THEMA Working Paper n°2020-12
CY Cergy Paris Université, France

Influence of Dynamic Congestion on Carpooling Matching

André de Palma, Patrick Stokkink, Nikolas Geroliminis

December 2020

Influence of Dynamic Congestion on Carpooling Matching

André de Palma[1], Patrick Stokkink[2], Nikolas Geroliminis[2*]
[1] CY Cergy-Paris University, Thema, France.
[2] École Polytechnique Fédérale de Lausanne (EPFL), Urban Transport Systems Laboratory (LUTS), Switzerland
[*] Corresponding author. E-mail address: nikolas.geroliminis@epfl.ch

Date: December 4, 2020

This is a draft version. Please do not quote or circulate without permission of one of the authors

---

Abstract

Carpooling is an efficient measure to fight car ownership and reduce vehicle kilometers travelled. By sharing their commutes, vehicle occupancy increases and congestion is reduced. We develop a dynamic ADL (Arnott, de Palma, Lindsey)-Vickrey approach for a corridor monocentric city à la Hotelling. First, we formulate the matching problem of heterogeneous users in carpooling as an MILP problem and we discuss its analytical properties. Next, we construct a bi-level optimization problem involving matching (first stage) and dynamic traffic congestion (second stage). We provide a heuristic to attain an optimal matching for a dynamic traffic equilibrium with congestion. Such a template allows studying the two-way causality between dynamic congestion and carpooling matching.

Keywords: Carpooling, Ride-sharing, Matching, Scheduling delay, Bottleneck congestion

JEL classification: C78, R40, R41

---

1 Introduction

Cars provide transportation between two points for a human activity. But, cars are parked 95% of the time, creating a huge waste of space; and when cars move, they mostly carry only one passenger. Collective transport has been realized for long as a proper solution to fight congestion. While collectivity has been successful with respect to subways (e.g., London Tube carries 5M passengers per day), it is not the norm in road traffic. Even if cars are more expensive than other alternatives like shared mobility or public transport, car ownership continues to increase (one car for every two people in EU or US according to Eurostat (2020)). Thus, collective travel obviously is not as generalized as it should be. The obstacles towards its general adoption are related to accessibility and increased travel times of other modes.

Over the last years, demand for mobility services such as vehicle-sharing and carpooling (or ride-sharing) has increased tremendously. One of the most popular examples of carpooling is BlaBlaCar (2020), with 100 million travellers per year in 22 countries. Carpooling initiatives are known to have many advantages such as increasing mobility and reducing congestion. Carpooling eases congestion by increasing average vehicle occupancy. According to BlaBlaCar (2020), they increased the average occupancy of their vehicles to 2.8 compared to an average of 1.6 in Europe. However, carpooling decisions are also influenced by congestion. Congestion brings forth significant economic costs both through extra fuel consumption and waste of time. Furthermore, it induces social costs as it is commonly associated with public health risks. These effects emphasize the need to introduce congestion in carpooling models (literature has remained silent on this issue) to better understand its potential benefits and costs.

We consider the commute of passengers and drivers in a parsimonious framework that will allow us to develop intuition for a challenging problem both in terms of mathematical formulation and economic insights. In our problem, passengers and drivers are located along a horizontal line and all have the same destination at the end of the horizontal line, the Central Business District (CBD). This framework forms a theoretical approximation of many real-life situations, where people live in the suburbs and work in the business district. Suburbs and the CBD are connected by a highway. Contrarily to a formally similar problem, the matching in marriage problem (Chiappori et al., 2012), in this paper a central operator matches the drivers and the passengers such that the total matching costs are minimized in the presence of untolled congestion. We refer to this as the *optimal matching problem*. It should be clear that we do not consider the market solution nor the relation between the equilibrium and the optimal solution, which has been extensively studied in applied mathematics and in economics (Galichon, 2018). The market solution is more relevant in, for example, labour economics or in the marriage market, where agents are free to choose their match. In urban economics, de Palma et al. (2007) have proposed an algorithm to match supply and demand in the housing market when demand exceeds supply. The matching algorithms considered here provide another allocation mechanism to address allocation between supply and demand. In our transport setting, the matching costs include the detour time to pick up a passenger but may also include scheduling delay (associated with early or late arrivals). Congestion can significantly influence travel time and tardiness of all agents and may therefore influence matching decisions. In this paper, we incorporate dynamic congestion in matching decisions as well as departure time decisions. In the same unified framework, drivers still determine their time of departure considering scheduling and congestion costs as per W. Vickrey (Vickrey, 1969).

The key element of carpooling is the matching process between drivers and passengers. This is a well-known mathematical problem (the optimal transport model reviewed by Galichon (2018)) that we will apply in this paper. In the carpooling context, since the quality of matching depends on travel times, which itself depends on how many matches have been performed, we have an extended version of the standard matching model, that we refer to as the *congested optimal matching model*. Other applications of this congested matching model can be studied, such as the housing market (in that case, matching will potentially change the market prices and the population mix which may feed-back to the matching process) or the taxi market (e.g. the operations research model of stable matching of Bai et al. (2014)), but this is clearly outside the scope of this paper.

A two-way causality exists between carpooling and congestion. One side of the implication is well understood and intuitive. If a fraction of the drivers starts to carpool, the number of cars on the road decreases, which eases congestion (Bahat and Bekhor, 2016; Li et al., 2016). The other side of the causality is more subtle and unexplored in the standard matching models. In such standard matching models, used in family economics and in labour market (Browning et al., 2014), the matching costs are exogenous in the sense that the match has no influence on the utility of other individuals, and in particular on those who do not decide to match. This is not the case here since matches do influence travel time, which are experienced by all drivers (solo drivers or not) and by all other carpoolers. Note that matching models in these two areas are not the only potential methodological approaches. Search models provide an alternative manner to describe how an individual can find his/her ideal spouse or how a firm can find an employee (Chiappori et al., 2012). This corresponds to the famous “secretary problem” in operations research; its solution is the optimal stopping rule (Fushimi, 1981). As noted by the Diamond paradox, such a model can lead to non-equilibrium when the prices (which would be similar here to congestion) are determined endogenously (Bagwell et al., 1992). The spatial search, not envisaged here but which provides an alternative approach for carpooling, has been studied in urban economics by, for example, Stoll (1999). However the search model, which could be potentially used for carpooling, seems less realistic for commuting than the matching model, which is explored here.

Here we consider an equilibrium-type of process in the sense that the congested matching cost depend on the level of congestion. This level of congestion depends on departure time decisions; drivers and passengers have specific preference over the optimal time to travel (i.e. about their optimal arrival time). Coordination over time (but not over space) was explored in a family context by de Palma et al. (2015), who study the joint morning departure time decision of couples. They showed that coordination of departure time reduces in aggregate the spread of the peak, thus leading to more congestion. The reverse causality also works in the sense that any policy (such as tolling) which spreads the peak, influences the intra-household decisions. Note that such processes are totally omitted by the current cost-benefit analysis. In this paper we envisage such issue, but not from the point of view of spouses within the same household but for a whole population of drivers and passengers. We are then able to analyze the impact of dynamic congestion on the matching costs and therefore on the quality of matching.

We model the two-way causality between carpooling and dynamic congestion with scheduling preferences using a bi-level optimization approach. We formulate the optimal matching problem (first stage) as an integer linear programming problem. Using this formulation, we can determine the matching that minimizes the sum of the cost of the detour as well as of the potential inconvenience costs while all passengers are matched to a driver. Inconvenience cost can be a psychological cost, and/or more specifically some extra schedule delay cost due to the need for coordination of the drivers' and passengers' departure times. Congestion is incorporated through a dynamic bottleneck model (for the second stage). Using an iterative approach, we obtain the optimal matching for a dynamic traffic equilibrium with congestion. By comparing multiple scenarios and performing extensive sensitivity analysis, we evaluate the effect of congestion on carpooling matching. We consider a fixed subsidy that is assumed to be sufficiently high to stimulate carpooling. It is possible for the operator to subsidise some drivers more than others to stimulate more matches, which is referred to as *targeted subsidies*.

The remainder of this paper is organized as follows. Section 2 provides a review of the relevant literature. The framework and corresponding theoretical results are described in Section 3. The framework is extended with bottleneck congestion in Section 4. Section 5 discusses the simulation results and in Section 6 the paper is concluded. Most formal proofs of the theorems are relegated to the Appendix.

2 Literature Review

Over the last years, a shift from privately-owned vehicles to shared mobility services has been observed. This is commonly reffered to as Mobility as a Service (MaaS). The main concept of MaaS is to offer on-demand transportation to satisfy the travel needs of customers. A review on MaaS is provided by Jittrapirom et al. (2017), who classify various MaaS frameworks based on their unique characteristics. In vehicle-sharing, individuals can rent a vehicle for a short period of time (Shaheen et al., 2012). This includes car-sharing (e.g. Boyacı et al. (2015), Bardhi and Eckhardt (2012)), bike-sharing (e.g. DeMaio (2009), Chemla et al. (2013)) and electric vehicle-sharing (He et al., 2017). In ride-sharing on the other hand, people share a single journey which prevents drivers from driving to their destination with (partially) empty vehicles.

Ride-sharing systems aim to connect various travellers with similar origins, destinations and time schedules. According to Sullivan and O'Fallon (2010), average vehicle occupancy in urban areas was approximately equal to 1.53 and the driver was driving alone for around 66% of the considered trips. This number is even lower for commuting trips to and from work and has increased even further since then. Ride-hailing companies such as Uber and Lyft also aim to connect drivers and passengers, but recent studies have shown that they can significantly influence congestion by having a large number of vehicles without passengers circulating in the network (Beojone and Geroliminis, 2020). Although some of their activities can be seen as ride-sharing, many drivers of ride-hailing companies act as taxi drivers, which is therefore not classified as ride-sharing. In this paper, we will use the terms carpooling and ride-sharing interchangeably.

For a review of the optimization challenges in ride-sharing, the reader is referred to Agatz et al. (2012). One of the most important optimization problems in ride-sharing is the matching of drivers and passengers. Matching algorithms can aim to find system optimal matching, as described by Özkan and Ward (2020), or a stable matching where no individual can improve their match as described by Wang et al. (2018). Various extensions to the traditional ride-sharing framework have been proposed. For example, Masoud and Jayakrishnan (2017) introduce a framework where passengers can transfer between multiple drivers. Matching is also considered in ride-sharing scenarios where more than two people share a ride. For example, Santi et al. (2014) consider the sharing of taxi services to reduce its negative effect at the cost of increased inconvenience perceived by the passengers. Alonso-Mora et al. (2017) consider high-capacity ride-sharing with dynamic trip-vehicle assignment and also display the trade-off between passenger inconvenience and negative externalities of commuting. The standard matching models are very much used in labor economics (Zenou, 2009) and in the economics of the family (Browning et al., 2014). In urban economics, matching is also used to justify the micro-foundations of agglomeration effects or wide economic effects (Duranton and Puga, 2004).

In this paper we consider the two-way causality between carpooling and congestion. The effect of carpooling on congestion is intuitive and has been well-studied. Among others Caulfield (2009), Li et al. (2016) and Gurumurthy et al. (2019) describe that carpooling can significantly decrease congestion in urban areas. Bahat and Bekhor (2016) incorporate ride-sharing in the static traffic assignment model. They examine the potential benefit of ride-sharing on congestion and emissions in road networks. Their results show that due to the introduction of ride-sharing, congestion as well as average ride time decreases.

To the best of our knowledge, the effect of dynamic congestion with scheduling preferences on carpooling has not been studied before. Dynamic traffic models (de Palma et al., 1987; Arnott et al., 1990) are commonly used to model the relation between congestion and traffic decisions. Smith (1993) presents a dynamic traffic flow model for peak period traffic flows in urban areas, where road capacity is tight. Morning and evening commutes are commonly modelled using a traffic bottleneck model (Arnott et al., 1993) based on the Vickrey (1969) model. Later, the corridor model was extended by Qian and Zhang (2011) to include multi-modal transportation, with among others carpooling. For a recent review of the bottleneck congestion model, the reader is referred to Li et al. (2020). The effect of fuel prices on carpooling behaviour has been studied by Bento et al. (2013). Their results indicate that flow on high occupancy vehicle lanes on average increases with fuel prices. However, they suggest that local traffic congestion such as bottlenecks may change the way these variables interact.

In this paper, we use a bi-level optimization approach to model the two-way causality between carpooling and congestion. The first level considers the optimal matching problem. We determine the matching that minimizes the total costs, which comprises detour, delay and inconvenience costs. Congestion is incorporated in the second level through a dynamic bottleneck model. Equilibrium departure times are determined given the optimal matching. Using an iterative process, we can solve the congested matching model to evaluate the effect of congestion on carpooling matching.

3 Matching without Congestion

In this paper, we consider the morning commute of passengers and drivers. In this case, we assume that the origins of all agents are distributed on Hotelling's line [0,1), while their destination is at the central business district (CBD), which is located at 1 and is therefore the same for all agents. An example of this with one customer and one driver is depicted in Figure 1. During the evening commute, this problem is reversed. In this case, the origin is the same for all agents, while their destinations are distributed on Hotelling's line. It is important to note that the morning and evening commute are similar but not identical, since in the evening there is an optimal desired departure time from the origin, rather than a desired arrival time at the destination.

Let the location of passenger j be denoted as y_j and the location of driver i as x_i. Let $\mathcal{C}(i,j)$ denote the (full) cost of a match between driver i and passenger j. A driver always picks up a passenger before driving to the destination. This means that if the location of the passenger is before that of the driver (i.e. driver is closest to the CBD), the matching distance is twice the distance between the customer and the driver. On the contrary, if the passenger is located closest to the CBD, there is no additional matching distance. We let α denote the value of time[1] of drivers. The cost of matching for a driver located at x_i and a passenger located at y_j is defined as follows:

$$\mathcal{C}(i,j) = 2\alpha \max(x_i - y_j, 0) \tag{1}$$

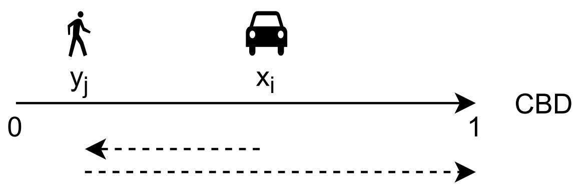

Figure 1: Example of a link with one passenger located further away from the CBD than the driver

In our framework, the set of drivers and passengers is assumed to be fixed and known beforehand. Thereby, we assume that a central operator determines the matching and this can be enforced to the participants. This assumption is reasonable as drivers and passengers are generally unaware of their alternatives in such a carpooling framework and they can inform the operator of their trip information (origin and desired arrival at destination) in advance. We start by assuming the matching costs are limited to the costs for making a detour (A1). This assumption is reasonable in case all individuals have identical desired arrival times, and this assumption will be relaxed in later sections.

Assumption A1 (Simple matching): *The cost of matching driver i to passenger j depends only on the detour driver i makes to pickup passenger j, as defined in Equation* (1).

[1]Different value of times can be considered easily without changing the mathematical nature of the problem. For simplicity, we assume here a single value of time.

The remainder of this section is organized as follows. First we discuss the basic model in Section 3.1. Then, we extend the model to incorporate unequal numbers of drivers and passengers in Section 3.2. We incorporate inconvenience and scheduling delay in Sections 3.3 and 3.4 respectively. Lastly, we consider the potential of targeted subsidies in Section 3.5.

3.1 Basic Carpooling Model: Equal Number of Drivers and Passengers

First we assume that the number of passengers and drivers are equal. We can formulate this as a linear assignment problem. Let I denote the set of drivers and J the set of passengers. To ease notation, the following assumption is made without loss of generality:

Assumption A2 (Sorting): *All passengers and drivers are sorted from left to right based on their location on the Hotelling line.*

Following this assumption, driver locations satisfy $x_i \leq x_{i+1}$ and passenger locations satisfy $y_j \leq y_{j+1}$. We study in this paper the matching process, which minimizes the total costs (which includes the potential schedule delay costs). This is thus an optimal assignment provided by a regulator under the following definition.

Definition 1 (Optimal matching without congestion): *The optimal matching minimizes the sum of detour cost and inconvenience cost (including scheduling delay costs) for all matched passengers and drivers.*

We use decision variables a_{ij}, which are equal to 1 if passenger j is matched with driver i. We then formulate this problem as a linear assignment problem (see Burkard and Cela (1999), Galichon (2018)). We determine the optimal matching, that is the matching minimizing total cost, given that all drivers are matched to a passenger and all passengers are matched to a driver.

$$\mathbf{P1} : \min \sum_{i \in I} \sum_{j \in J} \mathcal{C}(i,j) a_{ij} \tag{2a}$$

$$\sum_{i \in I} a_{ij} = 1, \qquad \forall j \in J, \tag{2b}$$

$$\sum_{j \in J} a_{ij} = 1, \qquad \forall i \in I, \tag{2c}$$

$$a_{ij} \in \{0,1\}, \qquad \forall i \in I, j \in J. \tag{2d}$$

This is a special case of an integer linear program which has a totally unimodular constraint matrix. This means that the linear programming (LP) relaxation of the problem will provide the optimal solution and therefore the problem can be solved efficiently. Even for a large number of passengers and drivers, the optimal solution can be found in seconds.

Theorem 1 (Optimal matching without congestion): *Let the number of drivers be equal to the number of passengers. All passengers and drivers have to be matched. The optimal matching (Definition 1) is the solution of problem P1.*

The optimal solution to the matching problem is not necessarily unique as multiple solutions can lead to the same optimal objective value. A trivial example is the case where all drivers are located before all passengers, in which every possible matching has an objective value of 0 as all cars are the same from the passengers' point of view. The optimal matching of drivers and passengers is formulated in the following lemma. A formal proof can be found in the Appendix.

Lemma 2 (Characterization): *Let the number of drivers be equal to the number of passengers. All passengers and drivers have to be matched. Under assumption A1, the matching where the i^{th} driver and the i^{th} passenger are matched is always among the set of optimal matches.*

Intuitively, drivers and passengers are matched based on the order in which they are ranked from furthest to closest to the CBD. We emphasize that this solution is not necessarily unique in the sense that two passengers can be switched without changing the value of the objective.

The optimal matching may not necessarily be stable. In a stable matching, there is no other match where both the driver and the passenger prefer each other over their current match. The advantage of such a matching is that no-one is able to improve herself by deviating from the current match. The concept of stability in matching is borrowed from marriage economics (Browning et al. (2014)). However, in a carpooling framework as we consider, it is reasonable to assume that individuals are unaware of their alternatives. In that case, a regulator can impose the optimal matching. Alternatively, the stable formulation by Wang et al. (2018) can be used to replace formulation $P1$. In that case, the matching will be stable but not necessarily optimal under our definition of optimality.

Another concern is stability in the choice of individuals to participate as either drivers or passengers in the carpooling system. Drivers will always agree to carpool if they receive a subsidy that is higher than their perceived inconvenience (i.e. detour and possibly scheduling delay). A passenger will always agree to carpool if their cost of the carpooling plus the perceived inconvenience are lower than the cost of their alternative, such as public transport and driving alone. In this paper, we assume these requirements are met, unless stated otherwise, such that we do not need to consider subsidies in our matching framework. Subsidies are considered explicitly in Section 3.5.

3.2 Unequal number of drivers and passengers

When the number of drivers and passengers are unequal, we can still construct and solve the model discussed in Section 3.1 in a similar way. In that case, we have to create dummy drivers or passengers such that every driver and passenger still has exactly one match. The same formulation is used, but costs for assigning dummy drivers or passengers have to be defined. We define a dummy driver/passenger as follows:

Definition 2 (Dummy driver/passenger): *Alternative matching option for passengers or drivers. If an individual is matched to a dummy driver/passenger, this means this individual is travelling alone and therefore not carpooling.*

We consider two cases. In the first, there are more drivers than passengers so we include a set of dummy passengers. In the second, there are more passengers than drivers so we include a set of dummy drivers.

We assume that drivers and passengers always prefer to carpool over travelling alone (either by private or public transportation). That is, matching a driver to a passenger is always preferred over matching a driver to a dummy passenger. Similarly, a passenger prefers to be matched to a driver than being matched to a dummy driver. For a driver, this means that the cost of carpooling is lower than the subsidy received. For the passenger, this means that the cost of carpooling is lower than the cost of public transportation. Note that the cost of carpooling experienced by drivers and passengers depends on the share of the matching costs $\mathcal{C}(i,j)$, which is assigned to either of them.

To ensure that carpooling is preferred for any distribution of the costs among the driver and passenger, we consider the two extreme cases. For a driver, the cost of carpooling is at most the total cost of the match. Therefore, the cost of matching driver i to a dummy passenger is defined as follows:

$$\mathcal{C}\left(i,d\right) = c^A \geq \max_{j \in J}\{2\alpha \max(x_i - y_j, 0)\} \qquad \forall i \in I, \tag{3}$$

where the cost of driving alone c^A can be interpreted as the lost subsidy that would have been received for carpooling.

For a passenger, in the worst case all costs of the match are assigned to them. The alternative would be public or private transportation. We assume all passengers have a fixed cost of alternative transportation, c^T, which may include possible inconvenience experienced for using public transport. To ensure that a passenger always prefers to carpool, the cost of public transportation should be at least as high as the cost of carpooling. This is satisfied for any distribution of the matching cost between the driver and the passenger if the following equation is satisfied:

$$\mathcal{C}\left(d,j\right) = c^T \geq \max_{i \in I}\{2\alpha \max(x_i - y_j, 0)\} \qquad \forall j \in J. \tag{4}$$

This suggests that the cost of alternative transportation is at least as high as the cost of carpooling. We emphasize that in general the cost and a possible subsidy are distributed between the driver and the passenger. In this case, it is reasonable to assume that carpooling is preferred by both the driver and the passenger over the alternative if the subsidy is sufficiently high. The optimal matching when the number of drivers is unequal to the number of passengers can be characterized as in Theorem 3. Dummy drivers are assigned to the left-most passengers whereas dummy passengers are assigned to the right-most drivers. For the remaining drivers and passengers, the same structure as described in Lemma 2 is observed.

Theorem 3 (Optimal matching without congestion): *Consider a sequence of m drivers and n passengers both ranked from left to right. Under assumption A1, the following match is optimal:*

1. *If $m < n$, the $n - m$ left-most passengers are not matched (i.e. matched to dummy drivers).*
2. *If $m > n$, the $m - n$ right-most drivers are not matched (i.e. matched to dummy passengers).*

3. The remaining $\min(m, n)$ *drivers and passengers are matched according to Lemma 2.*

If the requirements in Equations (3) and (4) are not satisfied, an individual may not be willing to carpool even though she is matched. The individual may then go alone if her carpooling costs are too high compared to the alternative. This can be incorporated in the model by adding a constraint on Equations (3) and (4) to the model. As total unimodularity of the constraint matrix is lost by adding these constraints, they increase computation time for large instances and are therefore omitted in this work. We also emphasize that a similar intuition applies to the cost structures that incorporate matching inconvenience such as scheduling delay.

3.3 Matching with Inconvenience

For some passengers or drivers, matching cost may not solely be based on the additional distance a driver has to drive to perform the pickup. Therefore, we add an additional term we refer to as inconvenience of a match. This inconvenience may either be separable, solely dependent on the individual (m_i), or not separable, dependent on the specific match (m_{ij}). The former case is straightforward under the assumption that every driver and every passenger has to be matched, as this does not influence the matching. This is summarized in the following result (of which the proof is left to the reader):

Lemma 4 (Inconvenience): *Let the number of drivers be equal to the number of passengers. All passengers and drivers have to be matched. If matching costs consist of detour and separable inconvenience, that is, individual specific and independent from the match (i.e.* $m_{ij} = m_i + m_j$*), the matching is the same as for the case without inconvenience costs (Lemma 2).*

The latter case, where inconvenience costs are not separable, may not be trivial. An example of such inconvenience is scheduling delay penalties, which depend on the matched pair rather than only the individual. In this case, the optimal matching problem may not have the same characterization as in Lemma 2. To investigate the effect of matching inconvenience we include inconvenience costs in the cost definition. The importance of inconvenience is determined through the parameter ρ, such that the relative importance of inconvenience to cost of travel time is given as $\frac{\rho}{\alpha}$. The matching cost including inconvenience are as follows:

$$\hat{C}\,(i, j) = C(i, j) + \rho m_{ij} = 2\alpha \max(x_i - y_j, 0) + \rho m_{ij} \tag{5}$$

where m_{ij} represents the non-separable inconvenience term experienced from matching driver i to passenger j.

3.4 Scheduling Delay

We consider below a specific type of inconvenience cost: the schedule delay cost. Until now, we assumed drivers and passengers had no scheduling delay preferences. However, in many scenarios the desired arrival time is different for every individual and individuals are penalized for early and late arrivals. Therefore, we present a model that incorporates time preferences and scheduling delay. Every individual i has a desired arrival time τ_i. In this case, drivers and passengers also need to agree on an arrival time when they are matched. The matching costs are therefore a trade-off between travel time and scheduling delay as first introduced by Vickrey (1969). If travel time is independent of departure time, the total matching cost including scheduling delay penalties is:

$$\mathcal{C}(i,j,t) = 2\alpha \max(x_i - y_j, 0) + \beta\Big((\tau_i - t)^+ + (\tau_j - t)^+\Big) + \gamma\Big((t - \tau_i)^+ + (t - \tau_j)^+\Big), \qquad (6)$$

where t is the agreed arrival time, β is the earliness penalty and γ is the lateness penalty. We assume $\gamma \geq \beta$ (lateness is weighted higher than earliness), consistent with the empirical literature. For a given match (i,j) the optimal joint arrival time t^* can be determined according to the following theorem (of which the proof can be found in the appendix) and illustrated graphically as in Figure 2. We have:

Theorem 5.1 (Linear scheduling delay): *Consider linear scheduling delay with $\gamma \geq \beta$. Driver i is matched to passenger j with desired arrival times τ_i and τ_j. The optimal arrival time t^* that minimizes the total earliness and lateness penalty as given in Equation (6) is given as $t^* = \min(\tau_i, \tau_j)$. In this case, the reduced cost of matching driver i to passenger j is $\tilde{\mathcal{C}}(i,j) = 2\alpha \max(x_i - y_j, 0) + \beta|\tau_j - \tau_i|$.*

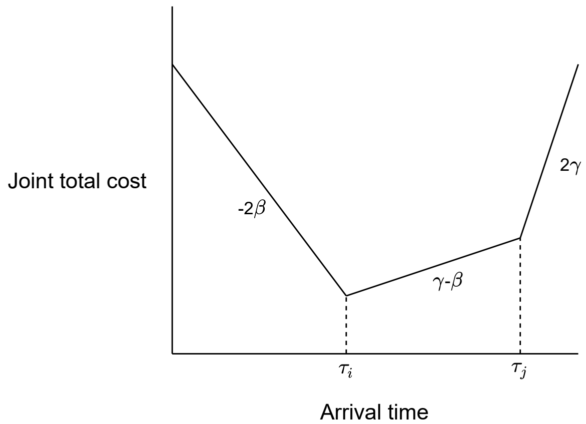

Figure 2: Scheduling delay cost for a given match

We emphasize that the reduced matching cost is independent of t^*. Using this property, the problem can be decomposed such that the optimal t^* can be determined for every potential match first and the optimal matching can be determined thereafter using the reduced matching cost. If earliness and lateness are not computed as piece-wise linear functions as in Equation (6), Theorem 5.1 does not necessarily hold. However, for many functions a similar closed form solution can be obtained. For any pair of functions satisfying the quasi-convex definition of earliness and lateness, the following theorem holds. When scheduling-delay penalties are quasi-convex, we have:

Theorem 5.2 (Quasi-convex scheduling delay): *Consider quasi-convex scheduling delay with $\gamma \geq \beta$. Driver i is matched to passenger j with desired arrival times τ_i and τ_j and $\tau_i \leq \tau_j$, without loss of generality. The optimal arrival time t^* that minimizes the total earliness and lateness is in the closed bounded interval $[\tau_i, \tau_j]$.*

This theorem can be used to efficiently evaluate the optimal arrival time for a given match. For example, if earliness and lateness are piece-wise quadratic instead of linear functions, $t^* = \frac{\gamma\tau_i + \beta\tau_j}{\gamma+\beta}$

where $\tau_i \leq \tau_j$ without loss of generality. Note that this only holds if travel time is independent of the departure time. In the presence on congestion, travel time and the optimal departure time are not independent and therefore Theorem 5.1 and 5.2 do not necessarily hold.

3.5 Targeted Subsidies

It is possible for the operator to subsidise some drivers more than others to stimulate more matches. In this section, we describe the potential of using targeted subsidies to reduce the total subsidy that needs to be paid. We offer a subsidy b_i to driver i and we assume the total budget is equal to B. This changes the matching costs and the formulation.

We emphasize that although we consider the basic carpooling model without inconvenience and scheduling delay here, inconvenience can be included in the matching costs without changing the formulation. Thereby, we add the following constraint to the formulation:

$$\sum_{i \in I} b_i \leq B. \tag{7}$$

This constraint defines the total budget of subsidies, B. Note that in this new formulation b_i is a decision variable and therefore the objective function is no longer linear but quadratic, which makes the problem more difficult to solve. Therefore, we reformulate the problem as follows, where $c_{ij} = 2\alpha \max(x_i - y_j, 0)$:

$$\textbf{P2:} \quad \max \sum_{i \in I'} \sum_{j \in J'} a_{ij} \tag{8a}$$

$$\sum_{i \in I} a_{ij} = 1, \qquad \forall j \in J', \tag{8b}$$

$$\sum_{j \in J} a_{ij} = 1, \qquad \forall i \in I', \tag{8c}$$

$$\sum_{j \in J} c_{ij} a_{ij} \leq b_i \qquad \forall i \in I', \tag{8d}$$

$$\sum_{i \in I'} b_i \leq B, \tag{8e}$$

$$a_{ij} \in \{0, 1\}, \qquad \forall i \in I, j \in J. \tag{8f}$$

The sets of drivers and passengers I and J include both actual and dummy drivers and passengers. The sets I' and J' are limited to actual drivers and passengers. The objective is to maximize the number of actual matches. Every driver and passenger has to be matched either to an actual or dummy driver or passenger. A match is only made if the subsidy b_i exceeds the cost of the match, which is imposed by Constraints (8d). Constraint (8e) is the budget constraint. This is summarized, without a proof, in the following theorem:

Theorem 6 (Targeted subsidies): *If subsidies b_i vary across drivers and there is a given budget B, the optimal matching and corresponding subsidies are given as the solution of P2.*

Other than the previous problem, the constraint matrix of $P2$ is not totally unimodular making it harder to solve this problem. The objective is to minimize the number of matches to dummy passengers, which is equivalent to maximizing the number of actual matches. As in $P1$, the optimal solution may not necessarily be unique, in the sense that multiple solutions attain the same objective value.

4 Matching in the presence of bottleneck congestion

During the morning commute, the access roads to the CBD are often heavily congested. Travellers cannot always arrive on time at their destination and they may experience both scheduling and travel time delays. The effect of congestion in road networks on departure time decisions is a well-studied problem (e.g. Arnott et al. (1993) and Small et al. (2007)). Here, we consider a basic Vickrey (1969) model of road congestion. We incorporate bottleneck congestion and scheduling delay in the matching decisions of the carpooling operator and the departure time decisions of the drivers. Bottleneck congestion is implemented in a similar way as described by Arnott et al. (1993). We assume the road is uncongested except at a single bottleneck at the entrance of the CBD. Here, at most s cars can pass per unit of time, suggesting that a queue is formed if this capacity is exceeded. The queue is served according to a first-come, first-served (FIFO) policy. Similar to Arnott et al. (1993), travel time from home to work is defined as $T(t) = T^f + T^v(t)$, where T^f is fixed travel time including a possible detour caused by matching, T^v is variable travel time and t is the departure time. Variable travel time depends on the queue length at the time of arrival at the bottleneck, $Q(t')$. A driver's (or carpooler's) queueing time is then equal to:

$$T^v(t') = \frac{Q(t')}{s}. \tag{9}$$

The length of the queue can be determined using the cumulative number of arrivals at the bottleneck ($r(t)$) and the departure rate according to the capacity. Let $\hat{t}$ be the last time the queue was empty. The queue length at time $t' > \hat{t}$ is given by:

$$Q(t') = \int_{\hat{t}}^{t'} r(u)du - s(t' - \hat{t}). \tag{10}$$

Matching costs in our problem now include possible delay at the bottleneck, detour to pickup a passenger, earliness and lateness. The relationship between travel time, earliness and lateness is similar to that described in Section 3.4.

Note that bottleneck congestion does not incorporate congestion on the corridor itself. However, as during the morning commute the traffic leaving the CBD is usually very limited, congestion during the pickup of passengers is unlikely. Congestion on the corridor towards the CBD tends to increase as the bottleneck is approached, due to the increased demand for road space. Therefore, the highest level of congestion is observed at the entrance of the CBD, which is incorporated in our bottleneck. More complex city structures are beyond the scope of the paper, but clearly a challenging future direction.

The joint problem of matching and departure time choices is solved in a sequential and iterative way. The departure time choices in equilibrium are determined using an equivalent optimization problem for the bottleneck model, as described by Iryo and Yoshii (2007). By solving an LP, the equilibrium departure times and corresponding cost functions can be determined. Variable x_{ti} describes the number of drivers in class $i \in I$ that choose time $t \in T$ to leave the bottleneck (and also enter it in case there is no congestion). We emphasize that for the congestion model, time is discrete and not continuous. Classes are defined as the sets of drivers with equal desired arrival times. In the case of carpooling, we use the matched couples to identify a class. The total number of drivers in class i is defined as D_i and the capacity at the bottleneck is given as s. The non-bottleneck costs (i.e. scheduling delay costs) of a driver (or match) in class i that choose to leave the bottleneck at time t is given as p_{ti}. Non-bottleneck costs p_{ti} are defined as the sum of earliness and lateness penalties for both the driver and the passenger. In the case of carpooling, we treat a couple as a single entity and therefore we take the average of their non-bottleneck costs.[2] This yields the following definition of non-bottleneck costs for a driver i arriving at time t. Note that, as the matching is fixed at this stage, every matched driver has a known passenger j.

$$p_{ti} = \begin{cases} \beta\Big(0.5(\tau_i - t)^+ + 0.5(\tau_j - t)^+\Big) + \gamma\Big(0.5(t-\tau_i)^+ + 0.5(t-\tau_j)^+\Big) & \text{if driver } i \text{ is matched to passenger } j \\ \beta(\tau_i - t)^+ + \gamma(t-\tau_i)^+ & \text{if driver } i \text{ drives alone} \end{cases} \tag{11}$$

The formulation is given as follows:

$$\textbf{P3:} \quad \min \sum_{t\in T}\sum_{i \in I} p_{ti} x_{ti} \tag{12a}$$

$$\sum_{i\in I} x_{ti} \leq s \qquad \forall t \in T \tag{12b}$$

$$\sum_{t\in T} x_{ti} = D_i \qquad \forall i \in I \tag{12c}$$

$$x_{ti} \geq 0 \qquad \forall t \in T, i \in I \tag{12d}$$

The objective is to minimize the total non-bottleneck costs given that the capacity is satisfied at every time interval and every driver (or match) chooses a time interval. According to Iryo and Yoshii (2007), the solution to this problem is equivalent to the equilibrium departure time choices of all drivers if the FIFO constraint is satisfied (FIFO is not enforced by the formulation proposed by Iryo and Yoshii (2007) and has to be checked separately). Note that the solution is not necessarily integer and therefore we obtain a continuous solution. This can be interpreted as a driver choosing to leave the bottleneck at various times with fixed probabilities. For example, a driver leaves at interval t with probability 0.2 and at interval $t+1$ with probability 0.8.

Iryo and Yoshii (2007) define the dual formulation, to obtain the delay w_t for a given departure time $t \in T$. We use dual variables w_t and θ_i to solve the following formulation:

$$\textbf{P4:} \quad \max \sum_{i\in I} D_i\theta_i - \sum_{t\in T} s w_t \tag{13a}$$

$$p_{ti} + w_t \geq \theta_i \qquad \forall t \in T, i \in I \tag{13b}$$

$$w_t \geq 0 \qquad \forall t \in T \tag{13c}$$

[2]Equal coefficients ensures optimality for driver and passenger if we assume costs can be redistributed between driver and passenger. The proof hereof is straightforward and left to the reader.

By solving the dual formulation, we can approximate the delay in the equilibrium state. The delay in equilibrium for every discrete time step t is given by w_t. As we consider a matched couple as a single agent, we are able to incorporate scenarios where a part of the drivers are driving alone, while another part is carpooling.

We use an iterative approach to determine the optimal matching decisions in equilibrium. This iterative procedure is summarized in Algorithm 1. The algorithm is initialized with the uncongested travel time (no delay at the bottleneck). The optimal matching and consequently the equilibrium departure times given this matching are determined using the aforementioned optimization problem. The matching costs are approximated using the estimated delay and for a fixed discrete departure time (i.e. the departure time that minimizes the matching costs) and may therefore slightly deviate from the matching costs using the equilibrium formulation where departure times may be probabilistic but discrete. We choose to keep departure time integer to restrict the computation time. The delay is updated using the estimated delay from the dual function as a weighted moving average. The predicted delay is updated with the delay from last iteration using a fraction λ^k. After every iteration, λ^k is decreased by multiplying it by a constant δ until λ_{min} is reached. This represents the increasing confidence in the estimated delay as the number of iterations increases. Delay estimates are updated as follows where $\hat{W}^k$ are the predicted delays for iteration k and W^k are the experienced delays for iteration k:

$$\hat{W}^{k+1} = \lambda^k W^k + (1 - \lambda^k)\hat{W}^k \tag{14a}$$

$$\lambda^{k+1} = \max(\delta\lambda^k, \lambda_{min}) \tag{14b}$$

At convergence, departure times are in equilibrium following the optimization problem defined by Iryo and Yoshii (2007) and the matching decisions are optimal given the equilibrium state. The generalized matching costs including bottleneck congestion and tardiness penalties are now as follows for driver i departing at time t:

$$\mathcal{C}(i,j,t) = \begin{cases} \alpha(D(i,j) + W(i,j,t)) + \beta E(i,j,t) + \gamma L(i,j,t), & \text{if matched to regular passenger } j \\ M, & \text{if matched to dummy passenger } j \end{cases} \tag{15}$$

Here, D represents the detour driver i makes to pickup passenger j, W is the delay at the bottleneck, E and L are earliness and lateness penalties respectively. As we only care about minimizing the costs for matched drivers and passengers, we set the costs for matching a driver to a dummy passenger to an arbitrarily large number M. Using this, we implicitly assign the drivers that form the worst match or those with the highest costs to dummy passengers. M is large enough such that all passengers are matched to a driver in the optimal solution. Bottleneck delay, earliness and lateness are experienced by both the driver and the passenger whereas the detour is only experienced by the driver. If a driver is matched to a dummy passenger, the inconvenience experienced by this driver is not incorporated in the matching costs and therefore set to M. To attain convergence, the matching only changes if the improvement compared to the previous matching is higher than 0.1%.

Algorithm 1: Iterative Matching Approach

Initialize the waiting time at every time interval to the uncongested waiting time
while *Stopping criterion is not met* **do**
    Determine the optimal matching using formulation P1
    Given the matching, determine the equilibrium departure times using the equivalent optimization problem defined by Iryo and Yoshii (2007) using formulation P3
    Given the matching and the equilibrium departure times, determine the delay at every time interval using the dual formulation P4
    Update the matching costs for every potential match given the estimated delays
    Gather statistics on travel time, delay, earliness and lateness based on current matching and actual waiting time.
end

In accordance with the updated matching costs, we redefine the notion of optimal matching as follows:

Definition 3 (Optimal matching with congestion): *The optimal matching minimizes the sum of detour costs, inconvenience costs (including scheduling delay costs) and delay at the bottleneck for all matched passengers and drivers.*

The updated definition of optimal matching requires to solve a dynamic bottleneck congestion problem. The optimal matching in equilibrium is such that the central operator does not want to change the obtained matching and the individuals do not want to change their departure times. As such, delay at the bottleneck and the departure time decisions are in equilibrium and therefore the corresponding optimal matching does not change. Due to the computational effort in this iterative framework, we do not consider targeted subsidies in the presence of bottleneck congestion here. This is left as a direction of future work.

4.1 Equilibrium Analysis

An analytical solution can be obtained for the case where the desired arrival time t^* is equal for every individual. This solution can be identified using the following theorem:

Lemma 7 (Dynamic matching with congestion): *Consider the dynamic carpooling model with n passengers and m drivers, each with $\alpha - \beta - \gamma$ (scheduling) delay penalties. If all users have the same desired arrival time t^*, then the equilibrium departure times can be determined using* $\min(n, m)$ *non-carpooling individuals as described by Arnott et al. (1993).*

Proof. Departure time is determined as a couple agreement between driver and passenger. Note that this is independent of the travel time before the bottleneck (including possible detour for pickup) as this is fixed for every departure time. Therefore, the cost function to determine the departure time only includes waiting time and scheduling delay. As all desired arrival times are equal, scheduling delay is independent of the matching. The matching and departure time decisions are therefore independent. Given that every individual has equal desired arrival time, the match can be considered to be a single agent. Using this, the departure times in equilibrium are similar to those of a non-carpooling setting with $\min(n, m)$ individuals and can be derived analytically as described by Arnott et al. (1993). □

This result does not necessarily hold if the desired arrival times are distributed. In this case, the cost functions are not equal and therefore driver and passenger may wish to respond differently to congestion. Any other kind of cost can be included, as long as the desired arrival time is equal.

Let us elaborate the above point with a numerical example. Algorithm 1 is used to obtain the delay at convergence for various t^* profiles. We consider a profile where every individual has identical desired arrival time t^* and a profile where desired arrival times are pseudo-randomly generated. We compare the well known scenario where every individual drives alone to the scenario where everybody carpools. Figure 3a displays the delay for all time intervals. The peak delay is at t^*, which is the same for all individuals. We observe that, in accordance with the aforementioned theorem, the shape of delay for carpooling mimics the shape of delay for no-carpooling with $\min(m,n)$ drivers. The slopes are equal for both curves and are equal to $\frac{\beta}{\alpha}$ before t^* and $-\frac{\gamma}{\alpha}$ after t^*. The beginning and end of the rush hour, t_q and $t_{q'}$ respectively, are also equal to those defined in Arnott et al. (1993) and can be defined as follows, where N is the number of vehicles on the road:

$$t_q = t^* - \frac{\gamma}{\beta+\gamma}\frac{N}{s} \tag{16a}$$

$$t_{q'} = t^* + \frac{\beta}{\beta+\gamma}\frac{N}{s} \tag{16b}$$

Figures 3b and 3c display the delay at the bottleneck if t^* is not identical. We evaluate the queue on a time horizon between 0 and 20, which is divided into 1000 intervals. The desired arrival times are pseudo-randomly drawn from a uniform distribution[3] between $t^*_s = 8$ and $t^*_e = 12$. The peak of congestion is observed at $\hat{t}^* = \frac{\beta}{\beta+\gamma}t^*_s + \frac{\gamma}{\beta+\gamma}t^*_e$, which is at 11.2. This result is in agreement with those defined in Lindsey et al. (2019). We also observe that the slopes are similar to those when every individual has an identical t^* and the start and the end of the rush hour can be approximated using (16a) and (16b).

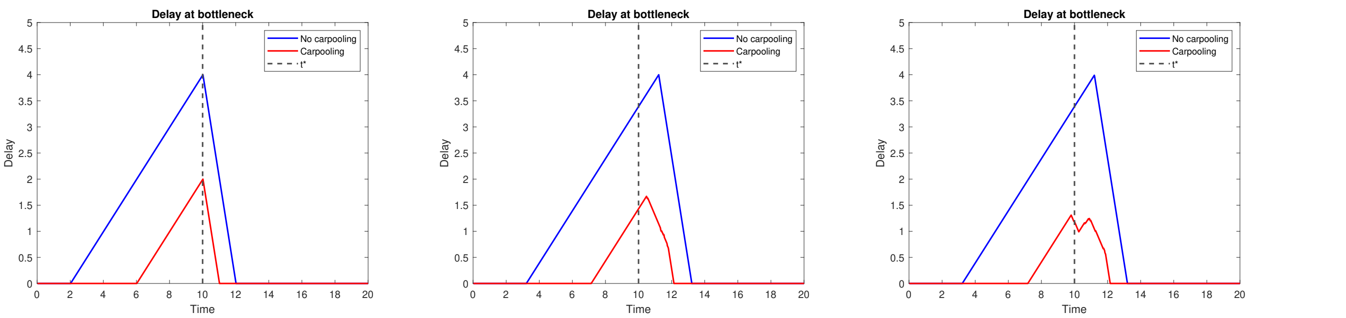

(a) Identical t^* (b) Pseudo-uniform t^* (c) Pseudo-uniform t^*

Figure 3: Delay at bottleneck

When matching exists, the uniform distribution of the desired arrival times are lost, as they are approximately pooled between the driver and the passenger. Therefore, the shape of the delay curve is lost and no theoretical value for the peak congestion time can be used. However, as the

[3]Desired arrival times are selected such that they are exactly uniformly distributed on the interval and are therefore not truly random

departure times of a matched couple are determined in a systematic way, we can still identify some descriptive properties. According to Theorem 5.1, in the absence of congestion, the optimal arrival time of a matched couple is the minimum of their two desired arrival times. We observed that in the presence of congestion, a matched couple still has a tendency towards the earliest arrival time, in agreement with this theorem. Following this, as displayed in Figure 3b, the peak of congestion is usually slightly earlier than $\hat{t}^*$. If the desired arrival times of the matched couples are close to a uniform distribution, there will be a single peak of congestion, as displayed in Figure 3b. However, if the uniformity of desired arrival times is lost after matching, the single peak congestion may also be lost. In that case, we can observe multiple peaks in the delay curve as displayed in Figure 3c. We emphasize that the delay curves in Figure 3a and those without carpooling (blue) in Figure 3b and 3c are theoretical. The delay curves with carpooling (red) in Figure 3b and 3c are simulated. A detailed analysis of delay and tardiness penalties is included in Section 5.2.

5 Numerical Results

In this section we provide numerical simulation results based on the methodology described in the previous sections. We first evaluate the theoretical results without congestion and the effect of targeted subsidies in Section 5.1. Thereafter we evaluate the optimal matching under congestion and perform extensive sensitivity analysis in Section 5.2.

5.1 Optimal Matching without Congestion - The Effect of Subsidies

As expected, simulation results have shown that in the absence of congestion, the optimal matching is in agreement with Theorems 1 - 7. While this is theoretically proven, it also allows to test that our simulations provide intuitive results. Results have also identified that, especially in large instances, the optimal matching is not necessarily unique. Consider the following example where driver i_1 and driver i_2 are both located earlier on Hotelling's line than passenger j_1 and j_2. Then, under Assumption A1, the relation in Equation (17) holds. This indicates that the optimal matching is not necessarily unique, as passengers j_1 and j_2 can be interchanged without changing the objective value.

$$\mathcal{C}(i_1, j_1) + \mathcal{C}(i_2, j_2) = \mathcal{C}(i_1, j_2) + \mathcal{C}(i_2, j_1) = 0. \tag{17}$$

To evaluate the effect of targeted subsidies instead of constant subsidies, we study various instances of Problem P2 as discussed in Section 3.5. In this experiment, drivers are only matched to passengers if this is profitable. That is, if the subsidy b is higher than the cost of the match. We compare untargeted and targeted subsidies for various budget values. Thereby, we evaluate the effect of inconvenience on the results. If inconvenience is considered, we assume $\alpha = \rho$. In addition to this, we assume $\alpha = 1$ without loss of generality. For $\alpha \neq 1$, the corresponding budget can be calculated as αB.

The results of this experiment are displayed in Figure 4. For untargeted subsidy, every driver is rewarded a fixed subsidy b. For targeted subsidy, every driver i is rewarded a subsidy b_i, which increases with the cost of matching.

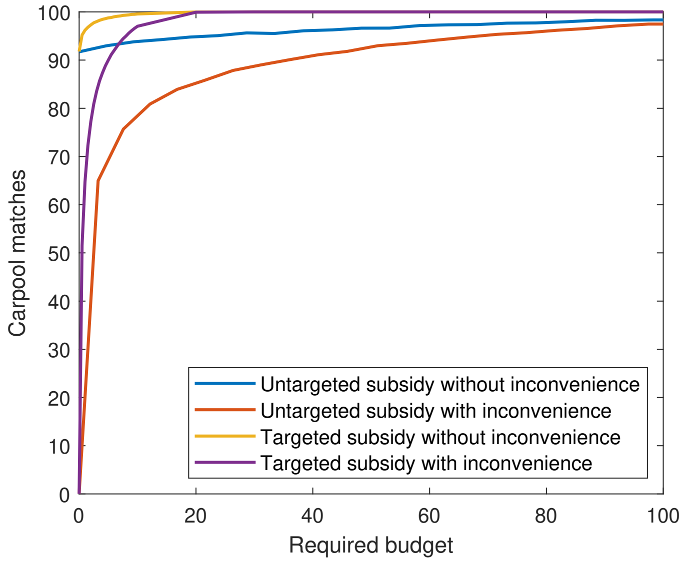

Figure 4: Effect of subsidy on carpooling activity

We observe that if inconvenience exists, the subsidy needs to be higher to achieve a similar carpooling activity. The total number of carpooling matches increases logartihmically with the budget. Furthermore, we observe that if a targeted subsidy is used, the total budget needed to achieve the same carpooling activity as with untargeted subsidy is significantly lower. For matches with low matching costs (i.e. driver is close to the passenger or inconvenience is low) the subsidy can be low, whereas the subsidy for matches with high matching costs should be higher to convince the driver to pickup a passenger. Therefore, using targeted subsidy significantly increases the carpooling activity. Considering matching inconvenience, for the same budget of 20, the number of matches can be increased by more than 15%.

While these results are based on hypothetical values, they carry the same intuition as for more complex cases. As drivers may not be willing to deviate too much, a higher subsidy needs to be awarded to drivers that make a bigger detour. Similar implications hold if drivers experience high inconvenience values. By targeting subsidies to those drivers that have the highest matching costs, more matches can be formed with the same budget. This signals the benefit of an adaptive pricing policy to stimulate carpooling. The development of such an adaptive pricing policy is clearly outside the scope of this paper and is therefore marked as a direction of further research. Several adaptive pricing policies have been proposed for the ride-sourcing problem by among others Yang et al. (2020) and Zha et al. (2018), which contains many similarities to the carpooling problem. Further research can shed light in this direction.

5.2 Optimal Matching under Bottleneck Congestion

In this section we evaluate the optimal matching under bottleneck congestion for various levels of carpooling and bottleneck capacity. We first evaluate the convergence of our algorithm. Given that we use an exact optimization problem to find the equilibrium departure times, we only need to evaluate the convergence of the matching decisions. In our analysis, we consider a total of 200 individuals. These individuals may either be passengers or drivers but we always consider to have at least as many drivers as passengers such that a passenger can always be matched. A driver may either drive alone or with a passenger. The length of the corridor is set equal to 10. We set $\lambda^0 = 1.0$, $\lambda_{min} = 0.5$ and $\delta = 0.8$ (for weighted moving average estimates in Equation (15)). An iteration limit is set at 100 iterations. Desired arrival times t^* are drawn from a uniform distribution between 8 and 12. This interval is divided into 1000 discrete time intervals.

Figure 5 presents the convergence of the algorithm for a single run. Although 100 iterations have been performed, only the first 20 are displayed as no significant differences are observed after 20 iterations. As a criterion, we use the improvement of the expected matching costs the new matching achieves over the previous matching. The expected costs are computed using the moving average estimation of waiting time as described in Equations (14a) and (14b). The results show that after 5 iterations, the objective is improved with at most 0.05% at every iteration compared to the previous iteration. This suggests that we have reached a state of approximate convergence. We emphasize that this may differ for various simulation configurations. The displayed graph considers a matching of 100 drivers and 100 passengers. In case the number of drivers is much higher than the number of passengers, absolute convergence can be reached within a few iterations. In general, the number of iterations required to attain convergence highly depends on the simulation settings.

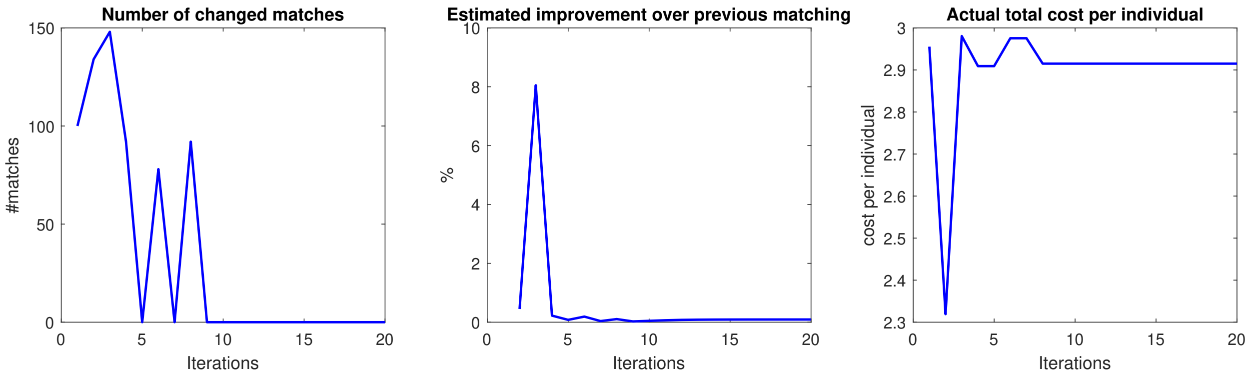

Figure 5: Convergence of matching decisions

Next, we compare various statistics for different population compositions (i.e. number of drivers and passengers). We first consider the special case where every individual has an identical t^*. The statistics are displayed in Figure 6. Looking at the first panel, we observe that, using the corresponding $\alpha - \beta - \gamma$ weights assigned to delay, earliness and lateness, costs related to scheduling delay penalties are equal to costs assigned to delay. This is in accordance with the conditions for a bottleneck equilibrium. We observe a trade-off between detour time versus delay and tardiness penalties. Observe the extreme increase of fixed travel time (due to the detour) as all passengers and drivers need to be matched. This leads to an optimal number of drivers of approximately 110 and 90 passengers. If the number of passengers is close to the number of drivers, detour is unavoidable if all passengers have to be matched. This explains the sharp increase in fixed travel time between 90 and 100 passengers. Intuitively, considering slightly more drivers than passenger allows the drivers that have the worst match with all passengers to drive alone. This gives a benefit over the case where all drivers need to be matched.

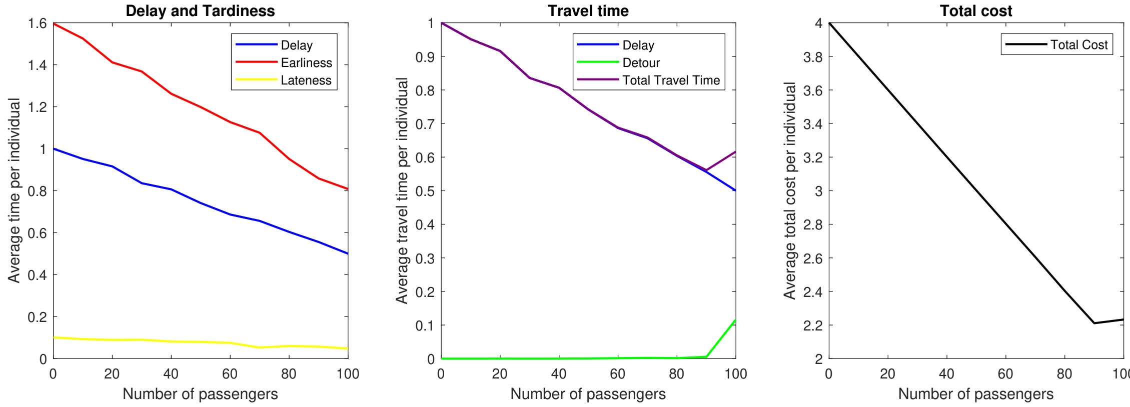

Figure 6: Cost components for varying number of passengers and drivers all with equal t^* and total number of commuters equal to 200

More realistically, we consider a pseudo-random uniform distribution of desired arrival times t^*, as defined before. We consider three levels of bottleneck capacity classified as high, medium and low (0.4, 0.8, 4.0 per time interval respectively). The results of this are displayed in Figures 7, 8 and 9 (note the different scales on the 3 y-axis). We clearly observe three distinct patterns for the different bottleneck capacities. If capacity is low, as displayed in Figure 7, delay and scheduling penalties are high whereas detour is low. This tendency is softened as the number of drivers decreases and the number of passengers increases and outweighs the increase in matching inconvenience (i.e. matching distance and trade-off in arrival time). Therefore, we observe that the optimum is reached if the number of passengers and drivers are equal ($n = m = 100$) and thereby the total flow of vehicles is minimized.[4] The results are rather similar to those where t^* is identical for all individuals. The main difference is that delay penalties are relatively high compared to scheduling delay penalties if t^* is varying. This is in line with Arnott et al. (1987), who state that bottleneck delay is equal to that with identical t^*, while scheduling delay penalties are smaller.

If capacity at the bottleneck is high, as displayed in Figure 8, no delay at the bottleneck is experienced. Contrary to low bottleneck capacity, high bottleneck capacity induces drivers to drive alone. If the number of matches increases, we observe that the fixed travel time increases as the drivers have to pickup the passengers. Thereby, we observe that scheduling delay penalties increase as a consequence of the matching, where the desired arrival time of the couple needs to be accommodated. In line with Theorem 5.1, we observe that matching increases early arrivals and decreases late arrival even in the presence of congestion. This is an unexpected consequence of carpooling from which firms may benefit.

[4]Note that this is the optimal solution given our definition of optimal but it is not necessarily the market solution nor the first-best solution.

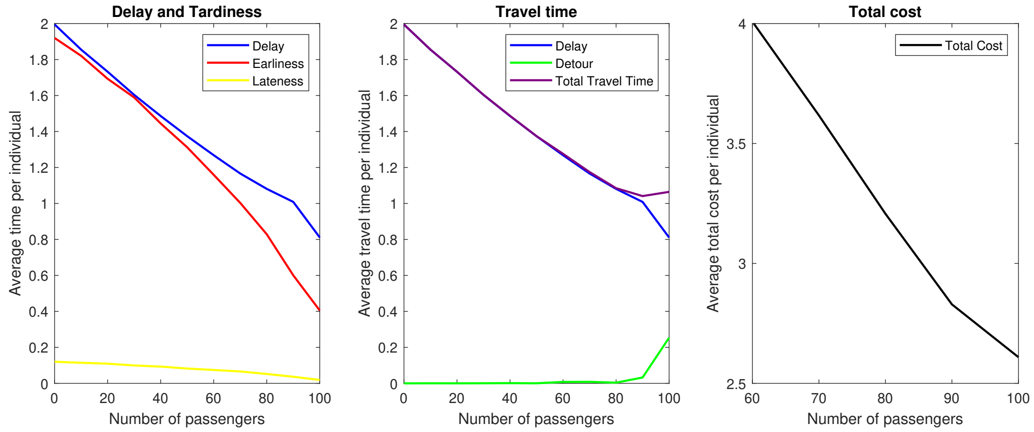

Figure 7: Cost components for varying number of passengers and drivers with pseudo-uniform t^*, total number of commuters equal to 200 and low bottleneck capacity

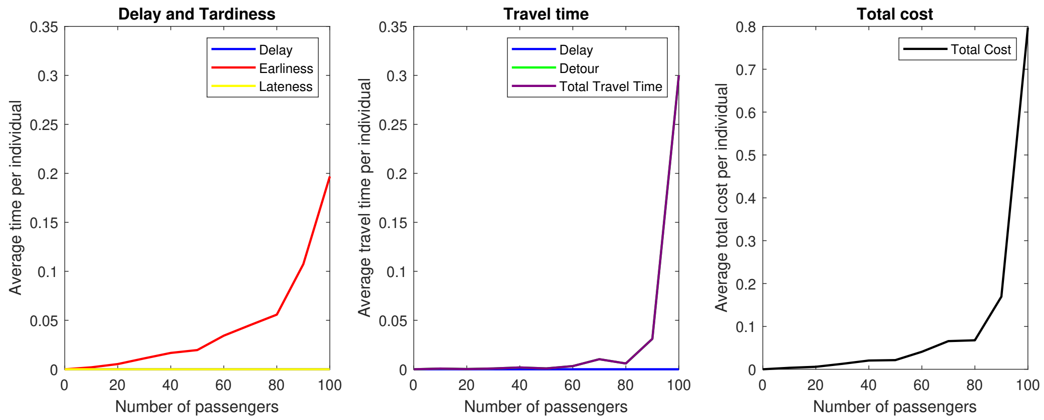

Figure 8: Cost components for varying number of passengers and drivers with pseudo-uniform t^*, total number of commuters equal to 200 and high bottleneck capacity

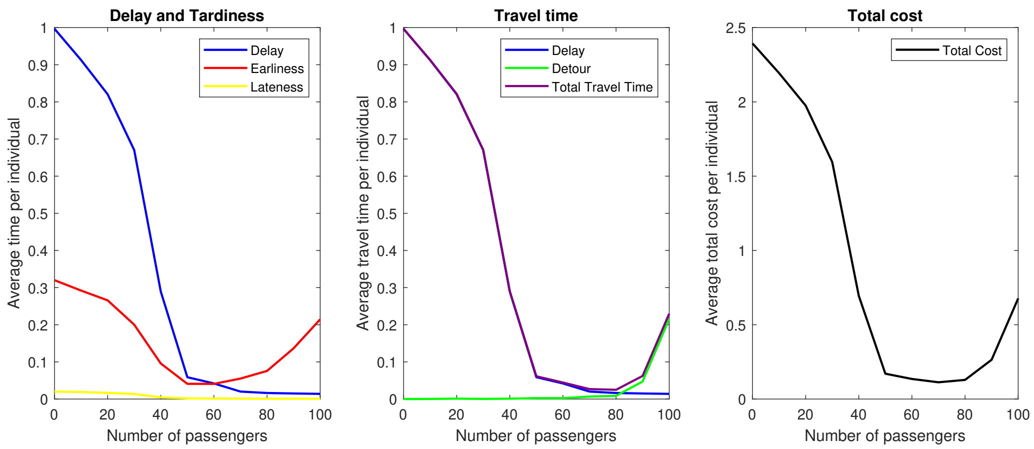

Figure 9: Cost components for varying number of passengers and drivers with pseudo-uniform t^*, total number of commuters equal to 200 and medium bottleneck capacity

For a medium bottleneck capacity, which is displayed in Figure 9, the two aforementioned effects construct a trade-off. By increasing the number passengers and thereby the number of carpooling matches, the demand for road space decreases and therefore delay decreases. However, matching inconvenience caused by increasing pickup distance for the drivers as well as earliness due to the coordination of arrival times increase and as a result the total costs increase. Therefore, the optimal number of drivers and passengers is somewhere between the two extreme cases (high and low bottleneck) we considered before. In our example, the optimal number of passengers is approximately 70 and therefore the optimal number of drivers is equal to 130.

The results indicate that carpooling induces earliness and reduces lateness. To quantify this result, we consider the scenario with 130 passengers and 70 drivers such that the number of drivers driving alone and those driving together with a passenger is approximately equal. Thereby, bottleneck capacity is equal to 0.4 cars per time interval, a corridor length of 10 and 50 runs are used to compute averages. Figure 10 displays desired and actual arrival times for carpoolers and individual drivers. The patterns of the two types of travellers are very different, whereas they are nearly identical (approximately uniform) when the bottleneck capacity is sufficiently high (no congestion).

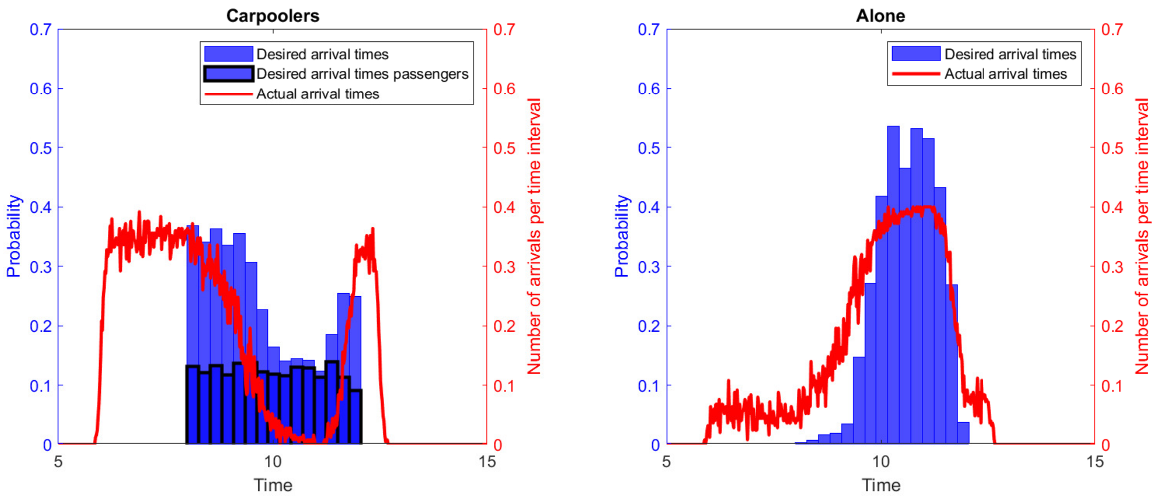

Figure 10: Pattern of desired and actual arrival times for carpoolers and individual drivers

As we have stated before, in general earliness increases as drivers tend to carpool. From this graph we observe that both the desired and actual arrival times of carpoolers are lower than that of individual drivers. We identify that individual drivers mainly arrive during the peak period of congestion, whereas carpoolers arrive either before or after this peak period. Note that, when aggregating the two types of commuters, the bottleneck operates at capacity from the moment a queue starts until the queue is completely depleted, in agreement with dynamic congestion equilibrium theory. Also note that the desired arrival times of carpoolers include those of both the driver and the passenger. As the desired arrival times are drawn from a pseudo-uniform, the arrival times of passengers are approximately uniformly distributed. It is clear that in the matching process, those drivers with the lowest bottleneck costs are chosen (i.e. those having off-peak desired arrival times) while those with high bottleneck costs are mainly left to drive alone. We emphasize the dependency of the results on the chosen parameters. For example, by changing the matching costs in Equation (15) such that the costs of unmatched drivers are also optimized, the observed pattern is lost and the arrival patterns of carpoolers and individual drivers are very similar and approximately uniform. In addition to this, a reverse pattern where carpoolers tend to arrive very late is observed in the unlikely case when $\gamma < \beta$.

The importance of matching inconvenience (detour) can be reinforced by increasing the length of the corridor. This increase brings forth an increase of the fixed travel time and of the detour. Therefore, it may change the matching accordingly. We emphasize that the longer length of the corridor does not have a direct effect on the congestion at the bottleneck, but only affects the bottleneck through the channel of the matching. To illustrate the effect of the corridor on the matching, Figure 11 displays the matches for a corridor of length 1, 10 and 100. The points mark the drivers (blue) and passengers (red) based on their location on the corridor and their desired arrival time. The line segments between the points indicate the matches. We observe that as the length of the corridor increases, distance between driver and passenger becomes increasingly important and therefore the graphs display more vertical distance between the matches and less horizontal distance. The horizontal distance where the driver (blue) is situated in front of the passenger (red) does not create any detour as the driver can pick up the passenger on his way.

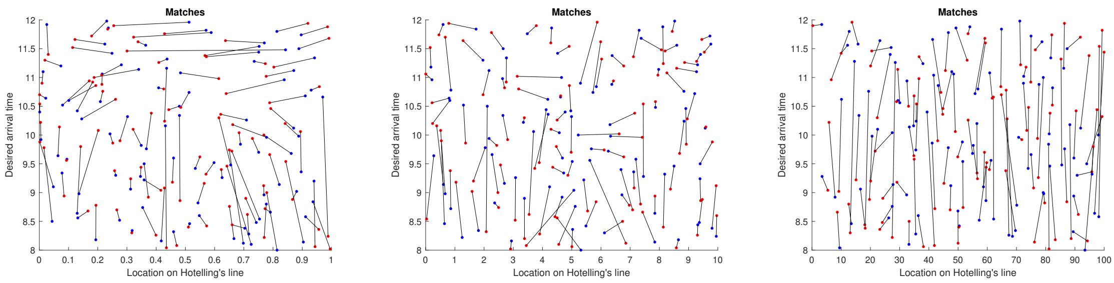

(a) Corridor of length 1 (b) Corridor of length 10 (c) Corridor of length 100

Figure 11: Matches of drivers and passengers for 3 different corridor lengths

A similar conclusion can be drawn from the total detour the driver makes to pickup the passenger and the total difference between desired arrival times of matches. Table 1 reports the total detour as a fraction of the corridor length and the total absolute difference between arrival times of drivers and passengers over all matches (time difference).

Table 1: Statistics of matching for three different corridor lengths

Corridor length	$\frac{\text{Total detour}}{\text{Corridor length}}$	Time difference
1	8.3	40.0
10	3.7	44.6
100	1.8	87.2

Time difference denotes the absolute difference between the desired arrival time of the driver and that of the matched passenger. Note that this induces tardiness penalties due to the coordination of arrival times.

We observe that as the length of the corridor increases, more weight is given to location rather than desired arrival time causing the total detour to decrease and the total time difference to increase. This suggests that the longer it takes to traverse the highway, the lower the effect of scheduling delay on the matching (although penalties will increase) and the higher the effect of the detour.

Carpooling has a positive effect on congestion, independent of the initial level of congestion. The magnitude of this effect, however, depends on the network configurations (most importantly the bottleneck capacity but also the length of the corridor, arrival time distribution, etc.). As carpooling increases inconvenience, whether carpooling is beneficial highly depends on the initial level of congestion. In highly congested networks (those with a small bottleneck capacity) carpooling is likely to reduce the costs per individual because the reduced congestion outweighs matching inconvenience. On the other hand, in networks where congestion is relatively small, reduced congestion does not compensate for matching inconvenience.

Carpoolers and non-carpoolers respond differently to congestion. Specifically, carpoolers have a tendency to depart earlier and avoid the peak period of congestion. The optimal matching and the corresponding optimal departure times are determined accordingly, such that carpoolers arrive in the wing of the bottleneck, whereas individual drivers arrive somewhat in the middle. The above analysis emphasizes the importance to model congestion in tandem with carpooling/matching.

6 Conclusion

We developed a framework for matching in carpooling or ride-sharing incorporating dynamic bottleneck congestion. We first described various fundamental properties of matching decisions in carpooling. Our theoretical results showed that, if the only matching inconvenience is a possible detour, a system optimal matching is obtained when the drivers and passengers are matched based on their location in the sequence of drivers and passengers. When individuals differ in their desired arrival time, matching induces tardiness penalties. For a given match, their jointly optimal desired arrival time is the lowest desired arrival time of the matched couple. Matching therefore induces earliness and reduces lateness.

In order to evaluate the two-way causal effect of dynamic congestion and matching decisions, we developed an iterative approach where matching decisions are adapted based on bottleneck congestion. In the presence of bottleneck congestion, we still observe a tendency towards early departure for carpoolers, thereby moving the peak of congestion forward in time. Our experimental results show that our algorithm tends to converge to a near-optimal solution within 20 iterations. Thereby, we observe that the optimal number of drivers and passengers depends on the trade-off between matching inconvenience induced by detour and tardiness penalties on the one hand and reduced congestion on the other. This suggests that the optimal matching as well as the optimal number of drivers and passengers in a system is dependent on the level of congestion, but also on the length of the corridor and the variability of desired arrival times.

We have proposed a first model on how matching models may be adapted in case of congestion. Congestion distorts the matching costs and may therefore distort the optimal matching. In that case, theoretical rules may no longer apply. A central operator that imposes the matching is more likely to know traffic conditions more precisely and in particular what the effect of the proposed matching on traffic congestion will be. As private agents are less likely to know about traffic conditions, in this case the market will lead to another solution, if congestion is not internalized via first-best road congestion. This justifies the choice for a central operator that imposes the matching. The implementation of our algorithm requires more testing for a realistic network, which will be the task of future work.

Note that if a car owner decides to carpool (instead of using his/her car), it is not so clear that mobility in the household will decrease. As an example in California (which may not be representative), according to Riggs (2020) while total work-based trip and VMT (vehicle-miles-travelled) declined slightly due to telework due to COVID-19, an increase in total trips is observed from 3.97 to 4.45. Given the high emissions during the cold start, this may have heavy environmental consequences, to be considered. Moreover, there is not a clear proportionality between car usage and energy saving. Fu et al. (2012) estimated that a 5 percent shift of the regular commuting workforce to full-time remote working would result in a net energy saving of 0.36 percent relative to the total transport energy use.

We have assumed so far that any commuter is either a driver or a passenger. This is not necessarily the case, as individuals with a car may still decide to carpool as a passenger. The matching problem can be extended to this more complex setting, which has to be studied in future research. As another direction of further research, we consider the incorporation of road pricing and road segmentation in our matching framework. Thereby, subsidies to stimulate carpooling and the corresponding decisions of individuals to carpool or not have to be considered. Furthermore, pricing policies that can enforce an equilibrium mode choice for individuals are an important and natural extension of this work.

Acknowledgements

The authors would like to thank Pierre-André Chiappori and Alfred Galichon for early discussions on the modelling of traffic congestion with matching. They both strongly encouraged us to persue this line of research. We also benefited from useful comments and references from Robin Lindsey and Jesse Jacobson. The first author was supported by the NSFC-JPI UE joint research project "MAAT" (project no. 18356856).

Appendix: Proofs

Lemma 2: *Let the number of drivers be equal to the number of passengers. All passengers and drivers have to be matched. Under assumption A1, the matching where the i^{th} driver and the i^{th} passenger are matched is always among the set of optimal matches.*

Proof. Using drivers at location x_i and x_j and passengers at location y_k and y_l where without loss of generality $x_i \leq x_j$ and $y_k \leq y_l$. Furthermore we assume $\alpha = 1$ without loss of generality. We show that the following inequality always holds:

$$\mathcal{C}(i,k) + \mathcal{C}(j,l) \leq \mathcal{C}(i,l) + \mathcal{C}(j,k). \tag{18}$$

To prove this, we consider six distinct scenarios that together comprise all possible realizations.

$$\begin{aligned}
x_i \leq x_j \leq y_k \leq y_l: &\quad \mathcal{C}(i,k) + \mathcal{C}(j,l) = 0 + 0 \leq 0 + 0 = \mathcal{C}(i,l) + \mathcal{C}(j,k) \\
x_i \leq y_k \leq x_j \leq y_l: &\quad \mathcal{C}(i,k) + \mathcal{C}(j,l) = 0 + 0 \leq 0 + 2(x_j - y_k) = \mathcal{C}(i,l) + \mathcal{C}(j,k) \\
x_i \leq y_k \leq y_l \leq x_j: &\quad \mathcal{C}(i,k) + \mathcal{C}(j,l) = 0 + 2(x_j - y_l) \leq 0 + 2(x_j - y_k) = \mathcal{C}(i,l) + \mathcal{C}(j,k) \\
y_k \leq x_i \leq x_j \leq y_l: &\quad \mathcal{C}(i,k) + \mathcal{C}(j,l) = 2(x_i - y_k) + 0 \leq 0 + 2(x_j - y_k) = \mathcal{C}(i,l) + \mathcal{C}(j,k) \\
y_k \leq y_l \leq x_i \leq x_j: &\quad \mathcal{C}(i,k) + \mathcal{C}(j,l) = 2(x_i - y_k) + 2(x_j - y_l) \leq 0 + 2(x_j - y_k) = \mathcal{C}(i,l) + \mathcal{C}(j,k) \\
y_k \leq x_i \leq y_l \leq x_j: &\quad \mathcal{C}(i,k) + \mathcal{C}(j,l) = 2(x_i - y_k) + 2(x_j - y_l) \leq 2(x_i - y_l) + 2(x_j - y_k) = \mathcal{C}(i,l) + \mathcal{C}(j,k)
\end{aligned}$$

This shows that for every pair of matches, it is always better to match them in order. Following this, we can start from any matching, iteratively select a pair of matches that are not matched in order and interchange them. According to (18) this will never increase the objective value. This will always lead to the matching where the i^{th} driver is matched to the i^{th} passenger which is therefore at least as good as any other matching. Specifically, the matching where the i^{th} driver is matched to the i^{th} passenger is at least as good as the optimal matching. Therefore, the matching where the i^{th} driver is matched to the i^{th} passenger is always among the set of optimal matchings. □

Theorem 3.1: *Given that the number of drivers exceeds the number of passengers ($m > n$), under assumption A1, it is always optimal to not match the last $m - n$ drivers and match the remaining drivers according to Lemma 2.*

Proof. We consider a driver i_2 and another driver i_1 located before driver i_2. Thereby j is a passenger and d is a dummy passenger. By definition, $x_{i_1} \leq x_{i_2}$. Furthermore we assume $\alpha = 1$ without loss of generality. We show that the following inequality holds:

$$\mathcal{C}(i_1,d) + \mathcal{C}(i_2,j) \geq \mathcal{C}(i_2,d) + \mathcal{C}(i_1,j) \tag{19}$$

To prove this, we consider three distinct scenarios that together comprise all possible realizations.

$$\begin{aligned}
x_{i_1} \leq x_{i_2} \leq y_j: &\quad \mathcal{C}(i_1,d) + \mathcal{C}(i_2,j) = c^A + 0 \geq c^A + 0 = \mathcal{C}(i_2,d) + \mathcal{C}(i_1,j) \\
x_{i_1} \leq y_j \leq x_{i_2}: &\quad \mathcal{C}(i_1,d) + \mathcal{C}(i_2,j) = c^A + 2(x_{i_2} - y_j) \geq c^A + 0 = \mathcal{C}(i_2,d) + \mathcal{C}(i_1,j) \\
y_j \leq x_{i_1} \leq x_{i_2}: &\quad \mathcal{C}(i_1,d) + \mathcal{C}(i_2,j) = c^A + 2(x_{i_2} - y_j) \geq c^A + 2(x_{i_1} - y_j) = \mathcal{C}(i_2,d) + \mathcal{C}(i_1,j)
\end{aligned}$$

This shows that for any matching, the matching where a later driver is assigned a dummy passenger is always an improvement. This can be applied iteratively until the last $m - n$ drivers in the sequence are assigned dummy passengers. Therefore, it is optimal to assign dummy passengers to the last $m-n$ drivers.

Next, we remove those last $m-n$ drivers from the considered set of drivers. We are therefore left with n drivers and n passengers. Those drivers and passenger are matched optimally as described by Lemma 2. $\square$

Theorem 3.2: *Given that the number of passengers exceeds the number of drivers ($n > m$), under assumption A1, it is always optimal to not match the first $n-m$ passengers and match the remaining passengers according to Lemma 2.*

Proof. We consider a passenger j_1 and another passenger j_2 located after passenger j_1. Thereby i is a driver and d is a dummy driver. By definition, $y_{j_1} \leq y_{j_2}$. Furthermore we assume $\alpha = 1$ without loss of generality. We show that the following inequality holds:

$$\mathcal{C}(i, j_1) + \mathcal{C}(d, j_2) \geq \mathcal{C}(i, j_2) + \mathcal{C}(d, j_1) \tag{20}$$

To prove this, we consider three distinct scenarios that together comprise all possible realizations.

$$\begin{aligned}
x_i \leq y_{j_1} \leq y_{j_2}: &\quad \mathcal{C}(i, j_1) + \mathcal{C}(d, j_2) = 0 + c^T \geq 0 + c^T = \mathcal{C}(i, j_2) + \mathcal{C}(d, j_1) \\
y_{j_1} \leq x_i \leq y_{j_2}: &\quad \mathcal{C}(i, j_1) + \mathcal{C}(d, j_2) = 2(x_i - y_{j_1}) + c^T \geq 0 + c^T = \mathcal{C}(i, j_2) + \mathcal{C}(d, j_1) \\
y_{j_1} \leq y_{j_2} \leq x_i: &\quad \mathcal{C}(i, j_1) + \mathcal{C}(d, j_2) = 2(x_i - y_{j_1}) + c^T \geq 2(x_i - y_{j_2}) + c^T = \mathcal{C}(i, j_2) + \mathcal{C}(d, j_1)
\end{aligned}$$

This shows that for any matching, the matching where an earlier passenger is assigned a dummy driver is always an improvement. This can be applied iteratively until the first $n-m$ passengers in the sequence are assigned dummy drivers. Therefore, it is optimal to assign dummy drivers to the first $n-m$ passengers.

Next, we remove those first $n-m$ drivers from the considered set of passengers. We are therefore left with m drivers and m passengers. Those drivers and passengers are matched optimally as described by Lemma 2. $\square$

Theorem 5.1: *Consider linear scheduling delay with $\gamma \geq \beta$. Driver i is matched to passenger j with desired arrival times τ_i and τ_j. The optimal arrival time t^* that minimizes the total earliness and lateness penalty as given in Equation (6) is given as $t^* = \min(\tau_i, \tau_j)$. In this case, the reduced cost of matching driver i to passenger j is $\tilde{\mathcal{C}}(i, j) = 2\alpha \max(x_i - y_j, 0) + \beta|\tau_j - \tau_i|$.*

Proof. Following Theorem 5.2, $t^* \in [\min(\tau_i, \tau_j), \max(\tau_i, \tau_j)]$. As we consider a closed bounded interval, we can use Weierstrass theorem to prove that $t^* = \min(\tau_i, \tau_j)$. For this, we consider the boundary points, all points that are non-differentiable and all points where the gradient is equal to zero. As the gradient is never equal to zero on the chosen interval, we consider $t = \tau_i$ and $t = \tau_j$.

Without loss of generality we assume $\tau_i \leq \tau_j$. Then,

$$E(t) + L(t) = \begin{cases} \beta(\tau_j - \tau_i) & \text{for } t = \tau_i \\ \gamma(\tau_j - \tau_i) & \text{for } t = \tau_j \end{cases} \tag{21}$$

Given that $\gamma \geq \beta$, the minimum is obtained at $t^* = \min(\tau_i, \tau_j)$. Filling this into Equation (6) yields $\tilde{C}(i,j) = 2\alpha \max(x_i - y_j, 0) + \beta|\tau_j - \tau_i|$. $\square$

Theorem 5.2: *Consider quasi-convex scheduling delay with $\gamma \geq \beta$. Driver i is matched to passenger j with desired arrival times τ_i and τ_j and $\tau_i \leq \tau_j$, without loss of generality. The optimal arrival time t^* that minimizes the total earliness and lateness is in the closed bounded interval $[\tau_i, \tau_j]$.*

Proof. Let the total earliness and total lateness for arrival time t be denoted by $E(t)$ and $L(t)$ respectively. Given the definition of earliness and lateness, the following inequalities hold:

$$E(t) \leq E(t') \text{ for } t \geq t',$$
$$L(t) \geq L(t') \text{ for } t \geq t'.$$

Furthermore, given the truncated nature of earliness and lateness, the following equalities hold:

$$E(t) = E(t') \text{ for } t, t' \leq \min(\tau_i, \tau_j),$$
$$L(t) = L(t') \text{ for } t, t' \geq \max(\tau_i, \tau_j).$$

Combining this, we observe that $E(t) + L(t) \geq E(t^*) + L(t^*)$ for $\min(\tau_i, \tau_j) \leq t^* \leq \max(\tau_i, \tau_j)$. $\square$

References

Agatz, N., Erera, A., Savelsbergh, M., and Wang, X. (2012). Optimization for dynamic ride-sharing: A review. *European Journal of Operational Research*, 223(2):295–303.

Alonso-Mora, J., Samaranayake, S., Wallar, A., Frazzoli, E., and Rus, D. (2017). On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment. *Proceedings of the National Academy of Sciences*, 114(3):462–467.

Arnott, R., de Palma, A., and Lindsey, R. (1987). *Schedule delay and departure time decisions with heterogeneous commuters*. University of Alberta, Department of Economics.

Arnott, R., de Palma, A., and Lindsey, R. (1990). Economics of a bottleneck. *Journal of urban economics*, 27(1):111–130.

Arnott, R., de Palma, A., and Lindsey, R. (1993). A structural model of peak-period congestion: A traffic bottleneck with elastic demand. *The American Economic Review*, pages 161–179.

Bagwell, K., Ramey, G., et al. (1992). *The diamond paradox: a dynamic resolution*. Department of Economics, University of California.

Bahat, O. and Bekhor, S. (2016). Incorporating ridesharing in the static traffic assignment model. *Networks and Spatial Economics*, 16(4):1125–1149.

Bai, R., Li, J., Atkin, J. A., and Kendall, G. (2014). A novel approach to independent taxi scheduling problem based on stable matching. *Journal of the Operational Research Society*, 65(10):1501–1510.

Bardhi, F. and Eckhardt, G. M. (2012). Access-based consumption: The case of car sharing. *Journal of Consumer Research*, 39(4):881–898.

Bento, A. M., Hughes, J. E., and Kaffine, D. (2013). Carpooling and driver responses to fuel price changes: Evidence from traffic flows in los angeles. *Journal of Urban Economics*, 77:41–56.

Beojone, C. V. and Geroliminis, N. (2020). On the inefficiency of ride-sourcing services towards urban congestion. *arXiv preprint arXiv:2007.00980*.

BlaBlaCar (2020). www.blablacar.com. `https://blog.blablacar.com/about-us`.

Boyacı, B., Zografos, K. G., and Geroliminis, N. (2015). An optimization framework for the development of efficient one-way car-sharing systems. *European Journal of Operational Research*, 240(3):718–733.

Browning, M., Chiappori, P.-A., and Weiss, Y. (2014). *Economics of the Family*. Cambridge University Press.

Burkard, R. E. and Cela, E. (1999). Linear assignment problems and extensions. In *Handbook of combinatorial optimization*, pages 75–149. Springer.

Caulfield, B. (2009). Estimating the environmental benefits of ride-sharing: A case study of Dublin. *Transportation Research Part D: Transport and Environment*, 14(7):527–531.

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

Chiappori, P.-A., Oreffice, S., and Quintana-Domeque, C. (2012). Fatter attraction: anthropometric and socioeconomic matching on the marriage market. *Journal of Political Economy*, 120(4):659–695.

de Palma, A., Lefèvre, C., and Ben-Akiva, M. (1987). A dynamic model of peak period traffic flows and delays in a corridor. *Computers & mathematics with applications*, 14(3):201–223.

de Palma, A., Lindsey, R., and Picard, N. (2015). Trip-timing decisions and congestion with household scheduling preferences. *Economics of Transportation*, 4(1-2):118–131.

de Palma, A., Picard, N., and Waddell, P. (2007). Discrete choice models with capacity constraints: An empirical analysis of the housing market of the greater paris region. *Journal of Urban Economics*, 62(2):204–230.

DeMaio, P. (2009). Bike-sharing: History, impacts, models of provision, and future. *Journal of Public Transportation*, 12(4):3.

Duranton, G. and Puga, D. (2004). Micro-foundations of urban agglomeration economies. In *Handbook of regional and urban economics*, volume 4, pages 2063–2117. Elsevier.

Eurostat (2020). Eurostat. `https://ec.europa.eu/eurostat/statistics-explained/index.php/Passenger_cars_in_the_EU`.

Fu, M., Kelly, J. A., Clinch, J. P., and King, F. (2012). Environmental policy implications of working from home: Modelling the impacts of land-use, infrastructure and socio-demographics. *Energy policy*, 47:416–423.

Fushimi, M. (1981). The secretary problem in a competitive situation. *Journal of the Operations Research Society of Japan*, 24(4):350–359.

Galichon, A. (2018). *Optimal transport methods in economics*. Princeton University Press.

Gurumurthy, K. M., Kockelman, K. M., and Simoni, M. D. (2019). Benefits and costs of ride-sharing in shared automated vehicles across Austin, Texas: opportunities for congestion pricing. *Transportation Research Record*, 2673(6):548–556.

He, L., Mak, H.-Y., Rong, Y., and Shen, Z.-J. M. (2017). Service region design for urban electric vehicle sharing systems. *Manufacturing & Service Operations Management*, 19(2):309–327.

Iryo, T. and Yoshii, T. (2007). Equivalent optimization problem for finding equilibrium in the bottleneck model with departure time choices. In *4th IMA International Conference on Mathematics in TransportInstitute of Mathematics and its Applications*.

Jittrapirom, P., Caiati, V., Feneri, A.-M., Ebrahimigharehbaghi, S., Alonso González, M. J., and Narayan, J. (2017). Mobility as a service: A critical review of definitions, assessments of schemes, and key challenges.

Li, Z., Hong, Y., and Zhang, Z. (2016). Do ride-sharing services affect traffic congestion? an empirical study of uber entry. *SSRN Electronic Journal*, (2002):1–29.

Li, Z.-C., Huang, H.-J., and Yang, H. (2020). Fifty years of the bottleneck model: A bibliometric review and future research directions. *Transportation Research Part B: Methodological*, 139:311–342.

Lindsey, R., de Palma, A., and Silva, H. E. (2019). Equilibrium in a dynamic model of congestion with large and small users. *Transportation Research Part B: Methodological*, 124:82–107.

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

Özkan, E. and Ward, A. R. (2020). Dynamic matching for real-time ride sharing. *Stochastic Systems*, 10(1):29–70.

Qian, Z. S. and Zhang, H. M. (2011). Modeling multi-modal morning commute in a one-to-one corridor network. *Transportation Research Part C: Emerging Technologies*, 19(2):254–269.

Riggs, W. (2020). Telework and sustainable travel during the covid-19 era. *Available at SSRN 3638885*.

Santi, P., Resta, G., Szell, M., Sobolevsky, S., Strogatz, S. H., and Ratti, C. (2014). Quantifying the benefits of vehicle pooling with shareability networks. *Proceedings of the National Academy of Sciences*, 111(37):13290–13294.

Shaheen, S. A., Mallery, M. A., and Kingsley, K. J. (2012). Personal vehicle sharing services in North America. *Research in Transportation Business & Management*, 3:71–81.

Small, K. A., Verhoef, E. T., and Lindsey, R. (2007). *The economics of urban transportation*. Routledge.

Smith, M. (1993). A new dynamic traffic model and the existence and calculation of dynamic user equilibria on congested capacity-constrained road networks. *Transportation Research Part B: Methodological*, 27(1):49–63.

Stoll, M. A. (1999). Spatial job search, spatial mismatch, and the employment and wages of racial and ethnic groups in los angeles. *Journal of Urban Economics*, 46(1):129–155.

Sullivan, C. and O'Fallon, C. (2010). Kilometres travelled and vehicle occupancy in urban areas: Improving evaluation and monitoring. *New Zealand Transport Agency Research Report*, (399).

Vickrey, W. S. (1969). Congestion theory and transport investment. *The American Economic Review*, 59(2):251–260.

Wang, X., Agatz, N., and Erera, A. (2018). Stable matching for dynamic ride-sharing systems. *Transportation Science*, 52(4):850–867.

Yang, H., Shao, C., Wang, H., and Ye, J. (2020). Integrated reward scheme and surge pricing in a ridesourcing market. *Transportation Research Part B: Methodological*, 134:126–142.

Zenou, Y. (2009). *Urban labor economics*. Cambridge University Press.

Zha, L., Yin, Y., and Xu, Z. (2018). Geometric matching and spatial pricing in ride-sourcing markets. *Transportation Research Part C: Emerging Technologies*, 92:58–75.