

Fair and Efficient Ride-Scheduling: A Preference-Driven Approach

Yi Cheng Ong, Nicos Protopapas, Vahid Yazdanpanah, Enrico H. Gerding, Sebastian Stein

University of Southampton, Southampton, UK

ARTICLE HISTORY

Compiled March 18, 2024

ABSTRACT

Smart mobility and, in particular, automated ridesharing platforms promise efficient, safe, and sustainable modes of transportation in urban settings. To make such platforms acceptable to the end-users, it is key to capture their preferences not in a static manner (by determining a fixed route and schedule for the vehicle) but in a dynamic manner by giving the riders the chance to get involved in the routing process of an upcoming journey. To that end, this work provides a toolbox of multiagent methods that enable different forms of active preference-awareness in ridesharing services. We capture riders' preferences (as end-users of a ridesharing service), preserve their privacy by avoiding expecting them to share preferences with other riders, and show the efficacy of the presented ridesharing algorithms using agent-based simulation and illustrating their utilitarian and fairness properties.

KEYWORDS

Smart Mobility; Multiagent Ridesharing Methods; Preferences; Agent-Based Modelling; AI for Social Good.

1. Introduction

Ridesharing is a promising means towards reducing carbon emissions and mitigating climate change effects (Furuhata et al., 2013). Integrating preference-awareness into ridesharing systems enhances the satisfaction of the end-users and, in turn, allows efficiently using the spare capacity in urban transportation services (Rigas, Ramchurn, & Bassiliades, 2014). However, although the potential gains are known, some traditional urban transportation systems—such as bus services—are not benefiting from ridesharing technologies. Our buses are still operating based on fixed schedules and, in the best case, use historical data on the behaviour of riders to improve their routes. However, determining routes based on data about *past* users does not necessarily fit how *present* riders want to use the service. For example, arguably, bus schedules that were generated based on travellers' behaviour in 2019 (i.e., before the COVID-19 pandemic) do not satisfy what riders want in 2024. While gathering data more frequently and then fixing a static schedule is an option, in this work, we go a step further and suggest dynamic schedules that are determined based on the preferences of the *current* users/riders. Doing this will allow buses to provide customised services to riders, avoid wasting resources by meeting unnecessary stations, and find compromises for pickup and drop-off locations that users see preferable.

CONTACT Nicos Protopapas. Email: n.protopapas@athenarc.gr. Nicos Protopapas is now with Archimedes/Athena RC in Athens, Greece.

Ridesharing can be considered a vehicle routing problem that is typically solved by optimising a given global objective function. Most studies focused on optimising operational-based objective functions (Mourad, Puchinger, & Chu, 2019), which usually benefit the ridesharing service provider rather than the passengers. Optimising based on the provider’s incentives would not lead to widespread adoption of this service, as passengers’ preferences are not taken into consideration. Thus, more recently, there have been studies regarding incorporating preferences or incentives of passengers into the ridesharing problem. These studies specified constraints, such as passengers’ maximum travel distance and maximum waiting time, to ensure some level of quality service, as well as incorporating new terms in the objective function, which encapsulates the overall satisfaction of passengers (Mourad et al., 2019). However, this method does not guarantee fairness among the passengers. For example, a ridesharing system might choose route *A* because it minimises the total waiting time and travelling time of all passengers. However, route *A* might lead to a longer travel time for a subset of passengers; for the sake of minimising the overall objective function, their preferences have been neglected, such that route *A* results in more inconvenience to them. In principle, assuming that the ideal route needs to be optimal merely based on the characteristics of the city in an objective sense (e.g., in terms of distances), ignores how satisfied riders are in a subjective sense. In such a view, ridesharing is approached and accordingly solved as a merely *technical* problem with no intention to take into account the *social* and preferential dimensions. In view of human-centred AI techniques and the need for developing trustworthy human-AI partnerships (Ramchurn, Stein, & Jennings, 2021), we see ridesharing as an inherently sociotechnical problem and argue that the acceptance of it by society depends on its ability to capture riders’ preferences throughout the journey.

In the line of work on the application of multiagent techniques to ridesharing and the so-called autonomous mobility-on-demand systems (Zhang, Spieser, Frazzoli, & Pavone, 2015), there are attempts to optimise the efficiency of using the shared space (Agatz, Erera, Savelsbergh, & Wang, 2012), to use ridesharing for transportation of goods and passengers (van der Tholen, Beirigo, Jovanova, & Schulte, 2021), and to analyse the environmental benefits of ridesharing (Jalali, Koohi-Fayegh, El-Khatib, Hoornweg, & Li, 2017). In comparison to techniques that build on a snapshot of the system with resource optimisation in mind (e.g., Agatz et al. (2012); Jalali et al. (2017)) and approaches on mixed rider-goods transportation settings (e.g., van der Tholen et al. (2021)), our focus is on passenger transportation and the use of ridesharing dynamically by capturing riders’ preferences. Here, it is critical to capture the preferences of passengers for a journey, and in particular, their satisfaction concerning their asserted urgency and ideal pickup/drop-off times. Moreover, it is unrealistic to assume complete knowledge about the preferences of riders (which would allow standard optimisation techniques to be applied). Instead, riders may have reservations about sharing their preferences and these may also change given additional knowledge they gain during the generation of routes. To address these challenges, our work is the first attempt to develop a safe framework (as we do not expect riders to share information between them) for preference-aware dynamic ridesharing. Moreover, in contrast to recent work that applies voting mechanisms to decide about destinations and accordingly form *like-minded coalitions* of riders for a journey (e.g., in Dennisen and Müller (2016)), we focus on a given set of users (e.g., employees of a company) and work on the satisfaction of users by allowing them to participate *in the planning*

of the journey.¹

Against this background, this is the first contribution that develops algorithms for determining ridesharing routes in participation with riders, allows dynamic routing by integrating voting mechanisms, and relaxes the expectation that riders need to compromise their privacy by sharing information with other riders.

On an abstract level, our approach integrates voting mechanisms into ridesharing to capture preferences. Note that the idea is not to ask all the (potential) riders to interact with their ridesharing application, e.g., on their personal mobile devices, and vote. In contrast, our approach allows providing minimal information to the service provider (via such an application) and then the voting procedure takes place in the background. To that end, we merely expect users to express their station of origin, their destination, and desired pick-up and drop-off times. Users share such information with the service (as a trusted agent), which respects users' privacy by avoiding expecting them to share such sensitive data with other users.

To that end, we take the following principles following the literature on ridesharing (Aïvodji, Gambis, Huguet, & Killijian, 2016; Hasan, Van Hentenryck, Budak, Chen, & Chaudhry, 2018) and trust in smart mobility (Zavvos et al., 2022):

- *Participatory*: in the sense that users can influence the operation of the service with respect to their pickup/drop-off preferences.
- *Generic*: to have a framework that is not fixed to a specific setting, e.g., to a particular urban area, but is aware of key concerns in the urban transportation context.
- *Modular*: in the sense that we can integrate different forms of voting mechanisms and utility evaluation functions.
- *Safe*: in the sense that riders do not need to share information between them, but only with a trusted ridesharing agent.

Contributions. To encourage more citizens to utilise ridesharing services, these services should not neglect passengers' preferences (Stein & Yazdanpanah, 2023) and should cope with the high computational complexity of finding an optimal (or close to optimal) solution. Our contribution revolves around solving these two problems within a static ridesharing context. We propose and experimentally analyse two families of algorithms to construct a common schedule for various riders using the same bus. Our first solution is grounded on a randomised greedy approach and allows the riders to build a schedule in a round-robin fashion collectively: the riders are visited in a random order and each rider expands the current solution to accommodate their needs. In the second family of algorithms, the schedule is expanded by voting on which partial solution should be inserted. Our proposed algorithms are evaluated using simulations for various spatial, temporal and preferential scenarios.²

Roadmap. The remainder of the work is organised as follows: in Section 2, we provide a formal modelling of the ridesharing problem using a graph-theoretical representation and show how preferences can be modelled using a utility function. In Section 3, different preference-aware ridesharing algorithms, the rationale behind each and how they capture aspects of the problem are presented. We evaluate the presented algorithms

¹A detailed discussion on how this work relates to different lines of research on ridesharing appears in Section 6.

²The initial idea behind this work is presented at the 6th International Workshop on Agent-Based Modelling of Urban Systems (Ong, Protopapas, Yazdanpanah, Gerding, & Stein, 2022).

Symbols	Description
r_i	Rider i
$l_d(r_i)$	Departure location for r_i
$l_a(r_i)$	Arrival location for r_i
$T_d^*(r_i)$	Preferred departure time for r_i
$T_a^*(r_i)$	Preferred arrival time for r_i
$T_d(r_i)$	Allocated departure time for r_i
$T_a(r_i)$	Allocated arrival time for r_i
$s(r_i)$	Status for r_i , which can be <i>onboard</i> or <i>onground</i>
β_{r_i}	Patience parameter for r_i
$G(L, E)$	Fully connected graph with locations L and edges E
$l_i \in L$	Location i
$e(l_i, l_j) \in E$	Edge between l_i and l_j
$w(l_i, l_j)$	Travel Time between station l_i and station l_j
\mathcal{S}	Sequence of TourNodes
$Pos(\mathcal{S})$	Insertion positions of \mathcal{S}
s_i	TourNode at the i -th position of \mathcal{S}
l_{s_i}	Location at s_i
$T_{s_i}^a$	Arrival time at s_i
$T_{s_i}^w$	Waiting time at s_i
$T_{s_i}^d$	Departure time at s_i
$R_{s_i}^P$	Set of riders picked up at s_i
$R_{s_i}^D$	Set of riders dropped off at s_i

Table 1. Overview of the model’s formal notation.

in Section 4 via simulation-based experiments and use real-life datasets from dense neighbourhoods to compare the efficacy of our suggested ridesharing algorithms.³ We discuss links to related work in Section 6 and provide concluding remarks in Section 7.

2. Formal Model

We consider the case of generating a route for a single bus that is part of a 24-hour ridesharing service. The route generated will be the order of stations to visit by taking into account the temporal preferences of the current riders. In the remainder of this section, we first define the map as a graph, and then we specify the riders in the system. Following this, we formalise how we represent a solution, i.e. the bus route.

Table 1 provides an overview of the notation.

2.1. Graph

The service region is represented by a fully connected graph $G(L, E)$, where L is the set of stations and E is the set of edges connecting any two stations. Each edge $(l_i, l_j) \in E$ has a corresponding weight $w(l_i, l_j)$ which is the shortest travel time between the two locations l_i and l_j .

³For this, we used the publicly available Great Britain’s national dataset of public transport access points (bus stops, rail stations, airports, ferry piers, tram/metro/underground stops) (Department for Transport, 2014).

2.2. Riders

The input also consists of a set of riders R such that each $r_i \in R$ has a departure station $l_d(r_i) \in L$ and an arrival station $l_a(r_i) \in L$. Moreover, each rider has their most preferred departure and arrival time, represented by $T_d^*(r_i)$ and $T_a^*(r_i)$ respectively.

To evaluate the satisfaction level of riders, riders' preferences for different temporal allocations should be represented. It can be reasonably assumed that, for every rider, as their departure and arrival time deviation increases, the satisfaction level decreases. Thus, a decreasing utility function with respect to increased time deviation is appropriate. Additionally, we do not assume that each rider is homogeneous (i.e., each rider has a different sense of urgency and patience). Furthermore, we expect the utility of the impatient riders to decrease rapidly and the utility of patient riders to have a much slower decreasing rate. Taking these points into consideration, we, therefore, assume the following utility function:

$$U_{r_i}(T_d, T_a) = U_{r_i}^d(T_d) + U_{r_i}^a(T_a) = \frac{\beta_{r_i}^{|T_d^*(r_i) - T_d(r_i)|} + \beta_{r_i}^{|T_a^*(r_i) - T_a(r_i)|}}{2} \quad (1)$$

The functions $U_{r_i}^d(\cdot)$ and $U_{r_i}^a(\cdot)$ measure the respective utilities obtained by rider r_i as he or she departs at time $T_d(r_i)$ and arrives at time $T_a(r_i)$. The total utility, measured by $U_{r_i}(T_d, T_a)$, is simply the sum of the two respective utilities. Value β_{r_i} is a scalar within the range $[0, 1]$ that captures the patience of rider r_i . Higher β values indicate greater patience. The distance $|T_d^*(r_i) - T_d(r_i)|$ is the time deviation from r_i 's most preferred departure time, and $|T_a^*(r_i) - T_a(r_i)|$ the time deviation from r_i 's most preferred arrival time. Since the β values are in the range $[0, 1]$, any increase in time deviation will cause a decrease in utility. Additionally, the rate of such decrease is dependent on the β value; the lower the β value, the greater the decrease rate. Finally, the utility obtained for each rider r_i will always be in the range $[0, 1]$.

To illustrate the working of the utility function, consider the following example: There are two riders r_1 and r_2 with the following beta values:

$$\beta_{r_1} = 0.95$$

$$\beta_{r_2} = 0.85$$

The two riders also have the same departure and arrival time deviation:

$$|T^*(r_1) - T_d(r_1)| = |T^*(r_2) - T_d(r_2)| = 20 \text{ minutes}$$

$$|T^*(r_1) - T_a(r_1)| = |T^*(r_2) - T_a(r_2)| = 10 \text{ minutes}$$

The resulting utilities obtained are:

$$U_{r_1} = \frac{0.95^{20} + 0.95^{10}}{2} \approx 0.48$$

$$U_{r_2} = \frac{0.85^{20} + 0.85^{10}}{2} \approx 0.12$$

We can see that rider r_2 has a lower utility than that of rider r_1 although both departure and arrival time deviations are the same. Essentially, time deviations will cause a greater decrease in utility for riders with lower patience (lower β value).

To generate a β_{r_i} value for each rider, we utilise the *beta distribution* (Gupta & Nadarajah, 2004), which is versatile enough to cover for multiple scenarios (see Section 4 for details). For clarification, “ β_{r_i} ” refers to the value that represents the patience value of r_i , while “*beta*” refers to the beta distribution used to generate the patience values of all riders R . We chose the beta distribution as it allowed us to sample values from the $[0, 1]$ range, and to simulate various interesting scenarios. Note that: (i) when $\beta_{r_i} = 0$, the utility function for r_i will be undefined. In this scenario, we set the utility value to be 0, indicating that r_i is indifferent to any time deviation, in the sense that they are incompletely satisfied with any time deviation; (ii) when $\beta_{r_i} = 1$, the utility of r_i is 1 given any time deviation. This also indicates that the rider is indifferent, but in the sense that r_i is perfectly satisfied with any time deviation; 3. If $\beta_{r_i} \in (0, 1)$ and the time deviations for both departure and arrival times are 0 minutes, the resulting utility will be 1 (the rider gets the maximum utility when the ideal pick up and drop off times are met).

2.3. Solution

We define the solution \mathcal{S} as a sequence of *TourNodes* where each TourNode s_i is represented by a tuple:

$$(l_{s_i}, T_{s_i}^a, T_{s_i}^w, R_{s_i}^P, R_{s_i}^D)$$

where l_{s_i} is the location at the i -th position of \mathcal{S} , $T_{s_i}^a$ is the arrival time at l_{s_i} , $T_{s_i}^w$ the waiting time duration at l_{s_i} , $R_{s_i}^P$ the set of riders to be picked-up at s_i and $R_{s_i}^D$ the set of riders to be dropped-off at s_i .

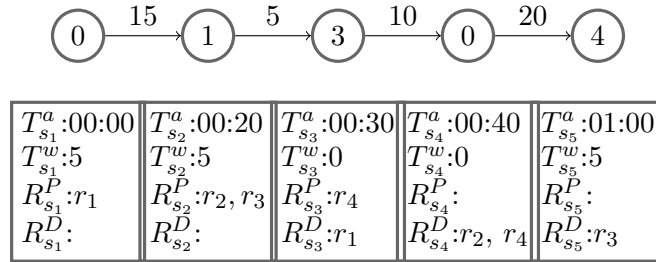


Figure 1. Solution Model.

Figure 1 illustrates an example of the solution model. The numbering on each node represents a station’s id. Note that there can be repeated locations in the solution. The edges represent travel time in minutes, calculated by taking the distance between the stations, and dividing it by the average bus speed. There are constraints imposed on the solution: (1) The arrival time of any node $T_{s_i}^a$ equals the arrival time of the previous node $T_{s_{i-1}}^a$ plus the waiting time of the previous node $T_{s_{i-1}}^w$, plus the travel

time between the stations $w(l_{s_{i-1}}, l_{s_i})$; (2) Riders can only be picked up and dropped off at their desired stations.

As depicted in Figure 1, starting from station 0, the bus waits for 5 minutes to pick up rider r_1 before departing to station 1. Once at station 1, the bus again waits for 5 minutes to pick up riders r_2 and r_3 before leaving for station 3. This continues until all riders have been dropped off. Note that, in this work, we are assuming that at any point in the journey, seat capacity always exceeds the number of riders. The bus will not accept any pick-up requests at any time range $[T_{s_i}^a, T_{s_i}^w]$ that has no more available seats.

The full list of model variables and their descriptions are summarised in Table 1.

3. Preference-Aware Ridesharing Algorithms

In this section, we present preference-aware ridesharing algorithms and discuss how they capture various aspects of the problem. The presented algorithms in this section will be evaluated in Section 4 using simulation-based experiments.

3.1. Randomised Greedy Algorithm

Our first approach, which we will also use as a benchmark for the following algorithms, is the Randomised Greedy Algorithm (RGA). The algorithm begins with an empty solution \mathcal{S} and a shuffled list of riders R . Then, for each rider $r_i \in R$, we find the best positions in \mathcal{S} to insert their departure and arrival stations to maximize the utility of rider r_i . When an existing TourNode $s_k \in \mathcal{S}$ better satisfies r_i than any valid, new TourNode, then r_i is included in $R_{s_k}^P$ or $R_{s_k}^D$ (depending if we are allocating arrival times or departure stations). In other words, we find positions (or existing TourNodes) in \mathcal{S} such that the utility of r_i is maximised given the current state of \mathcal{S} . The algorithm finishes when all riders are allocated to two TourNodes in \mathcal{S} : one for departure and one for arrival. The pseudocode for this algorithm is given in Algorithm 1.

As shown in Algorithm 1, we first shuffle the list of riders R . Then, for each rider r_i (examined according to the shuffled list), the **ValidTourNodes** function (at lines 4 and 5) iterates over all insertion positions $k \in Pos(\mathcal{S})$ and creates a TourNode s_k based on the rider's preferred departure/arrival time/stations. Additionally, **ValidTourNodes** also iterates over all existing TourNodes $s_i \in \mathcal{S}$ to find any s_i that has the desired departure and arrival stations. Then we have two sets, S_d and S_a , corresponding to the set of valid departure and arrival TourNodes. The **ValidTourNodes** function also enforces the following two constraints: When a new TourNode s_k is created at position k of \mathcal{S} , then:

- (1) the location l_{s_k} must not be in $\{l_{s_{k-1}}, l_{s_{k+1}}\}$. In other words, adjacent stations in \mathcal{S} must be different.
- (2) The resulting departure time of s_{k+1} should not exceed its original departure time.

The first constraint ensures that the final solution \mathcal{S} is a valid route (adjacent stations within the route should always be different). The second constraint was introduced as a means to reduce the time complexity of the algorithm, as well as ensuring that the resulting utilities remain stable throughout the insertion process. Consider that we have an existing solution \mathcal{S} with a large number of TourNodes and that the

Algorithm 1: Randomised Greedy Algorithm

input: Riders R

```
1  $S \leftarrow []$ ;
2  $R \leftarrow \text{Shuffle}(R)$  ;
3 foreach  $r_i \in R$  do
4    $S_d \leftarrow \text{ValidTourNodes}(T_d^*(r_i), l_d(r_i), S)$ ;
5    $S_a \leftarrow \text{ValidTourNodes}(T_a^*(r_i), l_a(r_i), S)$ ;
6    $s_d^* \leftarrow \arg \max_{s_d \in S_d} (U_{r_i}^d(T_{s_d}^d))$ 
7    $s_a^* \leftarrow \arg \max_{s_a \in S_a} (U_{r_i}(T_d(r_i), T_{s_a}^a))$  ;
8   if  $s_d^* \notin S$  then
9      $S \leftarrow \text{Insert}(s_d^*, S)$ ;
10  else
11     $R_{s_d^*}^P \leftarrow R_{s_d^*}^P \cup r_i$ 
12  end
13  if  $s_a^* \notin S$  then
14     $S \leftarrow \text{Insert}(s_a^*, S)$ ;
15  else
16     $R_{s_a^*}^D \leftarrow R_{s_a^*}^D \cup r_i$ 
17  end
18 end
```

best position for rider r_1 is position 10. If the second constraint is not enforced, inserting the new TourNode s_{10} would cause cascading updates on the departure, arrival and waiting times on the subsequent TourNodes after position s_{10} . Not only is this computationally expensive, but the utility for each rider would also fluctuate each time a new station is inserted. Finally, once we have the best departure and arrival tour node s_d^*, s_a^* , we either insert them into S if it does not exist in S , or add r_i into $R_{s_d^*}^P$ and $R_{s_a^*}^D$.

3.2. Fairness Algorithms

The Randomised Greedy Algorithm, although taking preferences into account, does not take into account any fairness aspects. Indeed, the riders that would benefit the most are those that get an allocation first; the remaining riders generally have less flexibility for their allocation. To encourage the use of such a system, some fairness policies need to be introduced. As such, we propose two algorithms which aim to provide fair solutions: Randomised Greedy++ (RGA++) and Iterative Voting (IV). Both of these algorithms utilise the ValidTourNodes function in Algorithm 1 in order to find (or create) the best s_d^* and s_a^* .

The first of these algorithms, Randomised Greedy++ is a variant of Randomised Greedy with just a simple change; for each rider, instead of allocating both departure and arrival times at the same time, we first only allocate a departure time for them. Next, we allocate the arrival times for each rider *in the reverse order*. This form of allocation is also known as *picking sequences* (see, e.g., Brandt, Conitzer, Endriss, Lang, and Procaccia (2016)). The main idea is that the earlier riders were allocated a departure time, the later they are allocated an arrival time. The algorithm is shown

Algorithm 2: Randomised Greedy ++

input: Riders R
 1 $S \leftarrow []$;
 2 $R \leftarrow \text{Shuffle}(R)$;
 3 **foreach** $r_i \in R$ **do**
 4 $S_d \leftarrow \text{ValidTourNodes}(T_d^*(r_i), l_d(r_i), S)$;
 5 $s_d^* \leftarrow \arg \max_{s_d \in S_d} (U_{r_i}^d(T_{s_d}^d))$;
 6 **if** $s_d^* \notin S$ **then**
 7 $S \leftarrow \text{Insert}(s_d^*, S)$;
 8 **else**
 9 $R_{s_d^*}^P \leftarrow R_{s_d^*}^P \cup r_i$
 10 **end**
 11 **end**
 12 $R \leftarrow \text{Reverse}(R)$;
 13 **foreach** $r_i \in R$ **do**
 14 $S_a \leftarrow \text{ValidTourNodes}(T_a^*(r_i), l_a(r_i), S)$;
 15 $s_a^* \leftarrow \arg \max_{s_a \in S_a} (U_{r_i}(T_a(r_i), T_{s_a}^a))$;
 16 **if** $s_a^* \notin S$ **then**
 17 $S \leftarrow \text{Insert}(s_a^*, S)$;
 18 **else**
 19 $R_{s_a^*}^D \leftarrow R_{s_a^*}^D \cup r_i$
 20 **end**
 21 **end**

in Algorithm 2.

Iterative Voting follows the previous idea of building a schedule iteratively but also allows for more interaction between the riders and the schedule-creating procedure. The algorithm works as follows: While there are still unallocated riders⁴, each unallocated rider proposes a TourNode (either of departure or arrival, depending on the rider's status) which encapsulates the desired station, the pick up time, drop-off time and waiting time (see lines 9 to 16 in Algorithm 3). Once the proposal phase ends, the riders rank their candidates according to their utility functions. The rankings are then passed on a voting rule, which selects a single TourNode as a winner. The process repeats until all riders are allocated. We should note here that all the riders' side actions can be implemented using a smart agent, which knows riders' temporal and spatial preferences and a strategy to rank TourNodes.

Iterative Voting is described in detail in Algorithm 3. The two constraints imposed by the ValidTourNodes function are also imposed in this algorithm, too. Hence, each proposing rider selects its proposal by a pool of available TourNodes S_d or S_a . Once the list of candidates is finalized, all riders need to rank all available TourNodes. To do this, we can compute the utility of the riders for arrival and departure TourNodes. This is straightforward when the departure or arrival locations of a rider, say rider r_i , coincide with the TourNode's s_k location (i.e. when either $l_d(r_i) = l_{s_k}$ or $l_a(r_i) = l_{s_k}$): the utility of rider r_i for s_k is computed as $U_{r_i}^d(T_{s_k}^d)$ or $U_{r_i}^a(T_{s_k}^a)$. Special treatment is,

⁴in the sense that not all riders have been allocated a departure and arrival time.

however, required when the riders need to rank TourNodes whose location is neither a rider’s departure nor arrival location. While the algorithm can handle many scenarios – the only requirement is that the riders should rank the TourNodes – in Section 4, we assume that each rider ranks candidate TourNodes according to their distance from their arrival or departure locations (depending on the rider’s status).

In Section 4, we evaluate our mechanism using some well-known voting rules. We start with the arguably more commonly used voting rule of *popularity*⁵ (IV-popularity): Given all rankings, the top choices of all riders are counted, and the most popular one is selected. One can argue that popularity ignores much information by just focusing on the top choice of each rider. To examine this, we consider two other rules: First, the *Borda* rule (IV-Borda), which assigns points to each candidate according to their position in the ranking. More specifically, given a ranking over m candidates, the candidate at the i -th position gets a score of $m - i$ points: hence, a candidate in the top position is assigned m points, a candidate in the second position is assigned $m - 1$, and a candidate in the last position gets 1 point. The candidate accumulating the highest score is the winner of the election. A similar rule is the *harmonic* rule (IV-harmonic), which assigns $1/i$ points to the candidate in the i -th position. Finally, we utilize the *instant runoff* (IV-IR)⁶ rule, a widely used voting procedure, which runs in phases: at each phase, the candidates are ranked according to their popularity scores, and the least favourable candidate is eliminated. The process continues until a single candidate survives, which is the winner. We should note, however, that Iterative Voting is not grounded on a single voting rule but is general enough to handle any voting rule based on candidate rankings. We refer the reader to Chapter 2 of (Brandt et al., 2016) for more information on the popularity, Borda and instant runoff voting rules, and to (Boutilier et al., 2012) for the harmonic rule.

3.3. Safety

To ensure a safe framework for the riders, where no rider can know another rider’s route, our algorithms can be employed by a trusted central authority, e.g., a cloud service. In this architecture, the trusted agent will handle necessary communications, i.e., our algorithm does not necessitate agents to share private information with one another, hence no potential breach of privacy. In the case of the RGA algorithms, the riders will send their preferences (the patience parameter, desired departure and arrival times, and preferred departure and arrival locations) to the cloud service. Then, the algorithm will be implemented, and the riders will be informed about their departure and arrival times and locations. The IV algorithms require even less sensitive information as there is no need to communicate the preferences to the central authority. The riders only need to send their proposed TourNodes, and a ranking of the candidate TourNodes. When a TourNode is proposed, the central authority informs the rider whether the TourNode is valid; if it is not valid, the rider will provide a new proposal. After the end of the algorithm, all riders are informed of the final schedule. Observe that during the communication of the proposed TourNodes and the final schedule (in both families of algorithms), there is no need to communicate the lists of riders for pick-up and drop-off, which can be held only by the central authority.

⁵This is also known as *plurality* voting.

⁶Also known as the *Single Transferable Vote* rule.

Algorithm 3: Iterative Voting

input: Riders R , Voting Mechanism V

```
1  $S \leftarrow []$ ;
2 foreach  $r_i \in R$  do
3    $s(r_i) = \text{onground}$ 
4 end
5  $\text{Unallocated} \leftarrow R$ ;
6 while  $\text{Unallocated} \neq \emptyset$ :
7    $\text{Candidates} \leftarrow []$ ;
8   foreach  $r_i \in R$  do
9     if  $s(r_i) = \text{onground}$  then
10       $S_d \leftarrow \text{ValidTourNodes}(T_d^*(r_i), l_d(r_i), \mathcal{S})$ ;
11       $s_d^* \leftarrow \arg \max_{s_d \in S_d} (U_{r_i}^d(T_{s_d}^d))$ ;
12       $\text{Candidates} \leftarrow \text{Candidates} \cup s_d^*$ 
13    else if  $s(r_i) = \text{onboard}$  then
14       $S_a \leftarrow \text{ValidTourNodes}(T_a^*(r_i), l_a(r_i), \mathcal{S})$ ;
15       $s_a^* \leftarrow \arg \max_{s_a \in S_a} (U_{r_i}(T_d(r_i)), T_{s_a}^a)$ ;
16       $\text{Candidates} \leftarrow \text{Candidates} \cup s_a^*$ ;
17    end
18  end
19   $\text{Rankings} \leftarrow []$ ;
20  foreach  $r_i \in R$  do
21     $\text{ranking}_{r_i} \leftarrow \text{Rank}(U_{r_i}, \text{Candidates})$ 
22     $\text{Rankings} \leftarrow \text{Rankings} \cup \text{ranking}_{r_i}$ ;
23  end
24   $s_k \leftarrow \text{Vote}(\text{Rankings}, V)$ ;
25  if  $s_k \notin \mathcal{S}$  then
26     $\mathcal{S} \leftarrow \text{Insert}(s_k, \mathcal{S})$ ;
27  end
28  Update  $\text{Status}$  and  $\text{Unallocated}$ 
29 end
```

3.4. Applicability and Practical Challenges

We end this section with a brief discussion on how our methods can be implemented in practice. All of our algorithms can be implemented through web-based (smartphone) applications. The application will ask for the spatial and temporal preferences of the riders and will use a discrete choice experiment⁷ to estimate the riders' β value. For the RGA family of algorithms, the applicability is straightforward. For the IV family of algorithms, the application will act as a proxy between the rider and the algorithm. Given the rider's preferences, it automatically votes on her behalf by optimising over the list of available choices. A few occasional questions might be asked during the process to fine-tune the procedure. Hence, the cognitive burden to the rider is limited to answering preference-elicitation questions. This is a clear benefit compared to the

⁷see, e.g., König and Gripenkoven (2020) on how the method has been implemented in the context of ridesharing.

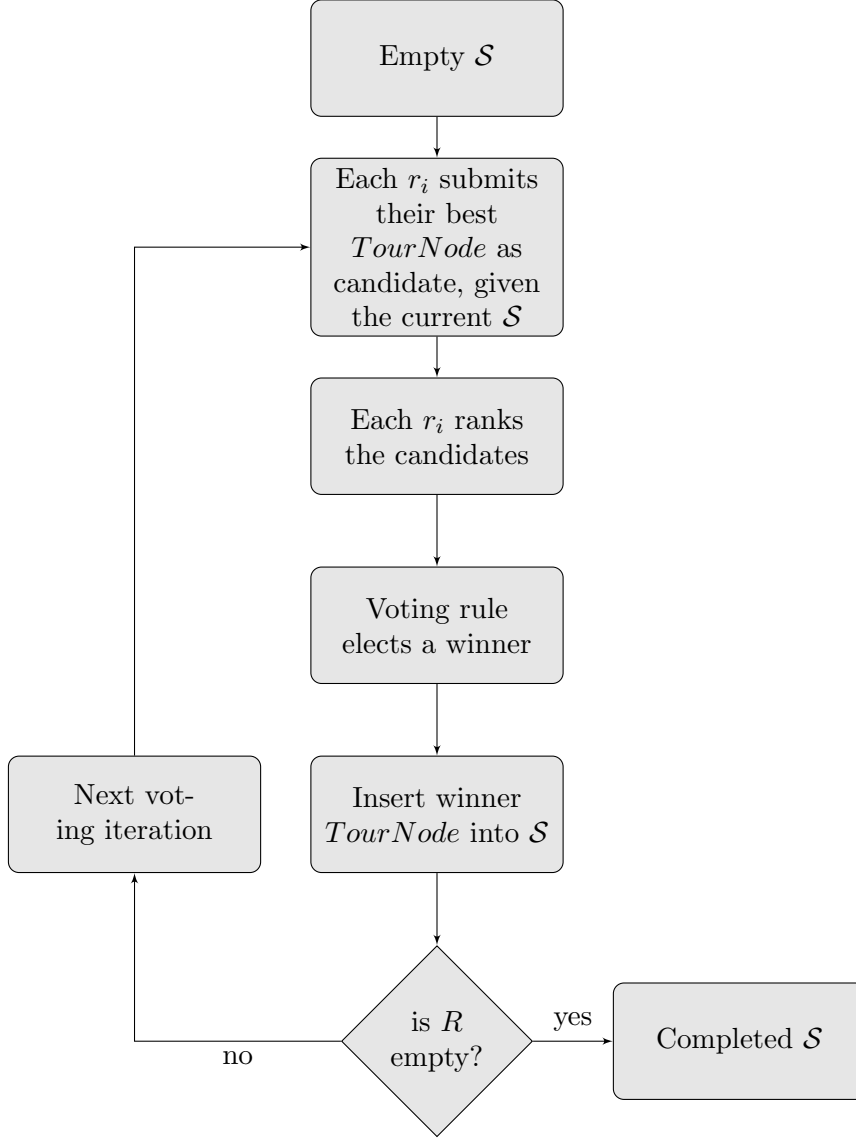


Figure 2. Iterative Voting Procedure.

fixed bus routes, which become extensively complex when needed to cover sizeable urban areas (Badia, Argote-Cabanero, & Daganzo, 2017).

The adoption of our approach is, of course, not coming without limitations. The most obvious one is its departure from the current norm of fixed bus schedules and our assumption of the social acceptability of such a technology. One possible way to implement our proposal is through a hybrid system, which will use the traditional approach of fixed routes to cover a stable and predictable demand and, in addition, smaller versatile buses using dynamic routes to cover unpredictable or less frequent demand. A careful implementation of this system will allow commuters to use the dynamically scheduled buses to connect easily between large transportation hubs or move towards areas with low or inconsistent demand, allowing for better adoption of the system.

4. Experimental Results

We conducted a range of experiments under different parameter settings in order to test the efficacy of the proposed algorithms in terms of their utilitarian and fairness properties under different real-life scenarios. We then suggest which type of algorithms would work best given the specific settings. The different parameter settings we consider for our simulation-based experiments include varying the number of riders and their similarity in terms of time sensitivity and station preferences. Each experiment is run 100 times to obtain the mean sum of utilities and mean Gini index, the metrics we use to evaluate our proposed algorithms, regarding efficiency and fairness. Next, we define various parameter settings used for our simulations.

Bus Service. Table 2 lists the bus service parameters used for our simulation experiments. To simulate an average performance of our algorithms, we assume that the buses travel with a constant average bus speed of 13 km/h, which has been chosen according to the “bus speed performance by boroughs” report from Transport for London (Transport of London, 2021). We evaluate our algorithms under perfect conditions where the buses are always on time to avoid distractions due to adverse effects like traffic.

Parameter	Value
Service Hours	24
Average Speed	13 km/h

Table 2. Bus Service Parameters.

Graph. We consider a high density geographical region where bus stations are near to each other. To this end, we utilised the Naptan dataset (Department for Transport, 2014), and obtained the latitude and longitudes of bus stations in Westminster, London. There are a total of 66 bus stations in this dataset. We consider the 66 bus stations as fully connected. We consider the shortest distance between each station, approximated using the great circle distance, calculated based on the *haversine* formula (Robusto, 1957).

Preference Generation. While generating preferences, it is assumed that riders want to get dropped off as early as possible and that they have knowledge of the travel time between the departure and arrival stations, such that their preferred arrival time is always greater than the sum of their departure time and travel time. This gives more meaning when reviewing the experimental results compared to having unrealistic preferences, which would result in very low utilities regardless of how the proposed algorithms perform. For the list of variable rider parameters, refer to Table 3.

Time-sensitivity. The alpha and beta parameters correspond to the parameters for the beta distribution (see, e.g., Gupta and Nadarajah (2004)). The beta distribution is quite versatile and allows us to model different scenarios regarding time-sensitivity. For our experiments, we use four specific alpha-beta pairs, to simulate four indicative cases: In the first case, we use $\alpha = 10$ and $\beta = 10$ to simulate a *balanced* population with mean patience factor equal to $1/2$, highly concentrated around the mean. The second case, with $\alpha = 10$ and $\beta = 2$, simulates a skewed, mostly *patient* (or *time-insensitive*) population, with a mean at $5/6$, while the case alpha

Parameter	Value
(alpha, beta)	$[(10, 10), (10, 2), (2, 10), (1/2, 1/2)]$
Spatial preferences	[uniform, 10 hotspots, 20 hotspots]
Number of Riders	[25, 50, 75, 100, 125, 150, 175, 200]

Table 3. Rider Parameters. The first row refers to the values of the Beta distribution; the second row describes the sampling method for the pick up and drop off locations; the third row shows the number of riders for each experiment.

$= 2$ and $\beta = 10$, simulates a symmetrical, mostly *impatient* (or *time-sensitive*) population with mean at $1/6$. We recall here that higher β values imply more patient riders. Finally, the case with $\alpha = 1/2$ and $\beta = 1/2$ is *bimodal*. While the mean of the distribution is $1/2$, the population is concentrated around the extreme values 0 and 1.

Temporal preferences. Given that this is a 24-hour service, riders’ time preferences are modelled as minutes in the interval $[0, 1440]$ with unit time steps. We consider peak-hours-based distribution when generating preferences. To encapsulate real-world demand, riders’ time preferences are first sliced into time frames as shown in Table 4.

Service Type	Time Frame
Midnight Rides	$[0, 420]$
Morning Peak	$[420, 560]$
Working Hours	$[560, 1020]$
Evening Peak	$[1020, 1140]$
Night Rides	$[1140, 1440]$

Table 4. Time Frames.

To simulate a higher percentage of riders during peak hours, we set a higher sampling probability for requests in peak time frames than non-peak time frames. We tested a range of probabilities $[0.6, 0.7, 0.8]$ (consequently, non-peak time frames have $[0.4, 0.3, 0.2]$ probability) and the results for the different experiments showed similar trends. As such, we use a peak-time probability of 0.8 for our simulations.

Spatial preferences. Lastly, the desired pick up and drop off stations V are sampled in two ways: First, by using *uniform sampling*, i.e., one location is sampled uniformly at random to act as a departure station, and another location is sampled, without replacement, to act as an arrival station. Our second approach uses a *hotspot based sampling*: in this scenario, specific locations are declared as hotspots during the morning and evening peak time windows, which get much higher demand than the other locations. In our experiments we select randomly a set of 10 or 20 locations from V . These locations act as morning and evening peak hotspots, and during these time windows, these locations are sampled with higher probability than the other locations. For simplicity, in our experiments, we assume that these are the only locations with positive demand during the morning and evening peak time windows.

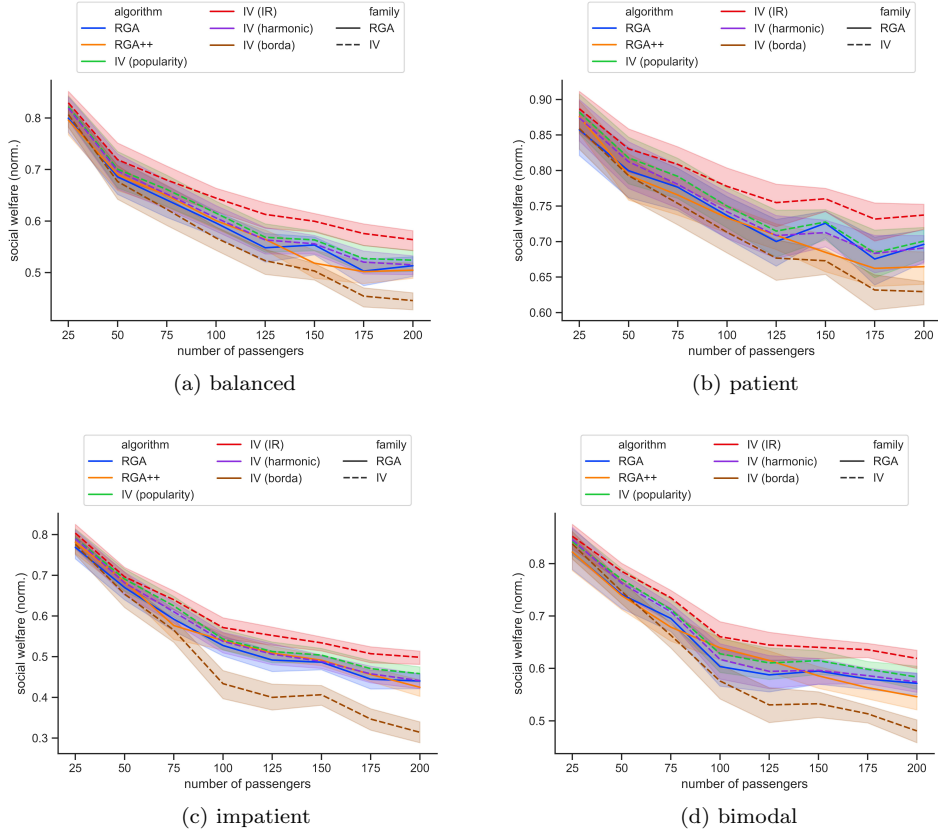


Figure 3. The social welfare, normalized by the number of passengers in the scenario without hotspots, for each algorithm. The plots show bands of 95% confidence intervals.

5. Result Discussion

We compare the performance of the proposed algorithms from a utilitarian and fairness perspective under different scenarios, where we vary (1) the number of riders, (2) their patience and (3) the similarity in their preferences in terms of desired pick-up and drop-off stations and times.

5.1. Utilitarian Perspective

We first examine our algorithms for efficiency, measuring their performance according to the mean *utilitarian social welfare*: This is the sum of the utility of the riders, normalized by the number of experiments. To allow for a better comparison between the algorithms, we also normalize by the number of riders. Formally, in each iteration, we measure the following:

$$\sum_{r_i \in N} U_{r_i}(T_d(r_i), T_a(r_i)), \quad (2)$$

where N is the number of passenger, and U_{r_i} follows equation 1.

Figure 3 depicts the normalized mean social welfare of the riders when riders' pref-

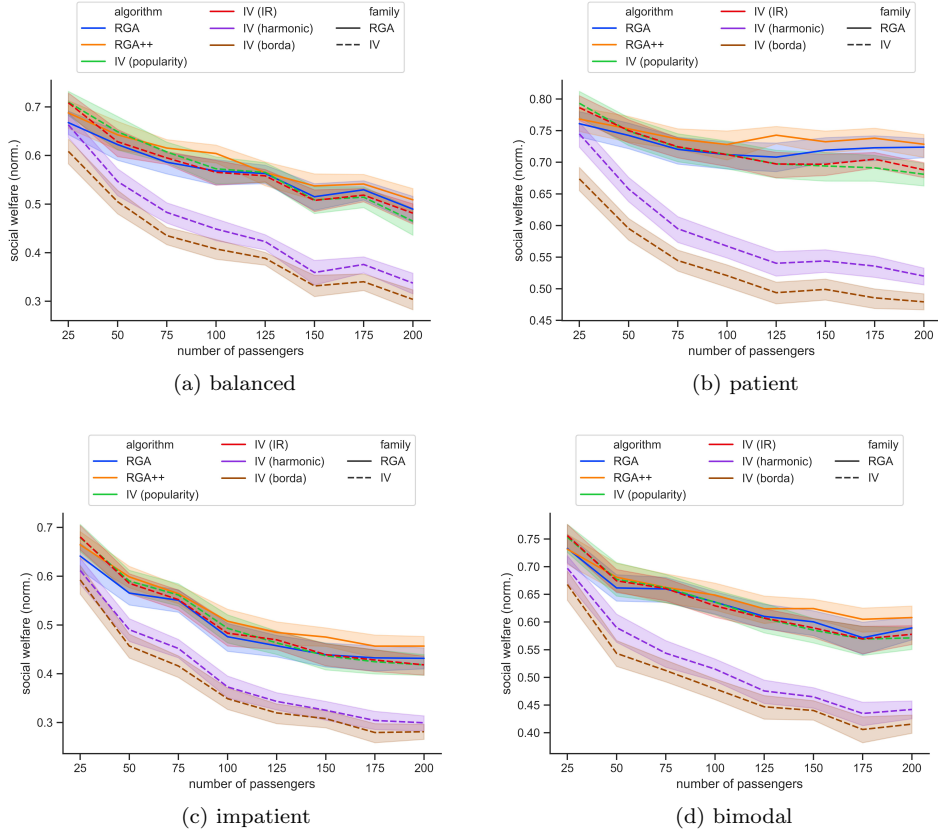


Figure 4. The social welfare, normalized by the number of passengers, in the scenario with 10 hotspots, for each algorithm. The plots show bands of 95% confidence intervals.

erences over the stations are uniform. In this case, the Iterative Voting algorithm, with the instant runoff voting rule, outperforms all others. The same algorithm, with the Borda voting rule, is clearly the worst case – as we will see later, this is a recurring theme. The remaining algorithms admit quite similar performances, with IV with popularity being the second-best algorithm in most cases.

Figures 4 and 5 present the same measure when the station sampling is based on the hotspot idea, as described above. In these scenarios, Borda and harmonic rules are clearly outperformed, while all remaining rules perform quite the same. IV-IR does not yield the best social welfare, as before, although its loss compared to the best case is not large. RGA++ performs the best in most cases now. IV-Borda again performs the worst, and in this case, IV-harmonic follows at a close distance.

5.2. Fairness Perspective

To measure fairness, we use the Gini Index (Gini (1912); see also, Moulin (1988)), a popular measure of wealth inequality, which is also used in the transportation literature (Yan & Howe, 2020), and has some precedence in the artificial intelligence

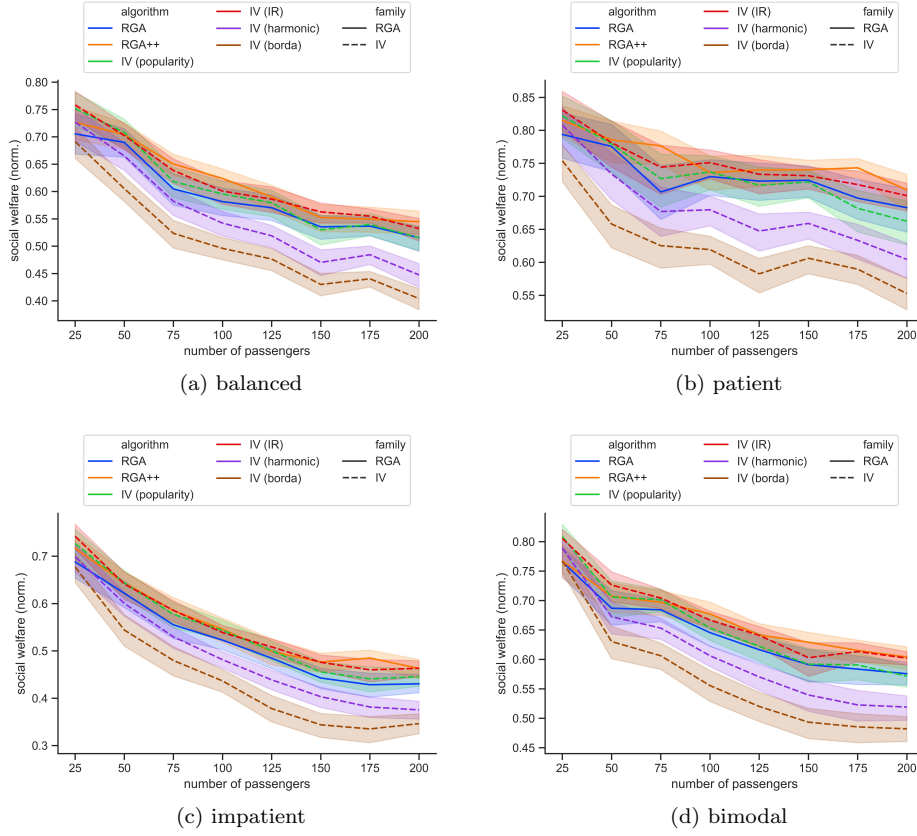


Figure 5. The social welfare, normalized by the number of passengers, in the scenario with 20 hotspots, for each algorithm. The plots show bands of 95% confidence intervals.

literature (Lackner, 2020). Formally, the Gini index is defined as the following ratio:

$$\frac{\sum_{r_i \in N} \sum_{r_j \in N} |U_{r_i}(T_d(r_i), T_a(r_i)) - U_{r_j}(T_d(r_j), T_a(r_j))|}{2 \sum_{r_i \in N} \sum_{r_j \in N} U_{r_j}(T_d(r_j), T_a(r_j))} \quad (3)$$

where N is the number of passenger, and U_{r_i} follows equation 1.

In each experiment, the Gini index is computed among the rider’s utilities, and the mean is presented in Figures 6-8. We recall that a Gini index of 0 implies total fairness in the sense that all riders receive exactly the same utility, while an index close to 1 implies total inequality. We also note that the Gini index should be considered in conjecture to our efficiency measures since a low Gini index implies nothing about the efficiency guarantees⁸.

For uniform spatial preferences (Figure 6) IV-IR returns fairer solutions. In conjecture to our efficiency results, IV-IR is clearly the best-performing algorithm from our proposals. The least fair algorithm is IR-Borda. The remaining algorithms define a middle ground with similar fairness performances.

When the spatial preferences are highly concentrated in a few hotspots, then RGA++ yields a lower Gini index in most cases. IV-Borda and IV-harmonic algorithms appear to be the least fair of all, with IV-harmonic performing slightly better

⁸Indeed, even if all utilities are negligible, then the Gini index would be 0 if their difference would be equal.

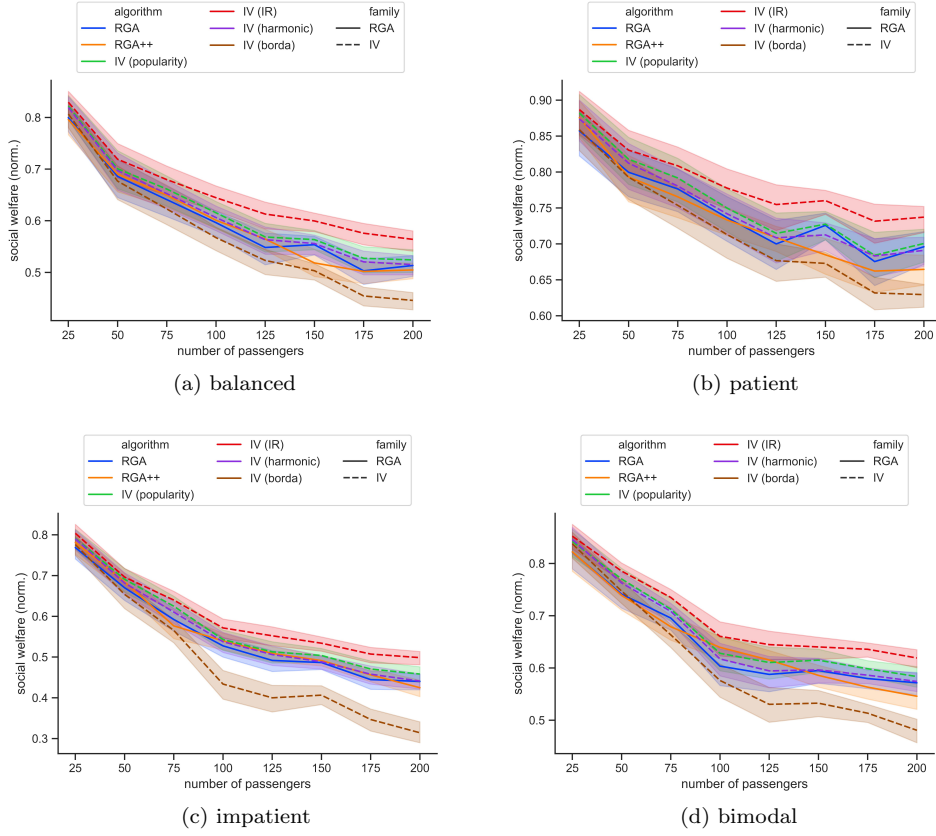


Figure 6. The mean Gini index in the scenario without hotspots, for each algorithm. The plots show bands of 95% confidence intervals.

between the two. In Section A in the Appendix, we employ a detailed analysis using Student’s t-test, where we compare each pair of algorithms along all possible scenarios we use.

5.3. Final Remarks

As a conclusion to this section, we give some final remarks:

- IV-IR appears to be the best option in the case without hotspots, both with respect to fairness and efficiency. Our detailed analysis in Section A in the Appendix further supports this. Whenever it is not the best algorithm, it is still a reasonable choice, and it performs quite well in the cases of 10 and 20 hotspots, although it falls behind the RGA++ algorithm. Due to this, we believe that this algorithm is a good candidate for further investigation. A possible explanation for this success is that IV-IR, at each point, clears up TourNodes, which are undesirable by many riders. Hence, at each point, most riders need to select two choices that are not very bad for them.
- When our experiments use hotspots, the RGA++ algorithm overperforms all the other algorithms, and RGA is a reasonable choice. This is justifiable since there are only a few locations where TourNodes can be scheduled, hence the greedy approach performs well.

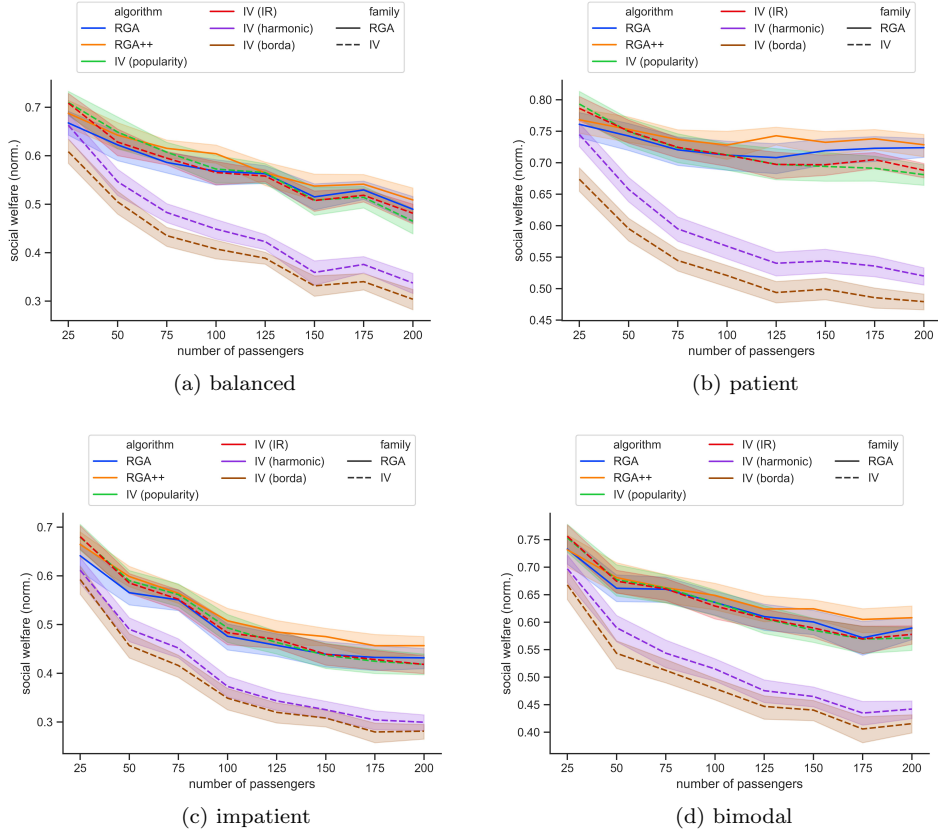


Figure 7. The mean Gini index in the scenario with 10 hotspots, for each algorithm. The plots show bands of 95% confidence intervals.

- In contrast, IV-Borda provides the worst results, again both with respect to fairness and efficiency, in nearly all cases. Due to these failures, this algorithm should not be considered for practical implementation.

Rather unsurprisingly, the IV-Borda algorithm fails to perform in both measures: the rule gives a relatively high score to TourNodes, which offers very small value to some riders. For example, consider an instance where roughly half the riders prefer station a for pick-up, half riders prefer station c , and a single rider prefers station b , and b is closer to a and b than a and b between them, hence this is the second choice for nearly all riders: under Borda rule, location b would be selected as the starting station, and almost all riders will not be fully satisfied. On the contrary, IV-IR, IV-popularity and IV-harmonic will never choose b . This decreases the overall sum of utilities and also increases the inequality between the riders. This phenomenon is more apparent in the cases of time-sensitive riders (see Figure 3c) and in the hotspot-based sampling, especially when the number of hotspots is small (see Figures 4 and 7).

- The situation is not as clear for the remaining algorithms. Most algorithms perform close to (or better than) RGA, our baseline algorithm. The RGA++ algorithm performs the best (or close to best) with respect to both objectives in highly concentrated spatial preferences. The IV-popularity algorithm is a reasonable choice, especially in the bimodal preferences example. The IV-harmonic algorithm might be useful in the uniform spatial preferences setting but suffers

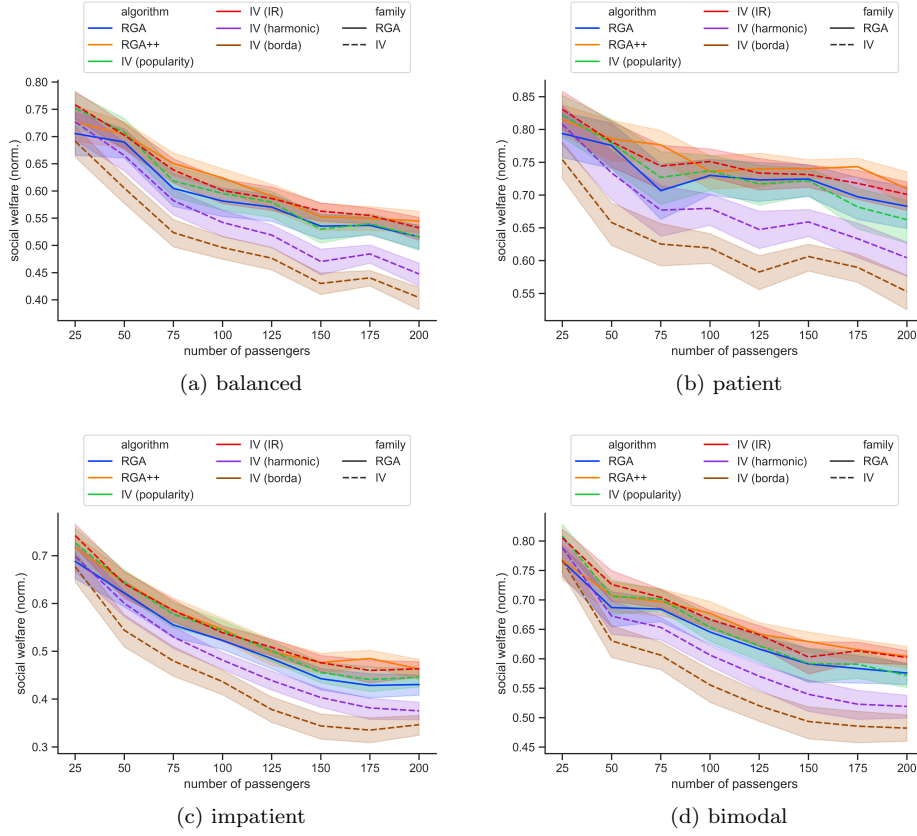


Figure 8. The mean Gini index in the scenario with 20 hotspots for each algorithm. The plots show bands of 95% confidence intervals.

- similarly to IV-Borda when a few stations dominate the demand. see Section A in Appendix for details.
- The RGA++ mechanism performs slightly better than RGA, both with respect to fairness and efficiency, in almost all cases explored. The same correlation is observed between IV-IR and IV-popularity.
 - As the number of riders increases, the performance decreases in most cases, although some notable exceptions are observed.
 - Finally, the changes in time sensitivity affect the relation between our algorithms, but not by much.

6. Related Work

This work relates to three lines of research on ridesharing: (1) optimisation-oriented approaches, (2) rider satisfaction and agent-based simulations for ridesharing, and (3) game-theoretical coalition formation and voting techniques in ridesharing. Below, we go through these lines and elaborate on how the presented approach in this work relates to each.

Building on the literature on operations research, ridesharing is traditionally understood as an optimisation problem (see the extensive review of such a view in (Agatz et al., 2012)). This perspective mainly focuses on improving the service from the ser-

vice providers’ point of view. For instance, the presented method in (Biswas et al., 2017) purely focuses on optimising the service with the main goal of profit maximisation, while the more recent work presented in (van der Tholen et al., 2021) looks at mixed-purpose autonomous vehicles (where goods and riders may share space in a transportation service). Then, the main focus is on the optimal use of the available space where riders and goods are both considered as preference-less entities. We deem that in such approaches, the preference of riders may be sacrificed as the course of a journey is optimised with respect to the service providers’ definition of optimality.

Another line of research (e.g., in Bistaffa, Rodríguez-Aguilar, Cerquides, and Blum (2018)) is focused on providing platforms for simulating ridesharing systems and modelling their behaviour. Such approaches are crucial as they provide a test bed for evaluating different ridesharing algorithms and comparing their performance. Recently, more attention has been given to the satisfaction of riders. For instance, in Levinger, Hazon, and Azaria (2020), the authors apply data-driven techniques to learn a satisfaction function that is concerned with how riders were satisfied in previously operated ridesharing journeys. Fitting to the curve of “rider behaviour/satisfaction” would be appropriate for predictably similar domains (e.g., when the service is operating in closed settings such as ridesharing within an industrial plant). However, in more unpredictable settings (such as an urban area), bounding the *current* users of a service to preferences of *past* users may result in lower satisfaction in comparison to capturing the preferences of current users themselves.

Simulation-based analysis has been used for various topics regarding public transportation and public buses in particular. We briefly present some indicative examples in the following. Dui and Zhang (2021) simulate urban taxi sharing systems. Legène, Auping, Correia, and van Arem (2020) use simulations to model the impact of automated vehicles in an urban area. Wang, Liu, and Xu (2023) use simulation to analyse the bus bunching phenomenon, where multiple busses arrive simultaneously at the same bus stop. Sung, Chu, Chang, Yeh, and Chou (2022) use simulations to tackle scheduling problems for electric buses. Moosavi, Ismail, and Yuen (2020) propose a simulation model to analyse and improve reliability factors, including waiting time. Multi-agent-based simulation has also been used to analyze and improve bus scheduling under different environments (Ap. Serratini, Liu, & Sinha, 2008; Meignan, Simonin, & Koukam, 2007; Urquhart et al., 2019).

Finally, our work is related to contributions that apply game-theoretical coalition formation techniques and voting theory to ridesharing problems (Bistaffa, Farinelli, Chalkiadakis, & Ramchurn, 2015; Cheng, Nguyen, & Lau, 2014; Dennisen & Müller, 2016; Nourinejad & Roorda, 2016). Such approaches are related to our use of voting mechanisms but mainly focus on orthogonal problems in ridesharing or zoom on a different decision point. For instance, in (Cheng et al., 2014), the authors focus on forming rider coalitions based on the amount they are willing to pay, or in (Bistaffa et al., 2015) on how to fairly share the service cost among the riders, while our focus is mainly on how the service can operate in order to maximise rider satisfaction. Then, in both (Nourinejad & Roorda, 2016) and (Dennisen & Müller, 2016), the focus is on how the journey should start and the initiation of the journey (which complements our focus on allowing riders to participate in the route generation process). In (Nourinejad & Roorda, 2016), an agent-based model is presented for matching rider groups to vehicles using auction-based techniques. Then, cost-revenue analysis is applied to evaluate matching mechanisms based on how they ensure the (financial) survival of service providers. Finally, in (Dennisen & Müller, 2016) voting methods are used to decide riders’ coalitions to share a ride. This is orthogonal to our use of voting mechanisms

to build a schedule for a ride.

7. Conclusions and Future Work

In this work, we presented algorithms for preference-aware dynamics ridesharing, used voting mechanisms to allow participation, and evaluated the performance as well as the utilitarian and fairness properties of the algorithms using simulation-based experiments. Such algorithms can be implemented in ridesharing services to improve the satisfaction of riders and, in turn, to foster the financial and environmental benefits of smart mobility.

Another extension to this work will be on how to nudge riders towards environmentally *sustainable* routes (e.g., to opt for less congested routes that may diverge from the shortest path). To that end, we envisage the applicability of incentive engineering methods (see, e.g., (Iwase, Stein, & Gerding, 2021; Protopapas, Yazdanpanah, Gerding, & Stein, 2024)). Such methods necessitate integrating a bargaining phase among computational agents that represent each rider to reach a consensus on a more sustainable trajectory that is acceptable by the riders (Baarslag, Kaisers, Gerding, Jonker, & Gratch, 2017).

Ethical Statement: Towards Sustainable and Trustworthy Smart Mobility.

We believe that ridesharing services are inherently sociotechnical practices. Thus, designing and developing autonomous systems for operating ridesharing should take into account both the technical and social aspects of the problem. Arguably, focusing merely on AI systems that optimise routes such that merely service providers' benefit is maximised is an unethical practice that results in lower user satisfaction and, in turn, lower acceptability of AI-based solutions. Our work follows the idea that designing a trustworthy partnership in systems that involve humans and AI systems needs to enable *dynamic* interactions between the two parts and has to provide means for capturing diverse *preferences* of humans in the loop (Ramchurn et al., 2021).

Data Access Statement. This study was a reanalysis of publicly available data from the NaPTAN dataset (Department for Transport, 2014). Implementations and data derived through the reanalysis undertaken in this study are available from the public GitHub repository at <https://github.com/MaxOng99/ECS-Ridesharing.git>.

Acknowledgements. We thank the anonymous reviewers and participants at *AB-MUS2022: The 6th International Workshop on Agent-Based Modelling of Urban Systems* for their incisive comments that were most useful in revising this paper. The authors acknowledge the use of the IRIDIS High Performance Computing Facility and associated support services at the University of Southampton in the completion of this work. This work was supported by the UK Engineering and Physical Sciences Research Council (EPSRC) through a Turing AI Fellowship (EP/V022067/1) on Citizen-Centric AI Systems (<https://ccais.ac.uk/>). The authors report there are no competing interests to declare. For the purpose of open access, the authors have applied a Creative Commons Attribution (CC BY) licence to any author-accepted manuscript version arising.

References

- Agatz, N. A. H., Erera, A. L., Savelsbergh, M. W. P., & Wang, X. (2012). Optimization for dynamic ride-sharing: A review. *European Journal of Operational Research*, 223(2), 295–303.
- Aïvodji, U. M., Gambs, S., Huguet, M.-J., & Killijian, M.-O. (2016). Meeting points in ridesharing: A privacy-preserving approach. *Transportation Research Part C: Emerging Technologies*, 72, 239–253.
- Ap. Sorratini, J., Liu, R., & Sinha, S. (2008). Assessing Bus Transport Reliability Using Micro-Simulation. *Transportation Planning and Technology*, 31(3), 303–324.
- Baarslag, T., Kaisers, M., Gerding, E. H., Jonker, C. M., & Gratch, J. (2017). Computers that negotiate on our behalf: Major challenges for self-sufficient, self-directed, and interdependent negotiating agents. In G. Sukthankar & J. A. Rodríguez-Aguilar (Eds.), *Autonomous Agents and Multiagent Systems - AAMAS 2017 workshops, visionary papers* (Vol. 10643, pp. 143–163). Springer.
- Badia, H., Argote-Cabanero, J., & Daganzo, C. F. (2017, September). How network structure can boost and shape the demand for bus transit. *Transportation Research Part A: Policy and Practice*, 103, 83–94.
- Department for Transport. (2014). *National public transport access nodes (NaPTAN)*. Retrieved 2021-10-05, from <https://data.gov.uk/dataset/ff93ffc1-6656-47d8-9155-85ea0b8f2251/national-public-transport-access-nodes-naptan>
- Transport of London. (2021). *Buses performance data*. Retrieved 2021-10-08, from <https://tfl.gov.uk/corporate/publications-and-reports/buses-performance-data>
- Bistaffa, F., Farinelli, A., Chalkiadakis, G., & Ramchurn, S. D. (2015). Recommending fair payments for large-scale social ridesharing. In H. Werthner, M. Zanker, J. Golbeck, & G. Semeraro (Eds.), *Proceedings of the 9th ACM Conference on Recommender Systems, RecSys* (pp. 139–146). ACM.
- Bistaffa, F., Rodríguez-Aguilar, J. A., Cerquides, J., & Blum, C. (2018). A simulation tool for large-scale online ridesharing. In E. André, S. Koenig, M. Dastani, & G. Sukthankar (Eds.), *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2018*, (pp. 1797–1799). ACM.
- Biswas, A., Gopalakrishnan, R., Tulabandhula, T., Mukherjee, K., Metrewar, A., & Rajasubramaniam, T. (2017). Profit optimization in commercial ridesharing. In K. Larson, M. Winikoff, S. Das, & E. H. Durfee (Eds.), *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, (AAMAS)*, (pp. 1481–1483). ACM.
- Boutilier, C., Caragiannis, I., Haber, S., Lu, T., Procaccia, A. D., & Sheffet, O. (2012). Optimal social choice functions: A utilitarian view. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)* (pp. 197–214).
- Brandt, F., Conitzer, V., Endriss, U., Lang, J., & Procaccia, A. D. (2016). *Handbook of computational social choice*. Cambridge, UK: Cambridge University Press.
- Cheng, S., Nguyen, D. T., & Lau, H. C. (2014). Mechanisms for arranging ride sharing and fare splitting for last-mile travel demands. In A. L. C. Bazzan, M. N. Huhns, A. Lomuscio, & P. Scerri (Eds.), *International Conference on Autonomous Agents and Multi-Agent Systems, (AAMAS)* (pp. 1505–1506). Richland, SC, USA: IFAAMAS/ACM.
- Dennisen, S. L., & Müller, J. P. (2016). Iterative committee elections for collective decision-making in a ride-sharing application. In *Proceedings of the ninth international workshop on agents in traffic and transportation (ATT)* (Vol. 1678, pp. 1–8).
- Dui, H., & Zhang, C. (2021). Simulations for urban taxi sharing system on routes and passengers with numerical experiments. *Journal of Simulation*, 15(4), 273–283.
- Furuhata, M., Dessouky, M., Ordóñez, F., Brunet, M.-E., Wang, X., & Koenig, S. (2013). Ridesharing: The state-of-the-art and future directions. *Transportation Research Part B: Methodological*, 57, 28–46.
- Gini, C. W. (1912). *Variabilità e mutuabilità. Contributo allo studio delle distribuzioni e delle relazioni statistiche*. Bologna: C. Cuppini.

- Gupta, A. K., & Nadarajah, S. (2004). *Handbook of beta distribution and its applications*. FL, US: CRC press.
- Hasan, M. H., Van Hentenryck, P., Budak, C., Chen, J., & Chaudhry, C. (2018). Community-based trip sharing for urban commuting. *Thirty-Second AAAI Conference on Artificial Intelligence (AAAI)*, 32(1).
- Iwase, T., Stein, S., & Gerding, E. H. (2021, 8). A polynomial-time, truthful, individually rational and budget balanced ridesharing mechanism. In Z.-H. Zhou (Ed.), *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21* (pp. 268–275). International Joint Conferences on Artificial Intelligence Organization.
- Jalali, R., Koohi-Fayegh, S., El-Khatib, K., Hoornweg, D., & Li, H. (2017). Investigating the potential of ridesharing to reduce vehicle emissions. *Urban Planning*, 2(2), 26–40.
- König, A., & Gripenkoven, J. (2020). Modelling travelers’ appraisal of ridepooling service characteristics with a discrete choice experiment. *European Transport Research Review*, 12(1), 1.
- Lackner, M. (2020). Perpetual voting: Fairness in long-term decision making. *The Thirty-Fourth AAAI Conference on Artificial Intelligence, (AAAI)*, 34(02), 2103–2110.
- Leg  ne, M. F., Auping, W. L., Correia, G. H. d. A., & van Arem, B. (2020). Spatial impact of automated driving in urban areas. *Journal of Simulation*, 14(4), 295–303.
- Levinger, C., Hazon, N., & Azaria, A. (2020). Human satisfaction as the ultimate goal in ridesharing. *Future Generation Computer Systems*, 112, 176–184.
- Meignan, D., Simonin, O., & Koukam, A. (2007). Simulation and evaluation of urban bus-networks using a multiagent approach. *Simulation Modelling Practice and Theory*, 15(6), 659–671.
- Moosavi, S. M. H., Ismail, A., & Yuen, C. W. (2020, May). Using simulation model as a tool for analyzing bus service reliability and implementing improvement strategies. *PLOS ONE*, 15(5), e0232799.
- Moulin, H. (1988). *Axioms of cooperative decision making*. Cambridge University Press.
- Mourad, A., Puchinger, J., & Chu, C. (2019). A survey of models and algorithms for optimizing shared mobility. *Transportation Research Part B: Methodological*, 123, 323–346.
- Nourinejad, M., & Roorda, M. J. (2016). Agent based model for dynamic ridesharing. *Transportation Research Part C: Emerging Technologies*, 64, 117–132.
- Ong, Y. C., Protopapas, N., Yazdanpanah, V., Gerding, E., & Stein, S. (2022). Preference-aware dynamic ridesharing. In M. L. Kieu, K. H. van Dam, J. Thompson, N. Malleison, A. Heppenstall, & J. Ge (Eds.), *Proceedings of The 6th International Workshop on Agent-Based Modelling of Urban Systems (ABMUS 2022)*.
- Protopapas, N., Yazdanpanah, V., Gerding, E., & Stein, S. (2024). Online decentralised mechanisms for dynamic ridesharing. In *Proceedings of the 2024 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2024, Auckland, New Zealand, to appear*. IFAAMAS.
- Ramchurn, S. D., Stein, S., & Jennings, N. R. (2021). Trustworthy human-ai partnerships. *iScience*, 24(8), 102891.
- Rigas, E. S., Ramchurn, S. D., & Bassiliades, N. (2014). Managing electric vehicles in the smart grid using artificial intelligence: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 16(4), 1619–1635.
- Robusto, C. C. (1957). The cosine-haversine formula. *The American Mathematical Monthly*, 64(1), 38–40.
- Stein, S., & Yazdanpanah, V. (2023). Citizen-centric multiagent systems. In N. Agmon, B. An, A. Ricci, & W. Yeoh (Eds.), *Proceedings of the 2023 international conference on autonomous agents and multiagent systems, AAMAS 2023, London, United Kingdom*, (pp. 1802–1807). ACM.
- Sung, Y.-W., Chu, J. C., Chang, Y.-J., Yeh, J.-C., & Chou, Y.-H. (2022). Optimizing mix of heterogeneous buses and chargers in electric bus scheduling problems. *Simulation Modelling Practice and Theory*, 119, 102584.
- Urquhart, N., Powers, S., Wall, Z., Fonzone, A., Ge, J., & Polhill, J. G. (2019). Simulating

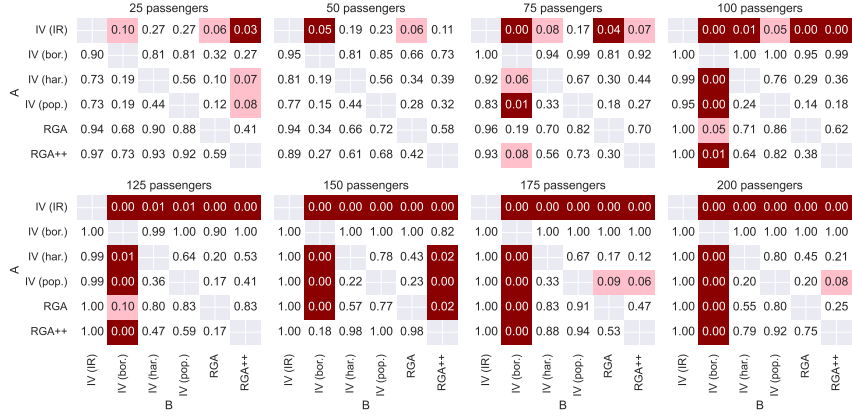
- the actions of commuters using a multi-agent system. *Journal of Artificial Societies and Social Simulation*, 22(2), 10.
- van der Tholen, M., Beirigo, B. A., Jovanova, J., & Schulte, F. (2021). The share-a-ride problem with integrated routing and design decisions: The case of mixed-purpose shared autonomous vehicles. In M. Mes, E. Lalla-Ruiz, & S. Voß (Eds.), *Computational Logistics* (pp. 347–361). Springer International Publishing.
- Wang, X., Liu, T., & Xu, M. (2023). Bus bunching simulation based on the lattice Boltzmann model of traffic flow. *SIMULATION*, 99(2), 183–199.
- Yan, A., & Howe, B. (2020). Fairness in practice: a survey on equity in urban mobility. *A Quarterly bulletin of the Computer Society of the IEEE Technical Committee on Data Engineering*, 42(3).
- Zavvos, E., Gerding, E. H., Yazdanpanah, V., Maple, C., Stein, S., & m. c. schraefel. (2022). Privacy and trust in the internet of vehicles. *IEEE Trans. Intell. Transp. Syst.*, 23(8), 10126–10141.
- Zhang, R., Spieser, K., Frazzoli, E., & Pavone, M. (2015). Models, algorithms, and evaluation for autonomous mobility-on-demand systems. In *American control conference, ACC 2015* (pp. 2573–2587). IEEE.

Appendix A. Statistical Significance Tests

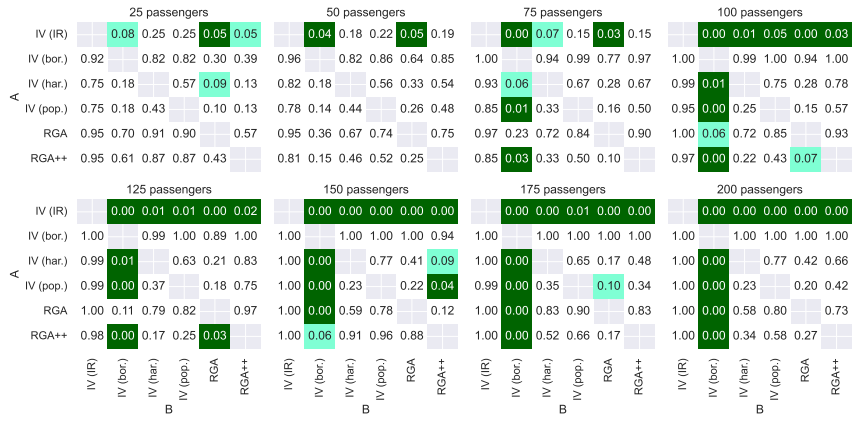
In this section, we complement our result from Section 5, by analysing the statistical significance between our algorithms’ computed average social welfare and average Gini indices. In particular, for each scenario (i.e., for each combination of parameters from Table 3), we employ a paired sample Student’s t-test comparison between pairs of algorithms. For the social welfare, given algorithm pairs A and B, we ask (as an alternative hypothesis) whether *the mean social welfare obtained by algorithm A is higher than the social welfare obtained by algorithm B*? The t-test returns a p-value. We can safely adopt the alternative hypothesis when the p-value is close to 0. Figure A1a shows the p-values obtained by each t-test for the case where the rider’s patience is balanced and no hotspots exist (corresponding to Figure 3a). A low p-value (e.g., at most 0.05) implies that the algorithm in the row yields a better outcome than the algorithm in the column (this applies to all figures in this section). In the case of Figure A1a, we can observe that IV-IR provides higher social welfare compared to all algorithms with confidence at least 95% with at least 100 passengers. This confidence increases to 99% for more than 125 passengers. On the contrary, as the number of passengers increases, all algorithms perform better than the IV algorithm with Borda count.

When we focus on fairness, we employ the alternative hypothesis *the mean Gini index of algorithm A is lower than that of algorithm B*. Recall that the optimal value for the Gini index is 0. A similar trend with social welfare follows here, with IV-IR being the algorithm with the lowest Gini index compared to all algorithms (see Figure A2b). As the number of hotspots increases, then RGA++ becomes both the most efficient and fairest algorithm (see Figures A2 and A3).

Figures A4, A5, and A6 focus on scenarios with mostly patient riders. Similar to the balanced patience case, IV-IR stands out as the most efficient and fair algorithm without hotspots while losing this status to RGA++ as hotspots are employed. Similar trends also appear in Figures A4, A5, and A6 with a focus on mostly impatient riders. Observe that RGA++ works much better as the demand becomes more predictable, while IV-IR is the clear winner for unpredictable demand.



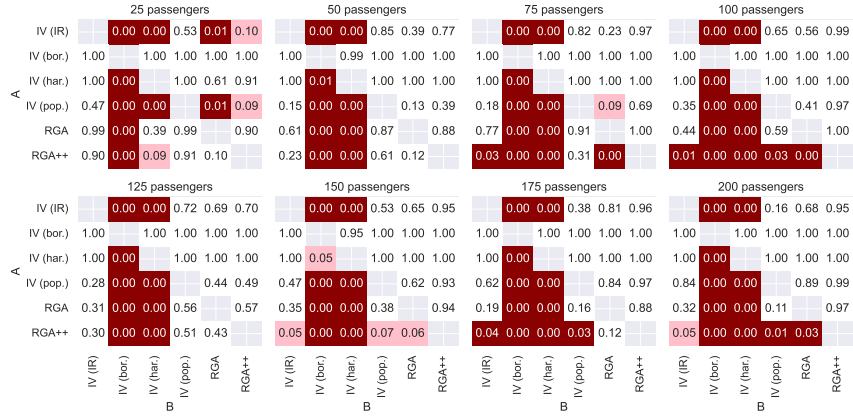
(a) social welfare



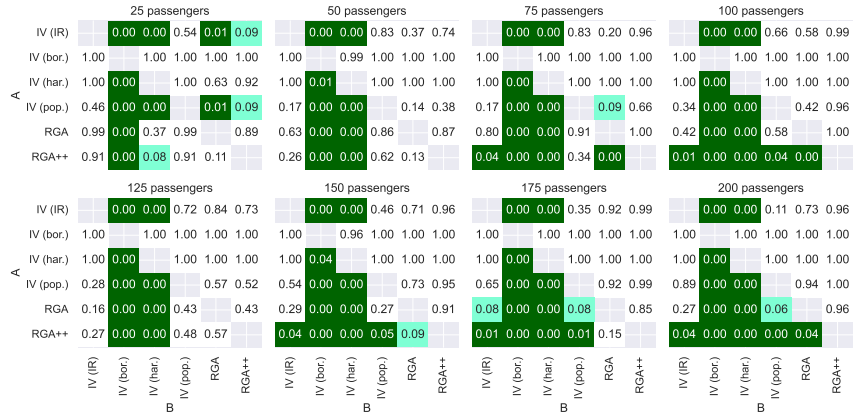
(b) Gini index

Figure A1. The p-values of our t-test comparisons. Scenario with balanced patience riders and no hotspots.

Finally, Figures A10, A11 and A12 analyse the case where the riders are a mixture of very patient and very impatient riders. While the already observed effects between IV-IR and RGA++ appear again, the IV-popularity algorithm stands out as a reliable third choice.

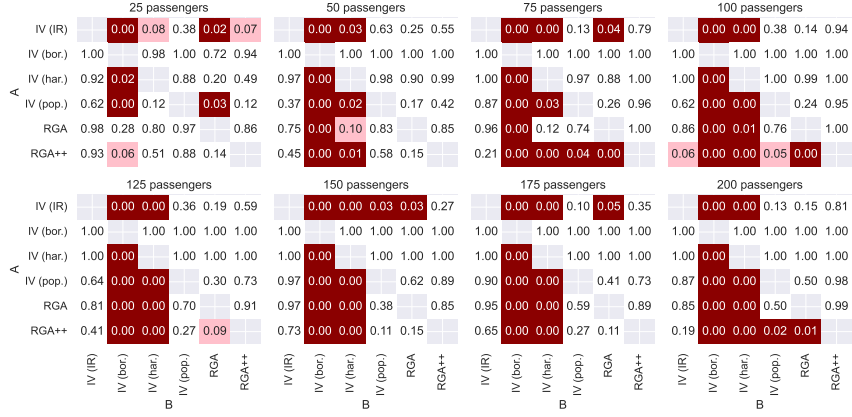


(a) social welfare

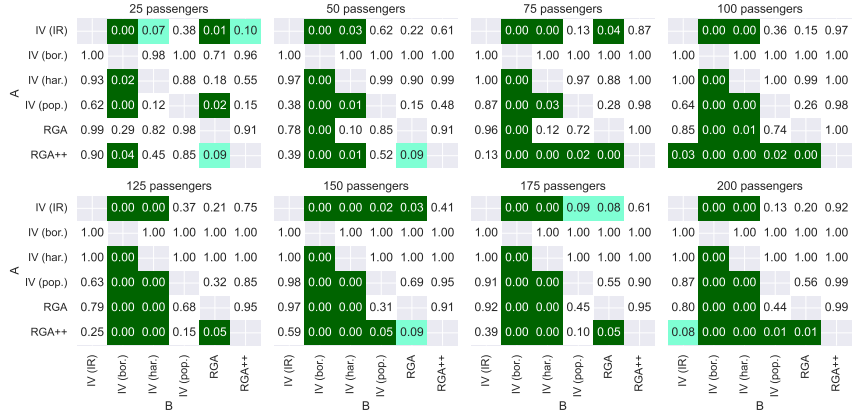


(b) Gini index

Figure A2. The p-values of our t-test comparisons. Scenario with balanced patience riders and 10 hotspots.

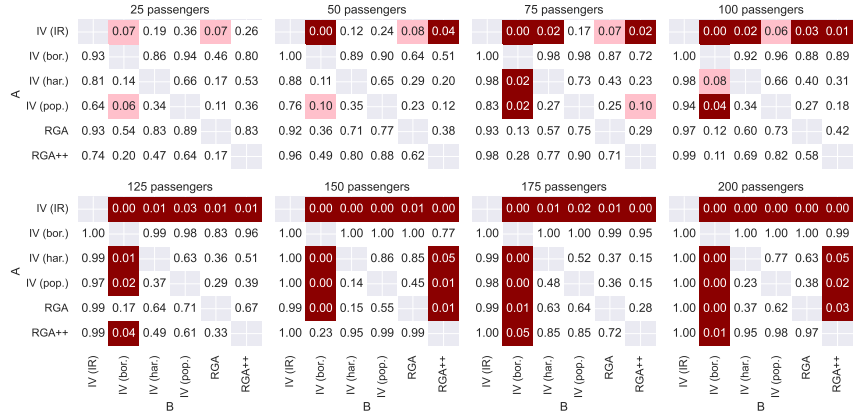


(a) social welfare

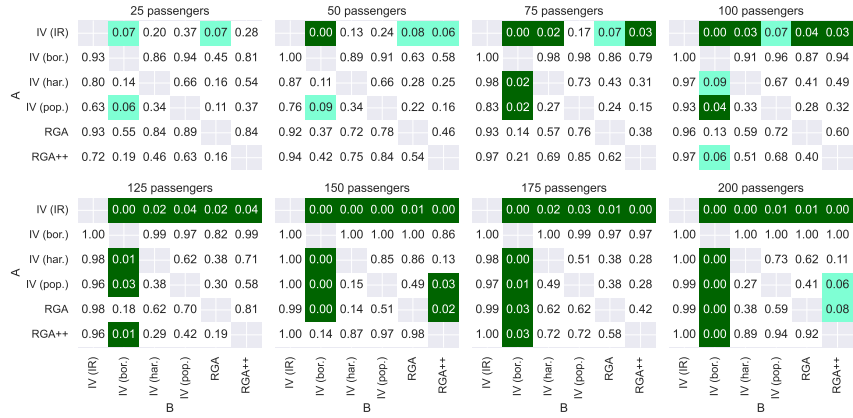


(b) Gini index

Figure A3. The p-values of our t-test comparisons. Scenario with balanced patience riders and 20 hotspots.

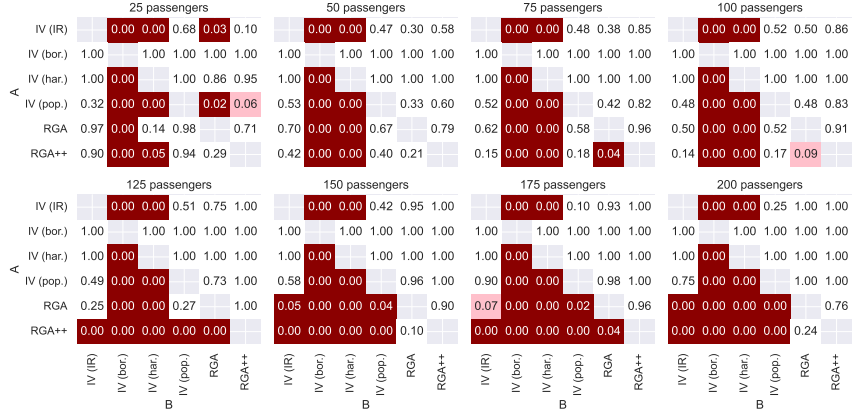


(a) social welfare

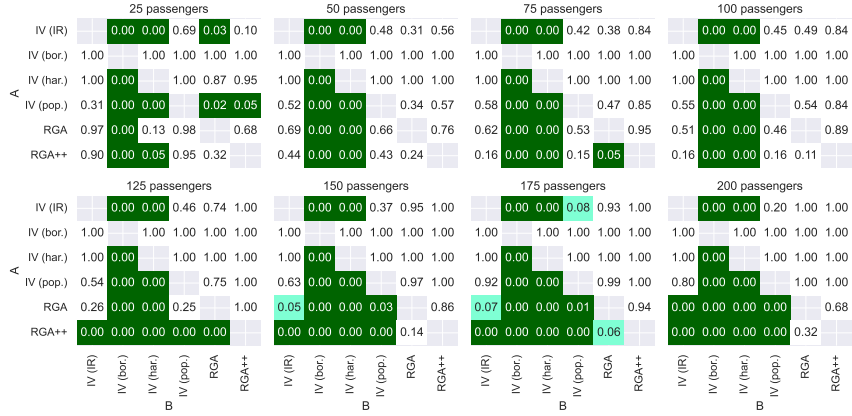


(b) Gini index

Figure A4. The p-values of our t-test comparisons. Scenario with mostly patient riders and no hotspots.



(a) social welfare



(b) Gini index

Figure A5. The p-values of our t-test comparisons. Scenario with mostly patient riders and 10 hotspots.

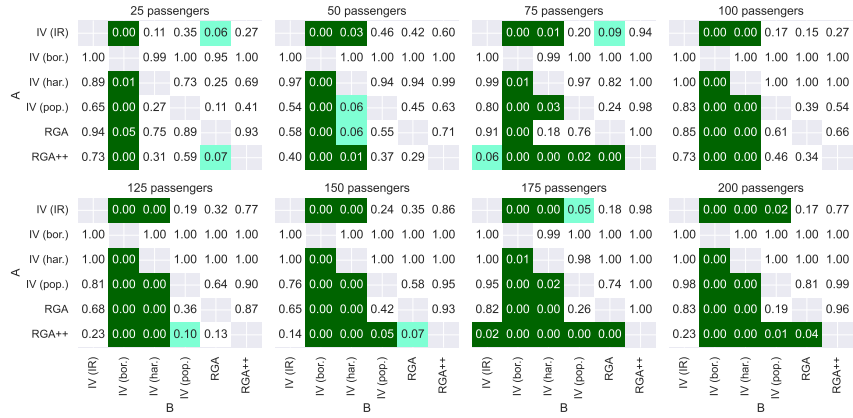
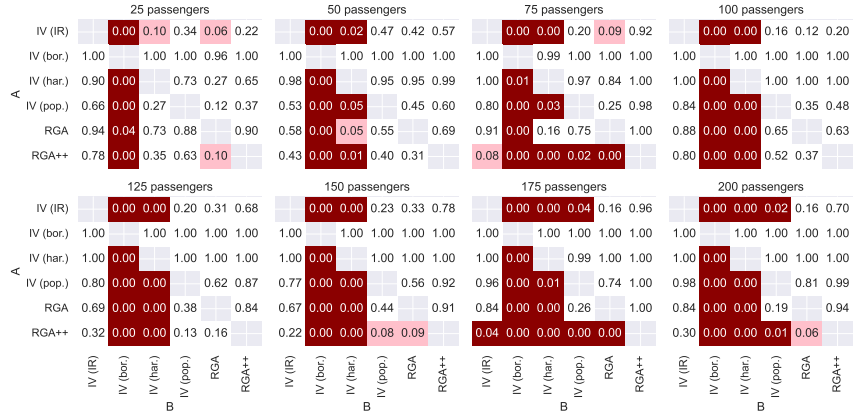
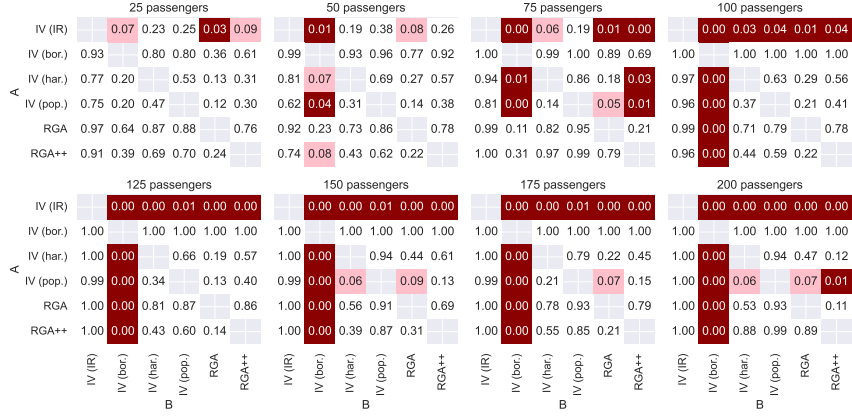
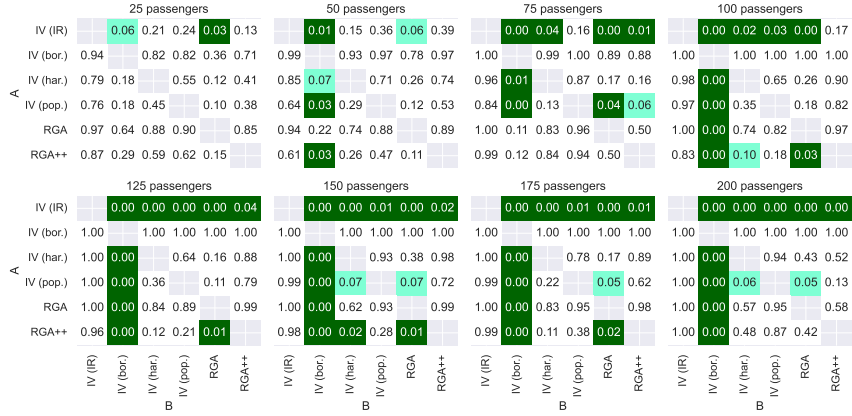


Figure A6. The p-values of our t-test comparisons. Scenario with mostly patient riders and 20 hotspots.



(a) social welfare



(b) Gini index

Figure A7. The p-values of our t-test comparisons. Scenario with mostly impatient riders and no hotspots.

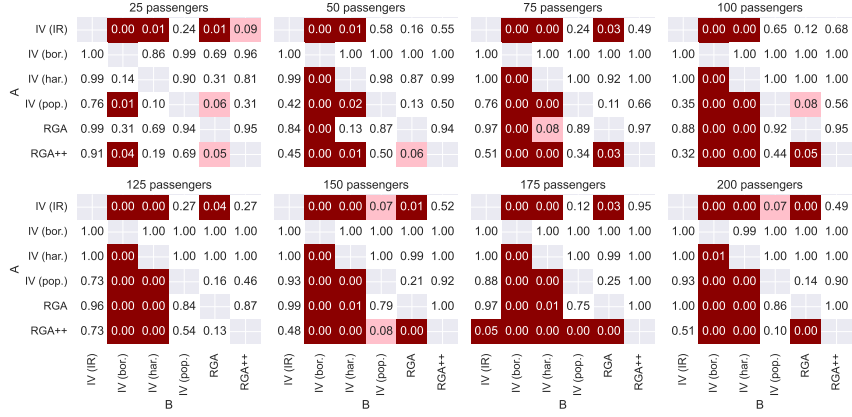
		25 passengers					50 passengers					75 passengers					100 passengers										
A	IV (IR)		0.00	0.00	0.47	0.02	0.19		0.00	0.00	0.64	0.08	0.83		0.00	0.00	0.71	0.47	0.79		0.00	0.00	0.69	0.35	0.91		
	IV (bor.)	1.00		0.86	1.00	0.99	1.00	1.00		0.98	1.00	1.00	1.00	1.00	1.00		1.00	1.00	1.00	1.00	1.00		0.93	1.00	1.00	1.00	
	IV (har.)	1.00	0.14		1.00	0.94	1.00	1.00	0.02		1.00	1.00	1.00	1.00	0.00		1.00	1.00	1.00	1.00	1.00	0.07		1.00	1.00	1.00	
	IV (pop.)	0.53	0.00	0.00		0.02	0.21	0.36	0.00	0.00		0.07	0.71	0.29	0.00	0.00		0.25	0.58	0.31	0.00	0.00		0.21	0.76		
	RGA	0.98	0.01	0.06	0.98		0.92	0.92	0.00	0.00	0.93		1.00	0.53	0.00	0.00	0.75		0.93	0.65	0.00	0.00	0.79		0.99		
	RGA++	0.81	0.00	0.00	0.79	0.08		0.17	0.00	0.00	0.29	0.00		0.21	0.00	0.00	0.42	0.07		0.09	0.00	0.00	0.24	0.01			
		125 passengers					150 passengers					175 passengers					200 passengers										
A	IV (IR)		0.00	0.00	0.33	0.20	0.84		0.00	0.00	0.44	0.49	0.99		0.00	0.00	0.42	0.61	0.96		0.00	0.00	0.50	0.83	1.00		
	IV (bor.)	1.00		0.96	1.00	1.00	1.00	1.00		0.94	1.00	1.00	1.00	1.00	1.00		1.00	0.97	1.00	1.00	1.00	1.00		0.95	1.00	1.00	1.00
	IV (har.)	1.00	0.04		1.00	1.00	1.00	1.00	0.06		1.00	1.00	1.00	1.00	0.03		1.00	1.00	1.00	1.00	1.00	0.05		1.00	1.00	1.00	
	IV (pop.)	0.67	0.00	0.00		0.36	0.89	0.56	0.00	0.00		0.54	0.99	0.58	0.00	0.00		0.66	0.96	0.50	0.00	0.00		0.80	0.99		
	RGA	0.80	0.00	0.00	0.64		1.00	0.51	0.00	0.00	0.46		1.00	0.39	0.00	0.00	0.34		0.98	0.17	0.00	0.00	0.20		1.00		
	RGA++	0.16	0.00	0.00	0.11	0.00		0.01	0.00	0.00	0.01	0.00		0.04	0.00	0.00	0.04	0.02		0.00	0.00	0.00	0.01	0.00			
		IV (IR)	IV (bor.)	IV (har.)	IV (pop.)	RGA	RGA++	IV (IR)	IV (bor.)	IV (har.)	IV (pop.)	RGA	RGA++	IV (IR)	IV (bor.)	IV (har.)	IV (pop.)	RGA	RGA++	IV (IR)	IV (bor.)	IV (har.)	IV (pop.)	RGA	RGA++		

(a) social welfare

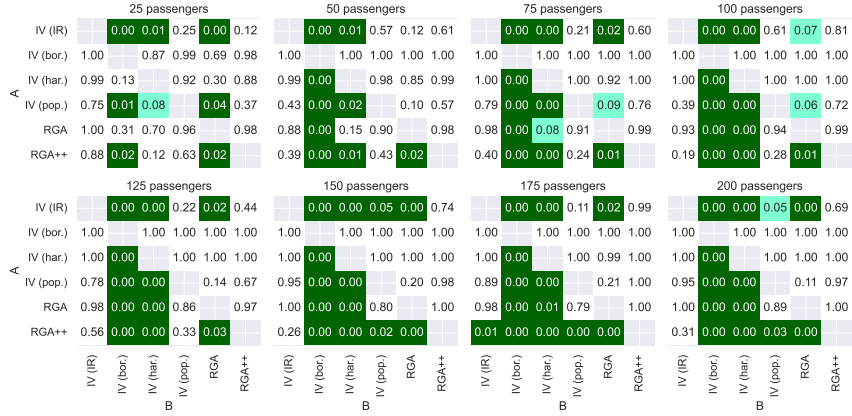
		25 passengers					50 passengers					75 passengers					100 passengers										
A	IV (IR)		0.00	0.00	0.46	0.01	0.19		0.00	0.00	0.62	0.05	0.83		0.00	0.00	0.69	0.43	0.80		0.00	0.00	0.70	0.32	0.92		
	IV (bor.)	1.00		0.87	1.00	0.99	1.00	1.00		0.98	1.00	1.00	1.00	1.00	1.00		1.00		1.00	1.00	1.00	1.00		0.94	1.00	1.00	1.00
	IV (har.)	1.00	0.13		1.00	0.93	1.00	1.00	0.02		1.00	1.00	1.00	1.00	0.00		1.00	1.00	1.00	1.00	1.00	0.06		1.00	1.00	1.00	
	IV (pop.)	0.54	0.00	0.00		0.02	0.23	0.38	0.00	0.00		0.06	0.72	0.31	0.00	0.00		0.23	0.61	0.30	0.00	0.00		0.18	0.77		
	RGA	0.99	0.01	0.07	0.98		0.94	0.95	0.00	0.00	0.94		1.00	0.57	0.00	0.00	0.77		0.95	0.68	0.00	0.00	0.82		0.99		
	RGA++	0.81	0.00	0.00	0.77	0.06		0.17	0.00	0.00	0.28	0.00		0.20	0.00	0.00	0.39	0.05		0.08	0.00	0.00	0.23	0.01			
		125 passengers					150 passengers					175 passengers					200 passengers										
A	IV (IR)		0.00	0.00	0.34	0.18	0.90		0.00	0.00	0.43	0.49	0.99		0.00	0.00	0.42	0.63	0.97		0.00	0.00	0.45	0.84	1.00		
	IV (bor.)	1.00		0.96	1.00	1.00	1.00	1.00	1.00		0.95	1.00	1.00	1.00	1.00		1.00	0.98	1.00	1.00	1.00	1.00		0.96	1.00	1.00	1.00
	IV (har.)	1.00	0.04		1.00	1.00	1.00	1.00	0.05		1.00	1.00	1.00	1.00	0.02		1.00	1.00	1.00	1.00	1.00	0.04		1.00	1.00	1.00	
	IV (pop.)	0.66	0.00	0.00		0.32	0.92	0.57	0.00	0.00		0.55	1.00	0.58	0.00	0.00		0.67	0.97	0.55	0.00	0.00		0.84	1.00		
	RGA	0.82	0.00	0.00	0.68		1.00	0.51	0.00	0.00	0.45		1.00	0.37	0.00	0.00	0.33		0.99	0.16	0.00	0.00	0.16		1.00		
	RGA++	0.10	0.00	0.00	0.08	0.00		0.01	0.00	0.00	0.00	0.00		0.03	0.00	0.00	0.03	0.01		0.00	0.00	0.00	0.00	0.00			
		IV (IR)	IV (bor.)	IV (har.)	IV (pop.)	RGA	RGA++	IV (IR)	IV (bor.)	IV (har.)	IV (pop.)	RGA	RGA++	IV (IR)	IV (bor.)	IV (har.)	IV (pop.)	RGA	RGA++	IV (IR)	IV (bor.)	IV (har.)	IV (pop.)	RGA	RGA++		

(b) Gini index

Figure A8. The p-values of our t-test comparisons. Scenario with mostly impatient riders and 10 hotspots.

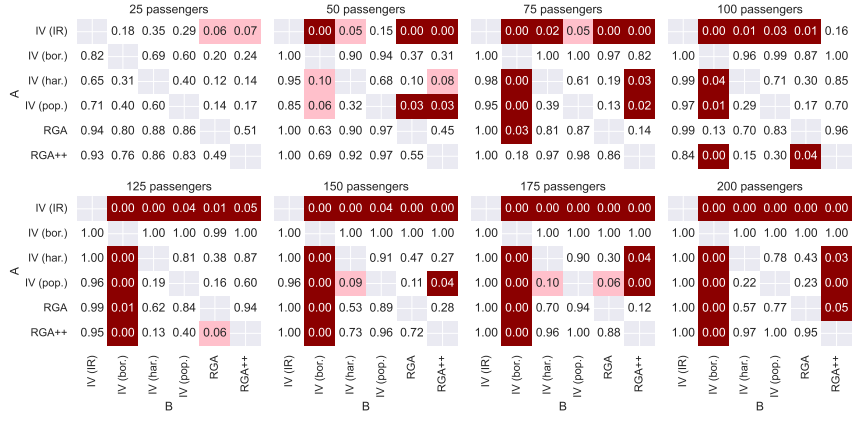


(a) social welfare

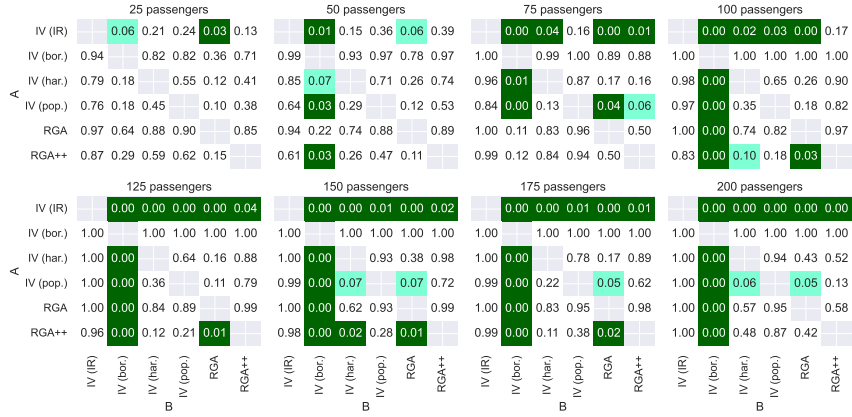


(b) Gini index

Figure A9. The p-values of our t-test comparisons. Scenario with mostly impatient riders and 20 hotspots.

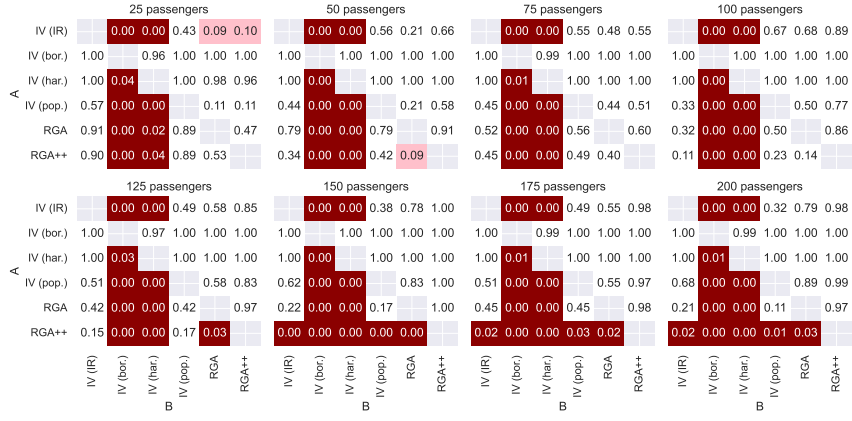


(a) social welfare

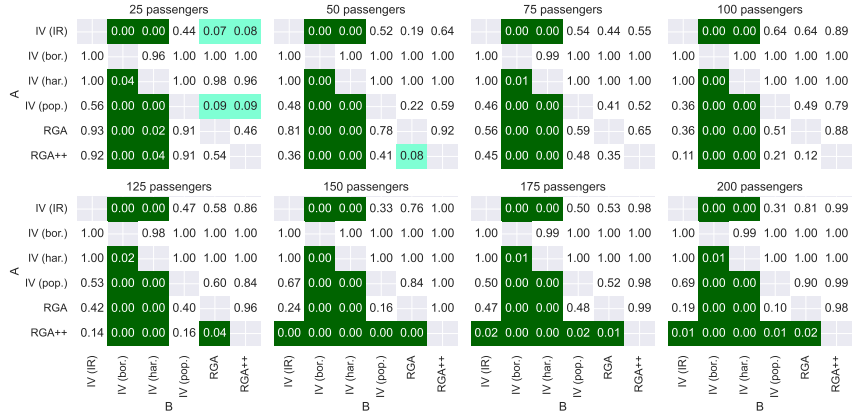


(b) Gini index

Figure A10. The p-values of our t-test comparisons. Scenario with mixtures of patient and impatient riders and no hotspots.

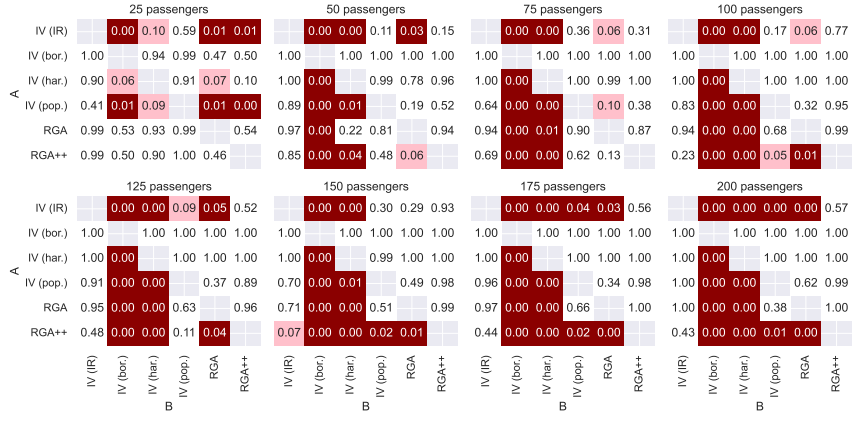


(a) social welfare

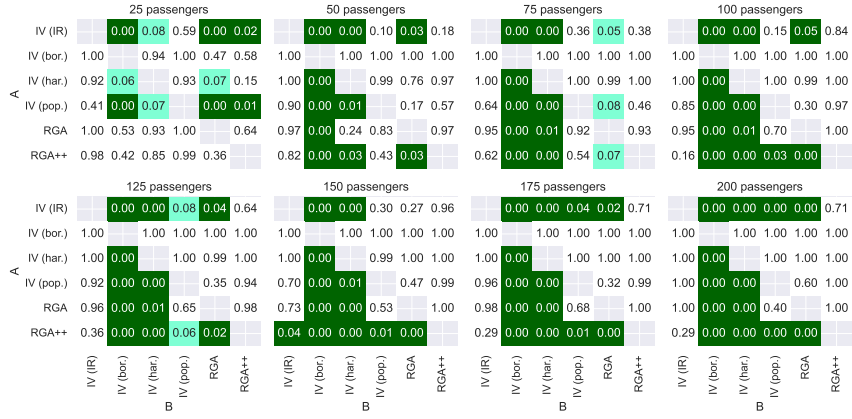


(b) Gini index

Figure A11. The p-values of our t-test comparisons. Scenario with mixtures of patient and impatient riders and 10 hotspots.



(a) social welfare



(b) Gini index

Figure A12. The p-values of our t-test comparisons. Scenario with mixtures of patient and impatient riders and 20 hotspots.