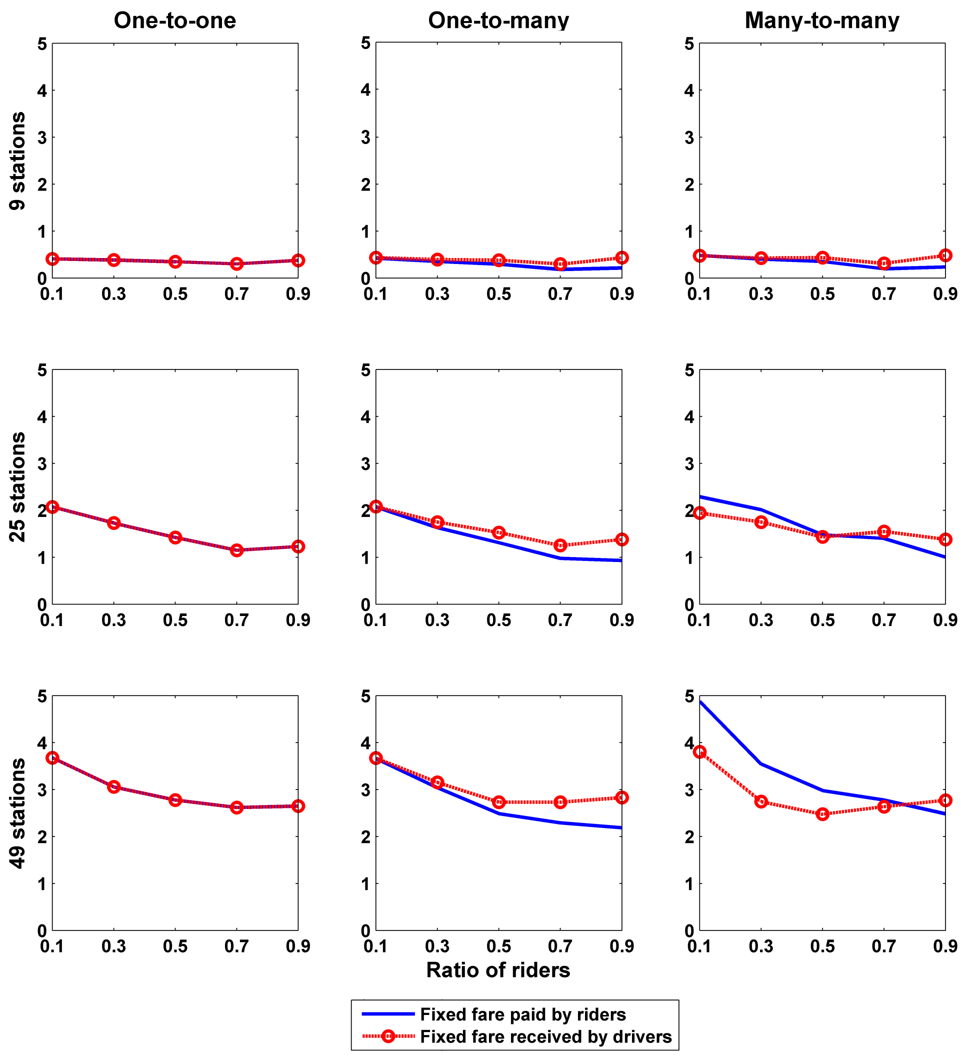

Figure 4: Base fares for ridesharing systems with different levels of spatiotemporal proximity between trips.

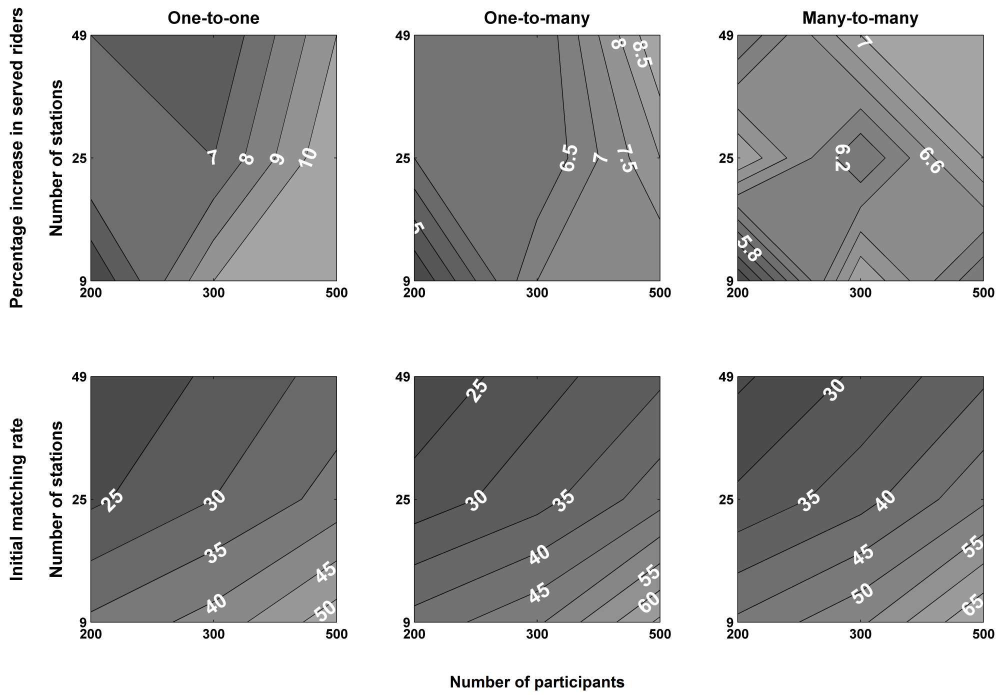

Figure 5: Initial rider matching and exchange rates under different number of stations and participants. Simulations are conducted using departure period of one hour.

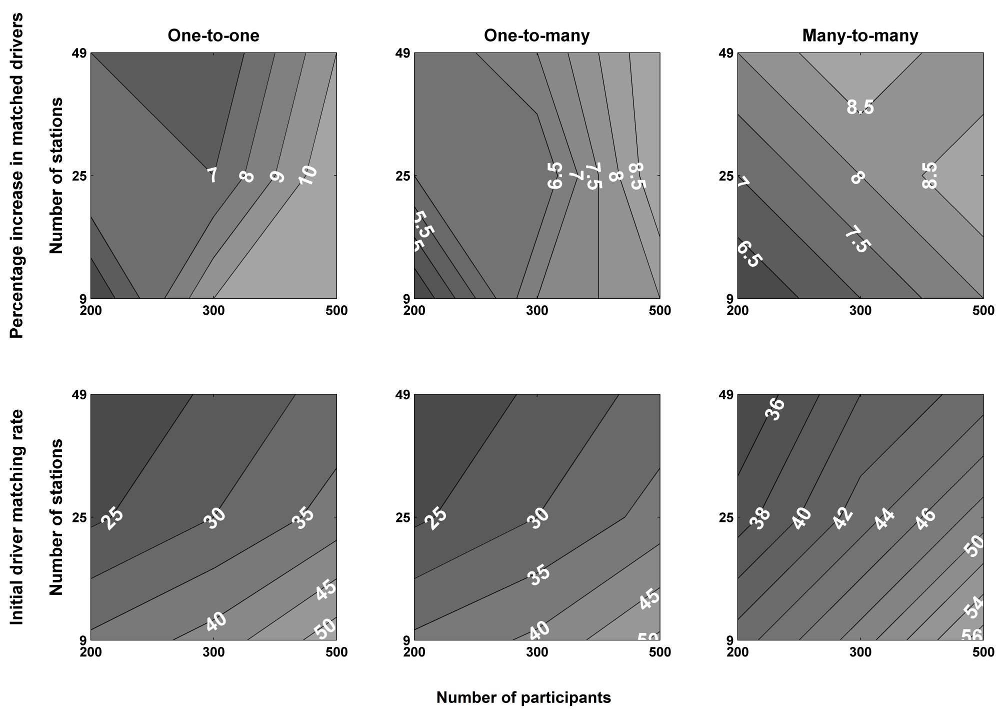

Figure 6: Initial driver matching and exchange rates under different number of stations and participants. Simulations are conducted using departure period of one hour.

riders to participants increases (and ratio of drivers to riders decreases), the rider matching rate experiences a declining trend. At rider ratio of 0.1 (50 riders and 450 drivers), most of the 50 riders can be matched, due to the abundance of supply (i.e., drivers). As we move along the horizontal axes, the increase in demand is met with a decrease in supply, which leads to a decreasing trend in the rider matching rate. At rider ratio of 0.9 (450 riders and 50 drivers) the matching rate of riders is the lowest, due to the lack of sufficient supply.

The trend in the driver matching rate is the opposite of the trend in the rider matching rate. At lower rider ratios where there are a few riders and many more drivers, the driver matching rate is low, since a low percentage of drivers is enough to satisfy the demand. As the rider ratio increases, the demand becomes higher than supply, and so the driver utilization rate increases.

Another general trend among matching rates is the bell-shaped form of the rider and driver exchange rates. In the beginning, when rider to driver ratio is low, a high percentage of demand is satisfied, and there is not much need for exchange, hence the low exchange rate. At high ratios of riders to system participants, the rider matching rate is small, but the driver matching rate is high, and therefore there are not many free drivers left to form alternative itineraries for the sellers. In case of a many-to-many system, an exchange mechanism that can extend to higher levels of trade could help to pick up the declining exchange rate. In the middle ranges, at rider ratio of 0.5 for the one-to-one and one-to-many systems and 0.3 for the many-to-many system, the exchange rate becomes the highest. In this range, the rider matching rate is not too high to eliminate the need for exchange, and the driver matching rate is not too high to decrease the chance of finding alternative itineraries for the sellers.

5.4 Departure Period

In this section, we study the impact of higher temporal proximity among trips on the performance of the exchange mechanism. Figure 8 shows that under all ridesharing implementations, both the initial rider and driver matching and exchange rates increase with the temporal proximity among trips (i.e., as the departure period becomes smaller.) This is not surprising, since higher temporal proximity among trips leads to higher probability of finding a match in the first place, and a higher probability of finding alternative itineraries for any seller, in case an exchange is required.

5.5 Travel Time Budget Factor

Experiments in this section demonstrate the impact of flexibility of participants in defining their travel time windows on the matching and exchange rates. Figure 9 demonstrates the initial rider matching rates and the percentage increase in the number of served riders under different travel time budget factors for participants. This figure suggests that in general the matching rate as well as the exchange rate increase with the travel time budget factor, regardless of the change in the proximity of trips. Figure 10 suggests the same results for drivers.

5.6 Statistical Analysis

Table 1 lists 10 ridesharing scenarios sorted in an increasing order of exchange rate. Scenarios have different levels of spatiotemporal proximity among trips, and are all generated for a one-to-one matching method, since it is the only matching method for which the exchange mechanism is optimal.

As demonstrated in the previous sections, when the spatiotemporal proximity of trips is too low or too high, the exchange rate is small. When the spatiotemporal proximity among trips is too low, the initial matching rate is too low, and so is the exchange rate due to lack of supply. When the

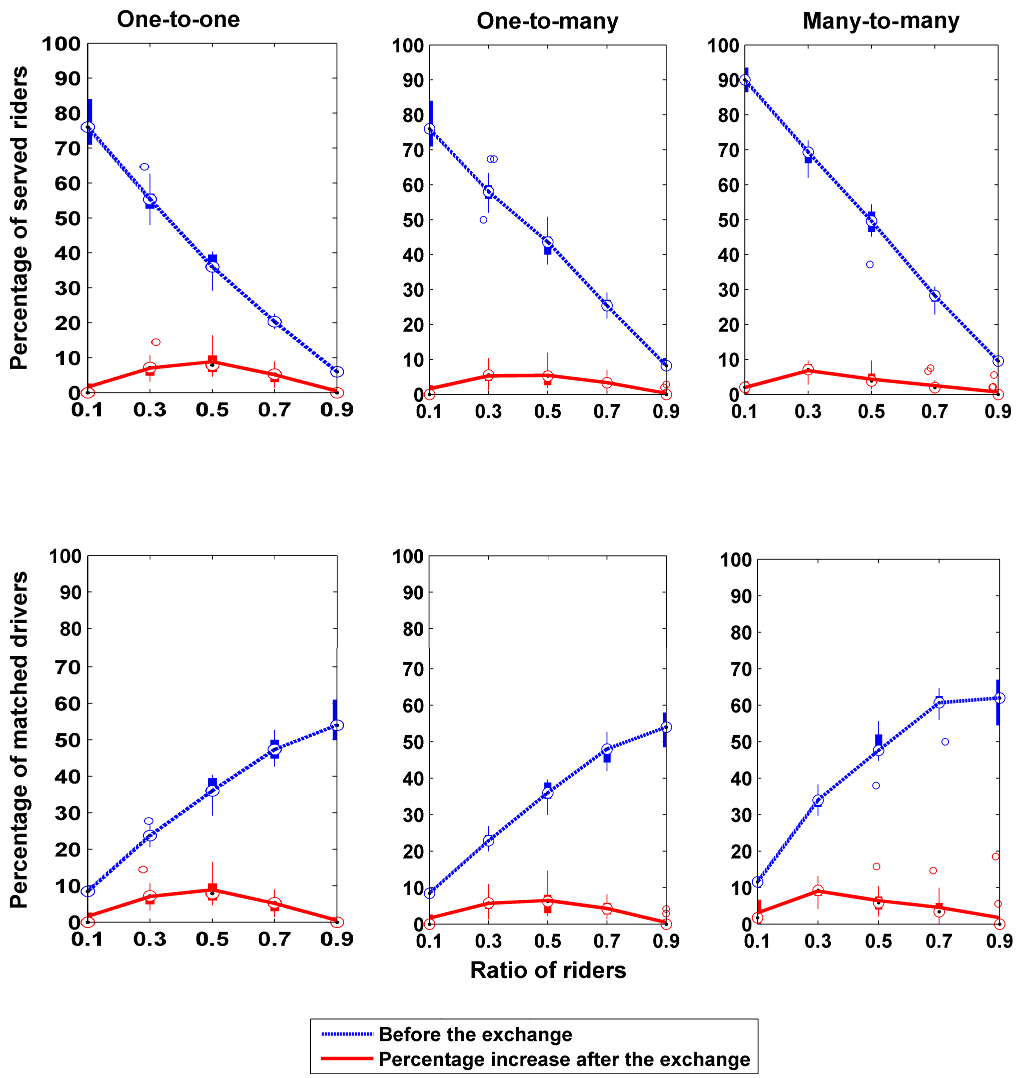

Figure 7: Matching and exchange rates under different ratio of riders. Problem instances include 500 participants, where the number of riders is changed from 50 to 450 in 100 increments. The departure period and number of stations are set as one hour and 25, respectively.

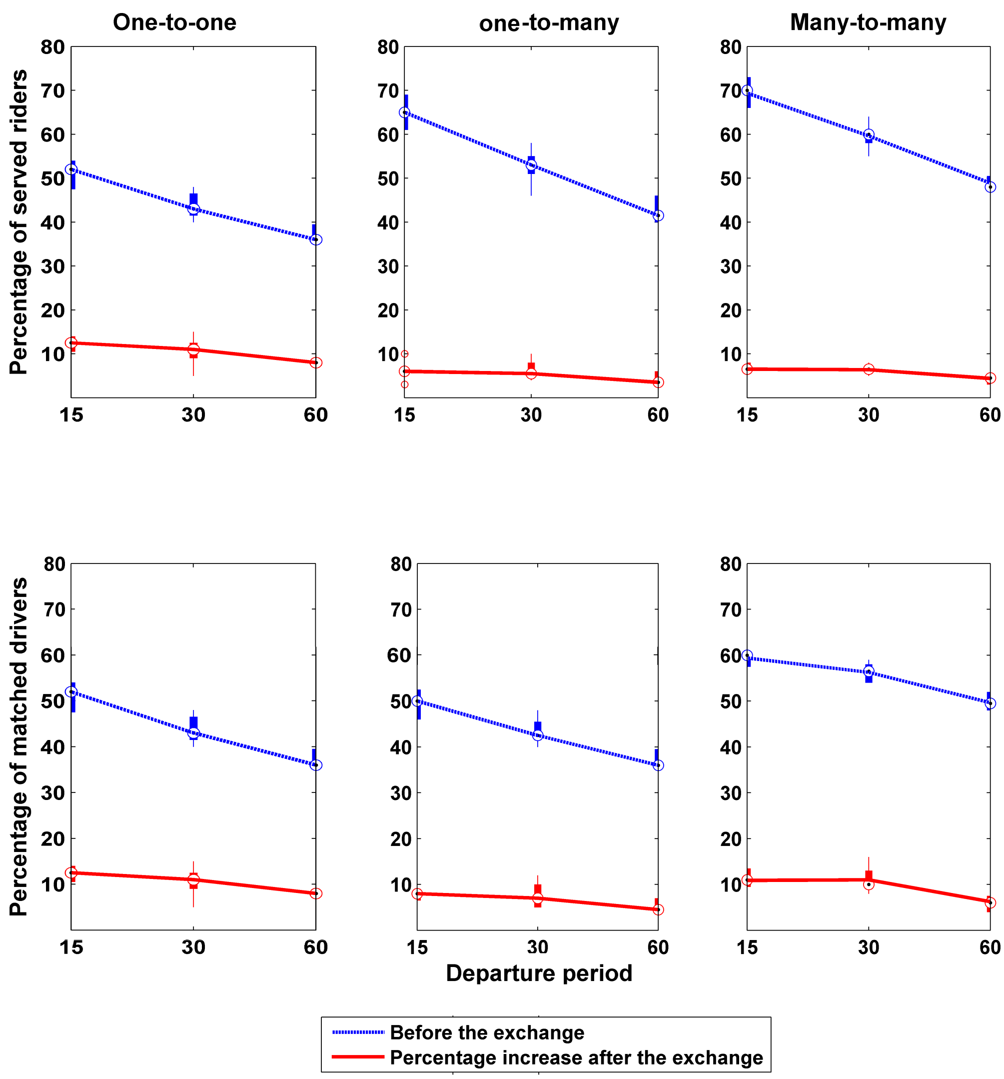

Figure 8: Percentage of riders and drivers matched before and by P2P ride exchange under different departure periods. Problem instances contain 500 participants, with equal number of drivers and riders. The travel time budget factor and number of statons are set as 1.5 and 25, respectively.

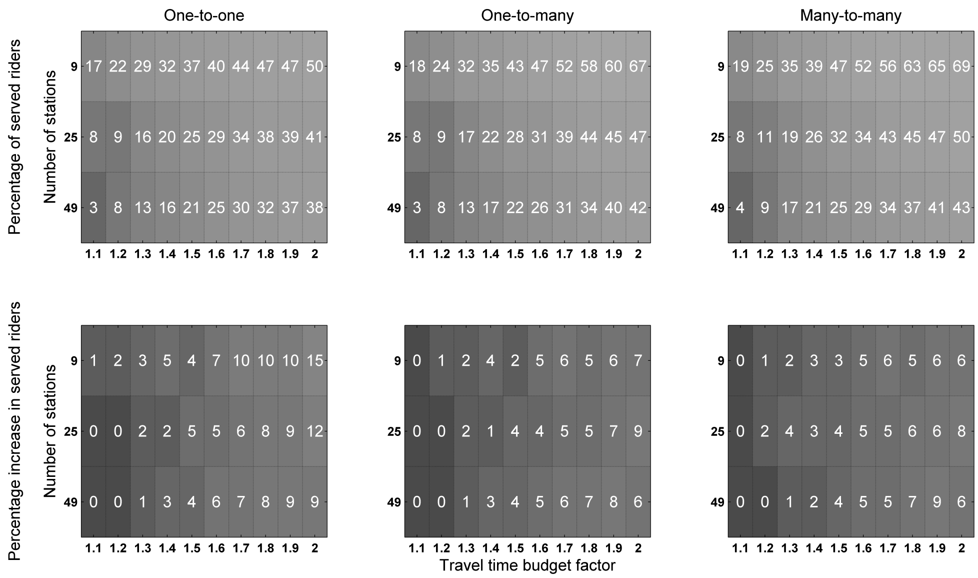

Figure 9: Rider matching and exchange rates as a function of travel time budget factor of participants. Problem instances contain 200 participants, with equal number of riders and drivers. The departure period is set as 60 minutes.

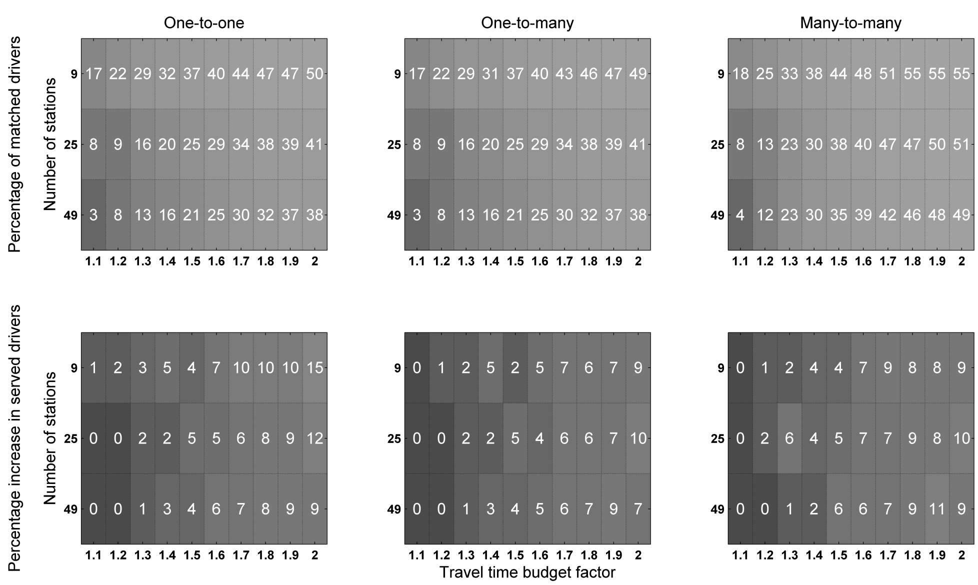

Figure 10: Driver matching an exchange rates as a function of travel time budget factor of participants. Problem instances contain 200 participants, with equal number of riders and drivers. The departure period is set as 60 minutes.

Table 1: Ridesharing instances. The scenario properties include (no. of participants, ratio of riders, departure period, no. of stations, travel time budget factor)

Scenario properties	Per. of served riders (mean,st.dev.)	Per. increase in served riders (mean,st.dev.)	Per. of retained riders	Average social surplus
(500,0.1,15,9,1.5)	(99,1.5)	(0.2,0.4)	0.6	11.33
(200,0.25,15,9,1.5)	(9,0.6)	(4.5,3.2)	11.5	10.92
(500,0.9,15,9,1.5)	(10,0.4)	(7.2,3.6)	18	8.82
(200,0.75,15,9,1.5)	(8,0.7)	(8,4.9)	20	8.66
(200,0.5,15,9,1.5)	(13,1.2)	(10,3.7)	26	9.69
(200,0.6,15,9,1.5)	(12,0.9)	(10,4)	25	9.47
(500,0.7,30,9,1.5)	(25,1.3)	(10,3)	28	9.01
(500,0.7,15,25,1.5)	(21,1.3)	(11,3.2)	28	14.34
(500,0.7,15,9,1.5)	(33,1)	(15,3)	38	9.97
(1000,0.5,30,25,1.5)	(60,1.6)	(15,1.2)	32	15.57

spatiotemporal proximity is too high, a high percentage of demand is served, and there is not much room left for improvement by making exchanges. In the middle ranges, where the spatiotemporal proximity among trips is moderate, is where the exchange rate is the highest. Table 1 lists scenarios that cover the entire range. The mean and standard deviation of the matching and exchange rates for each scenario are presented in this table. For each scenario, we have also noted the total social surplus of the system obtained due to exchange, averaged over the 30 generated instances for each scenario.

Figure 11 demonstrates the statistical significance of the difference in exchange rates among scenarios. For each pair of scenarios, we have used a two-sample t-test with the null hypothesis that the two scenarios come from independent random samples drawn from two normal distributions with equal means and equal but unknown variances. The alternative hypothesis is that the two scenarios come from populations with unequal means. The null hypothesis is rejected under a 5% significance level.

Figure 11 shows the p-values for the two-sample t-tests for each pair of scenarios. This figure suggests that under the 5% significance level each scenario is not statistically different from one or two scenarios with higher exchange rates, but as the difference in the exchange rates among scenarios increases, the null hypothesis is rejected, and the scenarios are shown to be statistically different. This figure suggests that the difference in the exchange rates observed for different ridesharing systems with different parameter values and levels of spatiotemporal proximity among trips is statistically significant.

5.7 Customer Retention

The higher number of served riders due to exchange does not have a one-to-one impact on the performance of the system. It is true that the number of served riders increases only by the number of successful exchanges, but the impact on customer retention and the reputation of the system should also be taken into consideration.

To study customer retention, we assume that a rider does not return to the system if he/she has three failed experiences. By simulating a three day experience for each rider, it is possible to compute the number of retained customers in a 3-day period. To conduct this analysis, we use

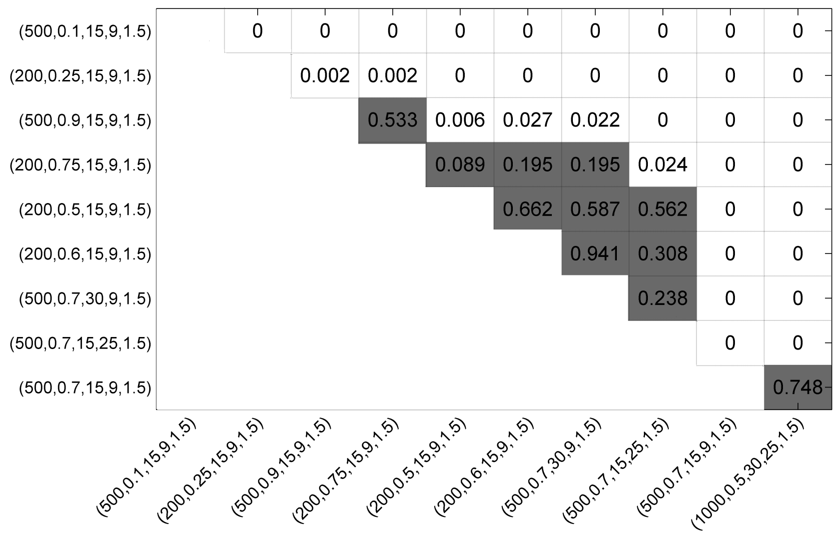

Figure 11: p-values between the exchange rates is scenarios listed in table 1

the same set of riders in each instance of each of the scenarios listed in Table 1 (i.e., we change only the driver set for each problem instance in each scenario). To compute customer retention, it is assumed that a rider will consider using the system again if he/she has at least one successful experience. The percentage of retained riders due to exchange are reported in Table 1.

In addition to creating a positive experience for these riders and increasing the probability of them returning to the system, using P2P exchange could eliminate the possible negative WOM that could have been generated by these riders had they not been served, and could even replace them with a positive WOM.

5.8 Higher Levels of Trade

This study concentrated on the simplest possible scenario for trade, where there is a single buyer and a single seller. We study the impact of higher levels of trade on the exchange rate by considering a many-to-many ridesharing system for scenarios in Table 1. The results show an average of 20% increase in the number of served riders, over all scenarios. Numbers in table 1 are generated solely based on the operational feasibility of exchange (e.g., making sure for a rider who sells his/her current itinerary, there are alternative non-conflicting itineraries available), and not the monitory transactions. In addition, note that although theoretically a multi-level trade can increase the percentage of served riders significantly, moving to lower levels of trade can make the trade harder to manage and less probable to succeed, due to the requirement for all individuals involved in the trade to agree to it.

6 Conclusion

In this paper, we introduced the P2P ride exchange mechanism to increase the performance level of a dynamic P2P ridesharing systems. Extensive numerical experiments suggest that this mechanism results in higher performance levels and customer retention rates than standard FCFS allocation. We demonstrated through experiments that the percentage of increase in the matching rate using

P2P ride exchange is positively correlated with the number of participants, travel time budget of riders and drivers, and spatial and temporal proximity of trips. Furthermore, we showed that the exchange mechanism is more effective in terms of increasing the percentage of served riders when the ratio of number of riders to system participants is within a moderate range (not too low or too high).

We showed that the mechanism employed is robust towards selfish manipulation from users and helps in increasing ridership with minimal information requirements from users. Despite these low information requirements, an increase in the probability of getting a ride is achieved, leading to further user permanence in the system, even in situations where the system operation has not yet started and no data is available. Once more data is available, a new mechanism that takes into consideration observed user interaction may be designed.

We conclude with a comment on the essential paradigm that lies behind the entire concept, namely that the design facilitates the easing of the traditional FCFS rule in the consumption of transportation supply by users. While FCFS is often considered as fair, it leads to economic inefficiency. New devices such as smartphones allow operators to incorporate a wider range of information about user preferences, allowing fairer and more efficient allocation schemes. This paper is an early attempt at analyzing such a paradigm for dynamic P2P ride-sharing systems.

References

Abou-Zeid, M., M. Ben-Akiva, M. Bierlaire, C. Choudhury, and S. Hess (2010). Attitudes and value of time heterogeneity. *Applied Transport Economics-A Management and Policy Perspective. De Boeck Publishing*, 523–545.

Agatz, N., A. Erera, M. Savelsbergh, and X. Wang (2009). Sustainable passenger transportation: Dynamic ride-sharing. *ERIM Report Series Reference No. ERS-2010-010-LIS*.

Agatz, N. A., A. L. Erera, M. W. Savelsbergh, and X. Wang (2011). Dynamic ride-sharing: A simulation study in metro Atlanta. *Transportation Research Part B: Methodological 45*(9), 1450 – 1464.

Alt, H., N. Blum, K. Mehlhorn, and M. Paul (1991). Computing a maximum cardinality matching in a bipartite graph in time o (n 1.5 mlog n). *Information Processing Letters 37*(4), 237–240.

Baldacci, R., V. Maniezzo, and A. Mingozzi (2004). An exact method for the car pooling problem based on lagrangean column generation. *Operations Research 52*(3), pp. 422–439.

Brereton, M. and S. Ghelawat (2010). Designing for participation in local social ridesharing networks: grass roots prototyping of it systems. In *Proceedings of the 11th Biennial Participatory Design Conference*, pp. 199–202. ACM.

Buttle, F. A. (1998). Word of mouth: understanding and managing referral marketing. *Journal of strategic marketing 6*(3), 241–254.

Cervero, R. and B. Griesenbeck (1988). Factors influencing commuting choices in suburban labor markets: a case analysis of pleasanton, california. *Transportation Research Part A: General 22*(3), 151–161.

Chan, N. D. and S. A. Shaheen (2012). Ridesharing in north america: Past, present, and future. *Transport Reviews 32*(1), 93–112.

Di Febbraro, A., E. Gattorna, and N. Sacco (2013). Optimizing dynamic ride-sharing systems. *Presented in the TRB 2013 annual meeting*.

Furuhata, M., K. Daniel, S. Koenig, F. Ordonez, M. Dessouky, M.-E. Brunet, L. Cohen, and X. Wang (2015). Online cost-sharing mechanism design for demand-responsive transport systems. *Intelligent Transportation Systems, IEEE Transactions on 16*(2), 692–707.

Furuhata, M., M. Dessouky, F. Ordóñez, M.-E. Brunet, X. Wang, and S. Koenig (2013). Ridesharing: The state-of-the-art and future directions. *Transportation Research Part B: Methodological 57*, 28–46.

Gaynor, T. (2015). Solving the critical mass problem. what is the simplest thing that could work? Accessed 7/29/2015.

Ghoseiri, K. (2013). Dynamic rideshare optimized matching problem. *Ph.D. Dissertation at University of Maryland*.

Gustafsson, A., M. D. Johnson, and I. Roos (2005). The effects of customer satisfaction, relationship commitment dimensions, and triggers on customer retention. *Journal of marketing 69*(4), 210–218.

Hagerty, K. M. and W. P. Rogerson (1987). Robust trading mechanisms. *Journal of Economic Theory 42*(1), 94–107.

Heinrich, S. (2010). Implementing real-time ridesharing in the san francisco bay area. *Ph.D. Dissertation at San Jose State University*.

Herbawi, W. and M. Weber. (2011a). Ant colony vs. genetic multiobjective route planning in dynamic multi-hop ridesharing. *in Tools with Artificial Intelligence (ICTAI), 2011 IEEE Conference*.

Herbawi, W. and M. Weber. (2011b). Comparison of multiobjective evolutionary algorithm for solving the multiobjective route planning in dynamic multi-hop ridesharing. *Proceedings of the 11th European conference on Evolutionary computation in combinatorial optimization, Torino, Italy*, 84–95.

Herbawi, W. and M. Weber. (2012). A genetic and insertion heuristic algorithm for solving the dynamic ridematching problem with time windows. *In Proceedings of the fourteenth international conference on Genetic and evolutionary computation conference*.

Kelly, K. L. (2007). Casual carpooling-enhanced. *Journal of Public Transportation 10*(4), 6.

Kleiner, A., B. Nebel, and V. Ziparo (2011). A mechanism for dynamic ride sharing based on parallel auctions. *In Proceedings of the 22th International Joint Conference on Artificial Intelligence (IJCAI)*, 266–272.

Mas-Colell, A., M. D. Whinston, J. R. Green, et al. (1995). *Microeconomic theory*, Volume 1. Oxford University press New York.

Masoud, N. and R. Jayakrishnan (2015). A real-time algorithm to solve the peer-to-peer ride-matching problem in a flexible ridesharing system. Technical report.

Masoud, N. and R. Jayakrishnan (2017). A decomposition algorithm to solve the multi-hop peer-to-peer ride-matching problem. *Transportation Research Part B: Methodological 99*, 1–29.

Morency, C. (2007). The ambivalence of ridesharing. *Transportation 34*(2), 239–253.

Myerson, R. B. and M. A. Satterthwaite (1983). Efficient mechanisms for bilateral trading. *Journal of economic theory 29*(2), 265–281.

Nisan, N. and A. Ronen (1999). Algorithmic mechanism design. In *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, pp. 129–140. ACM.

Pfeifer, P. E. (2005). The optimal ratio of acquisition and retention costs. *Journal of Targeting, Measurement and Analysis for Marketing 13*(2), 179–188.

Ranaweera, C. and J. Prabhu (2003). On the relative importance of customer satisfaction and trust as determinants of customer retention and positive word of mouth. *Journal of Targeting, Measurement and Analysis for marketing 12*(1), 82–90.

Raney, S. (2010). San francisco to silicon valley, california, instant ridesharing with transfer hub. *Transportation Research Record: Journal of the Transportation Research Board* (2143), 134–141.

Regue, R., N. Masoud, and W. Recker (2016). Car2work: A shared mobility concept to connect commuters with workplaces. *Transportation Research Record: Journal of the Transportation Research Board 2542*, 102–110.

Rust, R. T. and A. J. Zahorik (1993). Customer satisfaction, customer retention, and market share. *Journal of retailing 69*(2), 193–215.

Söderlund, M. (1998). Customer satisfaction and its consequences on customer behaviour revisited: The impact of different levels of satisfaction on word-of-mouth, feedback to the supplier and loyalty. *International Journal of Service Industry Management 9*(2), 169–188.

Teodorović, D. and M. Dell'Orco (2005). Bee colony optimization–a cooperative learning approach to complex transportation problems. In *Advanced OR and AI Methods in Transportation: Proceedings of 16th Mini–EURO Conference and 10th Meeting of EWGT (13-16 September 2005).– Poznan: Publishing House of the Polish Operational and System Research*, pp. 51–60.

Wang, X. (2013). Optimizing ride matches for dynamic ride-sharing systems.

Wolfler Calvo, R., F. de Luigi, P. Haastrup, and V. Maniezzo (2004). A distributed geographic information system for the daily car pooling problem. *Computers & Operations Research 31*(13), 2263 – 2278.