



The increasing gap dynamics in a general spatial matching model

Andres Fielbaum^{a,*}, Roberto Cominetti^{b,c}, Jose Correa^d

^a School of Civil Engineering, University of Sydney, Darlingtown, 2008, NSW, Australia

^b Institute for Mathematical and Computational Engineering, Faculty of Mathematics and School of Engineering, Pontificia Universidad Católica, Santiago, Chile

^c Department of Industrial and Systems Engineering, Faculty of Engineering, Pontificia Universidad Católica, Santiago, Chile

^d Department of Industrial Engineering, Universidad de Chile, Santiago, Chile

ARTICLE INFO

Keywords:

Increasing gap dynamics
Spatial matching
Wild goose chase
Ride-hailing
Greedy assignment

ABSTRACT

We study a representation of a problem that appears in numerous transport systems: N servers distributed over a given space (e.g., cars on an urban network), receive random requests from arriving users who get assigned to the closest server, after which this server is replaced by a new one at a random location. We show that this creates a negative feedback loop, which we call *Increasing Gap Dynamics* (IGD): when a server is assigned a spatial gap forms, which is more likely to attract new users that further widen the gap.

The simplest version of our model is a one-dimensional circle, for which we derive analytical results showing that the system converges to an inefficient equilibrium, worse than both balanced and fully random distributions of servers. We prove that an optimal assignment policy always matches the user to one of its two neighbouring servers so that long gaps tend to widen. Hence, the IGD persists even when assigning optimally rather than greedily. In two dimensions, the appearance of the IGD is illustrated through simulations on a square region. Finally, simulations of a proper ride-hailing system using real data from Manhattan confirm that the IGD arises and is responsible for the appearance of the well-known Wild Goose Chase when this phenomenon occurs.

1. Introduction: A negative feedback in spatial matching

Consider the following generic situation. There is a spatial domain, like a graph or an Euclidean region, and users who appear randomly in space at a roughly constant rate. These users must be matched with a server; the closer, the better. Concurrently, servers also appear randomly in space after finishing service for a previous user, potentially at a different location from where they started. This informal and abstract formulation can represent various transport systems, such as:

- Ride-hailing: Servers are drivers who appear upon completion of a ride and become available again.
- Dockless bike-sharing: Servers are the bicycles.
- On-street parking when delivering: Users are the households where items must be delivered, and servers are parking lots where the deliverer will park.

What is the evolution of these types of systems? Can we have a general model representing all of them? In particular, in all the examples above, it has been found that situations like the *Wild Goose Chase* (WGC) can occur (Xu and Sun, 2024; Castillo et al., 2024;

* Corresponding author.

E-mail address: andres.fielbaum@sydney.edu.au (A. Fielbaum).




Fig. 1. An example of the dynamic that prevents servers from remaining balanced in a spatial setting.

Ouyang and Yang, 2023; Emami and Ramezani, 2024; Li and Ge, 2025; Fielbaum, 2025). The WGC refers to a situation where the spatial distance between matched users and servers becomes inefficiently large. The WGC is undesirable but can be an equilibrium because the long distance implies that servers remain occupied for long periods, as servers are busy from the moment users are matched, thereby reducing the number of available servers. We refer to this negative feedback cycle between the number of idle servers and serving time as the *WGC cycle*. While this bad equilibrium has been clearly reported in the literature, the dynamics that explain why WGC appears in the first place have remained underexplored.

In this paper, we study the spatial evolution of this generic system. We propose a simple Markovian model and present several theoretical and experimental results, along with a few open questions. We also explain the mechanisms that make the WGC to appear. The reason is actually quite simple and rests on a fundamental limitation, namely the impossibility of keeping a balanced distribution of the servers (unless the system can be intervened and thus the servers rebalanced). Let us first explain it intuitively using Fig. 1:

- (a) A circle with servers (cars) uniformly distributed, where we highlight the *dominance zone* of every server s . That is, the region where s is the nearest compared to others so that a user appearing there would be assigned to s assuming a greedy rule.
- (b) A user appears, and it is assigned to the closest car.
- (c) The car is now taken so the two surrounding dominance zones are merged into a larger zone.

Even when a new car appears, the size of the dominance zones will now be heterogeneous. The larger areas will have a greater probability of receiving the next user, and when this happens, there are two negative effects:

1. The user will, on average, face an increased cost, as measured by the distance to its closest server.
2. The large area will increase even further, making the problem worse for future users. That is, there is a **negative feedback loop**, where the gaps created when one server disappears (gets assigned) tend to become bigger and bigger. We call this the *Increasing Gaps Dynamics* (IGD).

The IGD keeps worsening the servers' distribution until a probabilistic equilibrium is reached. As will be shown in the subsequent sections, this equilibrium is not only worse than a uniform distribution of servers but also worse than a fully random one.

We study and demonstrate these dynamics using three different techniques. First, we focus on what could be considered to be the simplest possible setting, namely a one-dimensional circle, where we provide a number of theoretical results. We then consider a unit square in two dimensions, where analytical results seem beyond reach but where we present visualisations to further develop the intuition. Finally, we include a set of microsimulation results using a real-world dataset from Manhattan (New York), showing that these dynamics also hold under realistic configurations. To enhance readability, the rest of this paper assumes we are describing ride-hailing, but the model remains general.

The rest of the paper is structured as follows. Section 2 gives an overview of previous studies related to this one. Section 3 introduces the formal mathematical model. Section 4 provides a detailed theoretical analysis of the dynamics on the one-dimensional circle. Section 5 focuses on the unit square and provides visualisations and a video to strengthen the understanding of these dynamics. Section 6 shows the results of several simulations on stylised settings and on Manhattan. Section 7 consolidates the two relevant dynamics that take place in these systems: IGD and WGC, through a schematic explanation of the feedback loops involved. Finally, Section 8 concludes and suggests directions for further research.

2. Related works

Our paper describes the evolution of dynamic spatial matching models. Therefore, two lines of research are the most relevant to review. First, the challenges of making decisions with partial information in these systems, and second, their analysis at equilibrium.

2.1. Forward-looking decisions in on-demand mobility

App-based on-demand mobility has become very popular in recent years. Different from traditional on-demand modes (e.g. taxis), the apps have the ability to collect in real time the information of all drivers and users, to then decide which is the best assignment.¹ However, the information they collect is only about the current users and drivers; and the optimal decision contingent on that information might differ from what would be decided if all the information were known beforehand (Berbeglia et al., 2010).

Researchers have proposed various approaches to address this challenge. First, there are several methods to decide how to *rebalance* the fleet, i.e., to instruct the idle vehicles to relocate at places that are expected to be undersupplied. Guo et al. (2021), Yang and Ramezani (2022) propose learning-based methods for rebalancing in ride-hailing systems, while Pan et al. (2019) also develops machine learning to rebalance dockless bike-sharing, a type of system where empirical studies show that only some stations can achieve “self-balancing” (Hua et al., 2025).

An interesting difference emerges when several passengers can share the vehicle at the same time. In this case, it might be optimal to follow a route that is not the shortest path, to try to find more passengers along the way; additionally, the decision of whom to pick up and drop off first can also include anticipatory techniques. These ideas are the focus of the papers by Alonso-Mora et al. (2017), Fielbaum et al. (2022).

Instead of instructing the vehicles to move in a certain way, similar behaviours can be induced through pricing. In ride-hailing, this is known as *surge pricing*, namely, when undersupplied zones see fare increases to attract more drivers and deter some users. The dynamics and equilibria have been studied by Besbes et al. (2021), Hu et al. (2022), Yan et al. (2020). Similar ideas have been explored in car-sharing (Zakharenko, 2023) and bike-sharing (Emami and Ramezani, 2024).

The uncertainty about the future state of the system not only affects its efficiency: in these systems, *unreliability* is often a critical issue, as users cannot know in advance the characteristics of their assignment. In the case of ride-hailing, Bansal et al. (2020) estimate that the demand could increase by 10% if waiting times could be predicted accurately; in ride-pooling, the routes can adapt dynamically to accommodate to the new requests, and this factor deters a more widespread adoption of these systems (Fielbaum and Alonso-Mora, 2020; Zhang et al., 2024; Alonso-González et al., 2020); regarding bike-sharing, in dockless systems the time until finding an available bicycle can be significantly less predictable than when systems are station-based (Beauvoir and Moylan, 2020).

2.2. Spatial equilibrium models

A different stream of research, like ours, proposes simplified models that allow for analytical treatment. Many of these works focus on identifying the types of equilibria that can arise. As a result, the models are typically *static*, describing how the equilibrium looks like rather than how it is reached.

The most relevant model in this context was developed by Castillo et al. (2024), who model the equilibrium between supply and demand. This framework was further extended by Yan et al. (2020) and Fielbaum et al. (2025). This model includes the analysis of the relationship between the number of idle vehicles and waiting time in ride-hailing, but can be easily applied to other similar modes, and it explains why the Wild Goose Chase can be an equilibrium. In this line of models, the spatial dimension is not explicitly considered, or, equivalently, the system is assumed to be homogeneous in space.

The WGC can be quite unintuitive and has relevant operational consequences, which is why several papers have delved into it. Castillo et al. (2024) present the basic equations to show that the same demand rate could be served with the same fleet in two different states, where one of them has a greater waiting time, and explain the corresponding negative feedback cycle. They also show how pricing can mitigate the WGC. Zha et al. (2018) provide mathematical tools to determine which state the system is in. Xu et al. (2020) study the relationship between the matching radius (i.e., the maximum admitted distance between a vehicle and its assigned user) and the appearance of the WGC, showing that a dynamic matching radius can avoid the WGC. Mo et al. (2022) show that the WGC can also appear when traffic congestion is considered and the ride-hailing fleet combines autonomous and human-driven vehicles. Ouyang and Yang (2023) model ride-hailing using differential equations, and propose a mitigation strategy that does not depend on pricing or matching radii, but on dynamic re-assignment of the drivers; their model has been recently extended by Beojone et al. (2025) to study how the system can switch between different equilibria, and analysing how long it stays in each state. Considering general transport systems, Fielbaum (2025) discusses how the WGC is a characteristic problem of modes where the matching between the vehicle and the user is done online (e.g., ride-hailing) as opposed to onsite (e.g., traditional buses).

Other papers have developed explicit physical models to analyse either the matching or the pickup times in the case of ride-hailing (Zhang et al., 2019, 2025; Zha et al., 2016; Yang et al., 2020; Li et al., 2024). These papers have proved useful in developing a theoretical understanding of several aspects, such as scale economies, comparison with traditional taxis, passengers’ competition for vehicles, and the role of walking. Finally, the papers by Kanoria (2025), Balkanski et al. (2023) are closer to ours as they develop mathematical models for spatial matching. Details are discussed later when introducing our model.

¹ Note that drivers are often independent decision-makers (Ashkrof et al., 2022, 2024). Nevertheless, most of the papers reviewed here assume that the platform can make centralized decisions.

3. A general model

Let (X, d) be a metric space (i.e., a set X endowed with a distance function d), and $S^1 = (s_1^1, \dots, s_N^1) \in X^N$ a tuple² representing the initial location of N vehicles in X . Consider two random sequences $(x^t)_{t \geq 1}$ and $(y^t)_{t \geq 1}$ in X , representing the arrivals of users and new vehicles, with i.i.d. distributions Q and P respectively. The process $(S^t)_{t \geq 1}$ goes as follows. At every stage $t \in \mathbb{N}$:

1. A newly arrived user at $x^t \sim Q$ is assigned to the closest vehicle, i.e., to $k^* = \operatorname{argmin}_{k=1, \dots, N} d(x^t, s_k^t)$, facing a cost of $c^t = d(x^t, s_{k^*}^t)$. This implies the assignment is done greedily, an assumption that will be discussed later in Section 4.1.
2. The assigned vehicle at $s_{k^*}^t$ is no longer idle and is replaced by a newly arrived vehicle at the random location $y^t \sim P$, so that the new vector of vehicle positions is $S^{t+1} = (s_1^t, \dots, s_{k^*-1}^t, y^t, s_{k^*+1}^t, \dots, s_N^t)$.

Typical examples of X are the nodes of a graph with d the length of the shortest paths, or a subset of $\mathbb{R}$ or $\mathbb{R}^2$ (as in Fig. 1) using Euclidean distance. The case of $X = [0, 1]^d$ with Euclidean distance has been previously proposed by Kanoria (2025) under different assignment rules, focusing on the asymptotic values of c^t . Borovkov and Odell (2007) studied a generalised version of this problem in $\mathbb{R}^d$, where the probability of selecting driver y^t is an increasing function of the area of its dominance zone (with our case being proportional), although without considering a related cost function.

We remark that in this model, the WGC cycle cannot appear: the new vehicle appears immediately once the previous one gets assigned, keeping the size of the available fleet constant. By these means, for now we focus solely on the IGD.

The sequence $(S^t)_{t \geq 1}$ describes the evolution of the positions of the N drivers and is an homogeneous Markov chain with state space X^N . This state space is finite if X is finite, countable if X is countably infinite, and uncountable if X is uncountable. Throughout the paper we consider both discrete (finite) and continuous (uncountable) cases and clarify when needed. For now, because Markov chains are much simpler when the state space is finite, we include this extra assumption for X . This is largely harmless from an interpretation perspective, as any transport scenario will ultimately involve a finite (potentially discretised) set of feasible locations. Nevertheless, in Section 4 we analyse the one-dimensional circle considering both discrete and continuous cases, whereas in Section 5 we present visualisations on the continuous two-dimensional square. Also, in the Appendix we provide some theoretical support for models with a continuous and compact ground set $X \subseteq \mathbb{R}^d$.

In this discrete scenario for X , we assume that users and drivers may appear anywhere so that for every location $u \in X$ we have $Q(x^t = u) > 0$ and $P(y^t = u) > 0$. This implies that the Markov chain is:

1. *Irreducible*: This property means that if $A, B \in X^N$ are two N -tuples in the state space X^N , there exists some integer k such that $\mathbb{P}(S^{t+k} = B | S^t = A) > 0$. In other words, regardless of the current state of the system, any other state can be reached in finite time. In our setting, this can always be achieved in N steps. Indeed, let us show how the vector of drivers can evolve from $S^t = A$ to $S^{t+N} = B$ with positive probability. We begin focusing on iteration t , where two things need to happen. First, that the new user x^t appears exactly in A_1 so that it is matched with the driver S_1^t . Second, that the new driver y^t appears in B_1 . We note that this has positive probability equal to $Q(x^t = A_1) \cdot P(y^t = B_1)$. If both things happen, then we have that $S^{t+1} = (B_1, A_2, A_3, \dots, A_{N-1}, A)$. In the subsequent $N-1$ iterations, the same would happen but with the remaining coordinates one by one, until reaching the vector B at iteration $t+N$.
2. *Aperiodic*: This follows from the fact that for all $A \in X^N$ we have $\mathbb{P}(S^{t+1} = A | S^t = A) > 0$. The latter is true by a similar argument as above: it suffices that both the user x^t and driver y^t appear where the first driver used to be, i.e. $x^t = y^t = A_1$.

As a direct application of the convergence theorem for finite ergodic Markov Chains, these two properties imply the following lemma.

Lemma 1. *If X is finite, the chain has a unique invariant probability measure π such that $\sum_{A \in X^N} \pi_A = 1$ and regardless of the starting point $S^1 = s$, we have*

$$\lim_{t \rightarrow \infty} \mathbb{P}(S^t = A | S^1 = s) = \pi_A.$$

Note that the final equilibrium is not a specific location tuple of the drivers, but probabilities over all possible location tuples. We refer to this ergodic equilibrium as the *IGD Equilibrium*. Using Lemma 1, standard results from Markov Decision Processes imply the following theorem (Puterman, 2014):

Theorem 1. $\mathbb{E}[c^t]$ converges as $t \rightarrow \infty$.

Under mild regularity assumptions, these results can be extended to the case where X is a compact subset of $\mathbb{R}^d$ and the probabilities governing users and drivers have a strictly positive density with respect to the Lebesgue measure. The proofs in this case are more technical and follow by combining results from Meyn and Tweedie (2012) and Lasserre (2000), which ensure the convergence in total variation of the distribution of the chain towards a unique invariant measure³. The details are presented in the Appendix.

² Throughout the paper, we use superscripts to denote iterations, and subscripts to denote coordinates of a vector.

³ A relevant difference lies in the convergence rates. While the finite case exhibits a geometric rate of convergence (Levin and Peres, 2017), for uncountable state spaces convergence (or mixing) rates can only be established under more stringent conditions (Meyn and Tweedie, 2012), such as the existence of petite sets and geometric drift. These conditions are not guaranteed to hold for our model in continuous space.




Fig. 2. Two interpretations of the metric space, either as a circle (left), or as the $[0, 1]$ interval where 0 and 1 are identified as the same.

Simplifications: Before studying some specific settings for X , let us remark that our model contains a number of simplifying assumptions: the fleet size is kept constant, the demand and supply are completely independent, and there is no rebalancing (which is relevant for ride-hailing but could be impossible for other transport systems as discussed in Section 1). As such, our paper is similar to the ones by, e.g., Castillo et al. (2024), Beojone et al. (2025), Li et al. (2024), in that it considers a stylised representation rather than a fully realistic and detailed one (with the crucial difference that we focus on the spatial heterogeneity of the drivers instead of assuming a uniform one).

All of these simplifications are needed to keep a tractable model that admits analytical results. We recall that in Section 6 we run simulations under more realistic settings to complement the theoretical analysis, including the Manhattan network with real-world demand.

4. The one-dimensional circle

In order to get a better intuition of the dynamics, as well as various theoretical results, we now focus on a simple case, which is still challenging to analyse. We consider the interval $[0, 1]$, and $d(u, v) = \min(|v - u|, 1 - |v - u|)$, representing that the points 0 and 1 are merged, or that the process takes place in a circle of length 1 (just the border). We assume both drivers and users appear following a uniform distribution and independently. Throughout this section, we will detail when we are considering the full interval (continuous), or a discretisation where $X = \{M/N : M = 1 \dots, N\}$.

Note that we can assume that drivers are sorted, simply by reassigning the subindices. That is, we assume $s_1^t \leq s_2^t \leq \dots \leq s_N^t$, and thus define $\ell_k^t = d(s_k^t, s_{k+1}^t)$, with the convention that $s_{N+1}^t = s_1^t$. In plain words, ℓ_k^t is the length of the interval I_k^t between the two consecutive drivers s_k^t, s_{k+1}^t . Note that ℓ is a function of s , but we omit that in the notation to keep it lighter. This is all illustrated in Fig. 2. These intervals constitute the fundamental object in this single-dimensional model, and the Markov chain can be described in terms of $\ell^t = (\ell_1^t, \dots, \ell_N^t)$ as the state.

Let us focus now on the continuous setting. In this case, we can easily compute the expected value of the costs given the positions of the drivers. First, the probability that the new passenger appears in the k -th interval I_k^t is exactly ℓ_k^t . Second, it is easy to see that its expected cost conditional on appearing in the k -th interval is $\ell_k^t/4$. Therefore, defining the function $V(\ell) = \sum_{k=1}^N \ell_k \cdot \ell_k$ for any vector $\ell = (\ell_1, \dots, \ell_N)$, and conditioning on the possible intervals where the new user can appear, it follows that the expected cost at stage t is

$$\mathbb{E}[c^t | \ell^t] = \sum_{k=1}^N \ell_k^t \cdot \frac{\ell_k^t}{4} = \frac{V(\ell^t)}{4}. \quad (1)$$

Note that $V(\ell)$ is minimised when $\ell_k = \frac{1}{N}$ for all k . That is, the better distributed the drivers, the lower the expected cost. This is consistent with the intuition that an equidistant distribution of the drivers is the optimal situation.

In order to understand the dynamics of the system, let us analyse what is expected to happen in the next iteration given the current one. Concretely, given $\ell^t = \ell$ we analyse the expected cost in $t + 1$. When a driver is assigned and thus removed, a new *merged* interval appears, as illustrated in Fig. 3. Two cases can happen, as the new driver might appear in this same merged interval or in a different one. A careful computation of both cases allows us to compute the function $\Delta V(\ell)$, defined as the expected change in $V(\cdot)$ from ℓ to the next random state ℓ' that arises after the arrival of a user and the replacement of the assigned vehicle by a new vehicle at a random position.

Lemma 2. *In the continuous case:*

$$\Delta V(\ell) := \mathbb{E}[V(\ell') - V(\ell) | \ell] =$$




Fig. 3. When a new user arrives, the two intervals that surround the assigned vehicle s_{i+1} get merged.




Fig. 4. In the left, two intervals had the same length. A new user and a new driver appear, each on one of these intervals. As a result, the interval where the user appeared merges with another one and becomes wider, while the interval where the driver appeared becomes subdivided. The “After” situation is therefore worse than the “Before”.

$$\underbrace{\sum_{k=1}^N \frac{1}{3} [4\ell_k \ell_{k+1} - \ell_k^2 - \ell_{k+1}^2] \frac{(\ell_k + \ell_{k+1})}{2} (\ell_k + \ell_{k+1})}_{(*1)} + \underbrace{\sum_{k=1}^N \sum_{j \neq [k, k+1]} 2 \left[\ell_k \ell_{k+1} - \frac{\ell_j^2}{6} \right] \frac{(\ell_k + \ell_{k+1})}{2} \ell_j}_{(*2)} \quad (2)$$

In Eq. (2), the term (*1) represents the case where the new driver appears in the merged interval, and (*2) when it appears in a different interval. The expectation is taken over all possible appearances of the new user and vehicle. The proof of Lemma 2 is provided in the appendix.

Eq. (2) can be used to prove one of the fundamental insights of this paper, namely that **a fully random distribution tends to worsen under the IGD**. To be precise:

Theorem 2. *In the continuous case, if the vector S follows a uniform distribution (each of its coordinates is uniformly and independently drawn from $[0, 1]$), then*

$$\mathbb{E}_{\ell} [\Delta V(\ell)] > 0 \quad (3)$$

Where $\mathbb{E}_{\ell}[\cdot]$ represents that the expectation is taken over ℓ .

The proof is provided in the appendix. Its core is to show that:

- (*1) is zero. This has a direct interpretation. The merged interval is everything that matters for the change, as all the other terms in $V(\ell)$ will remain the same. The merged interval used to be divided by a random driver and now is again randomly divided by a new driver, which is why the expected change is zero.
- (*2) is positive, which can also be interpreted. Ideally, the new user would appear in a short interval and the new driver in a long one. That would result in an elongation of the short interval and a subdivision of the long one, improving the overall situation. The opposite case, where a short interval is subdivided and a long interval is further increased, would be worse.

To be neutral, let us focus on the case where both intervals were of the same size. In that case, the interval of the new user is merged with a different one, increasing its length, whereas the interval of the new driver is divided into two smaller ones. In total, the distribution of the interval becomes more heterogeneous, i.e., worse. This is illustrated in Fig. 4.

A crucial point is that the reasoning behind steps (*1) and (*2) follows the exact same logic illustrated in Fig. 1. Theorem 2 suggests that the IGD limit distribution should perform worse than a fully random distribution of drivers, an observation that is empirically confirmed below. This outcome stems directly from the IGD dynamics: when a user appears, it often leads to the creation of a long interval without drivers, as discussed when analysing (*2). Such long intervals have a significant effect on $V(\ell)$ because (i) they have a high probability of containing a future user, and (ii) users appearing within those intervals incur a high expected cost.

The result from Theorem 2 also holds empirically in the discrete case. However, providing a formal proof is significantly more challenging, as the closed-form expression used for the probabilistic distribution of the ℓ_i is only valid in the continuous setting. In




Fig. 5. Comparison of the IGD with having the drivers randomly or uniformly located.

the discrete case, Eq. (1) requires a slight correction –specifically, subtracting $O/4M^2$, where O is the number of intervals with odd length. This correction is always small, as it is bounded above by $1/4M$.

Theorem 2 suggests a trend, namely that the first step after a fully random start makes the drivers' distribution worse. However, this theorem does not imply that the limit is necessarily worse. This is now explored empirically. In Fig. 5, we compare the IGD against two benchmarks, namely the static expected cost if drivers are either random (red curves) or uniformly distributed (green curves). On the left, we show, for different numbers of drivers N , the limit of $\mathbb{E}[c']$ as $t \rightarrow \infty$, that is, the average cost faced by users once the drivers distribution has converged to the IGD limit. This is computed through Monte Carlo simulations. For the benchmarks, we directly compute the expected cost a user would face when the drivers either follow a fully random (i.i.d.) distribution or are uniformly distributed. Therein, it becomes clear that IGD yields significantly worse results than the two benchmarks. While the random and uniform curves decrease proportional to $1/N$, our simulations suggest that the IGD equilibrium decreases at a rate that lies between $\log(N)/N$ and $\log^2(N)/N$.

The dynamic component can be seen on in the right panel of Fig. 5, where we fix $N = 50$. For IGD, we compute (through Monte Carlo simulations) the evolution of the expected cost faced by each user if drivers start uniformly distributed. The increasing trend until convergence is also evident. Theorem 2 can be observed by noting that the limit of the blue curve is significantly greater than what would be achieved if the drivers remained randomly distributed (red horizontal line).

To compute the blue curve, we simulated the process 10,000 times and averaged the resulting costs. To verify that $\mathbb{E}[c']$ had effectively converged, we compute a moving average of the blue curve using a window of five consecutive passengers, where each point already corresponds to the average over the 10,000 simulations. We then identify the first passenger for which this moving average changes by less than 0.5% over the next 15 consecutive passengers. This criterion is satisfied after 1799 passengers, indicating that the expected cost had stabilised. It is worth recalling that convergence is not expected in a single simulation. For example, it may happen that users repeatedly appear very close to a driver location, yielding almost zero cost. Theorem 1 guarantees convergence of the expected costs (represented here as the average of 10,000 simulations), not of individual sample paths.

We can further illustrate the source of why the IGD equilibrium is different than a random distribution through the following hypothetical example. Imagine that instead of assigning the users greedily, we do it randomly. In this case, the users' position would play no role, or in other words, we replace one random driver with another random driver. Therefore, that system would converge to a fully random distribution. The difference between assigning greedily and randomly lies precisely on the gaps: in a greedy assignment, the drivers that already cover a large area have a greater probability of being selected, thus increasing their area even further. This is the ultimate reason why the IGD equilibrium is worse than random.

A clarification is relevant in the example above. Assigning randomly would obviously yield greater costs, because every assignment would be very costly: If the first t iterations are assigned randomly, and the passenger $t + 1$ is assigned greedily, then the latter would be benefited since it would face a random distribution of the drivers rather than one approaching the IGD equilibrium. However, the first t users would likely be worse off.

Similarly, there is a noteworthy analogy with the well-known *inspection paradox* (Stein and Dattero, 1985). This refers to a statistical bias that arises when the sampling method is related to the unit being sampled. The most famous example in transport is the difference between time-mean-speed TMS and space-mean-speed SMS (Knoop et al., 2009). In TMS, fast vehicles are sampled more frequently due to their higher speed, resulting in an overestimation of the true mean speed SMS. In the IGD, we are not sampling or estimating; instead, we are choosing a driver in every iteration. The probability of selecting a given driver is greater if the driver covers a larger dominance zone, which in turn increases the resulting cost. This analogy extends beyond the qualitative mechanism. In traffic




Fig. 6. The policy ρ assigns u to s_a , so ρ' assigns its neighbour s_n instead.

flow theory, $TMS = SMS + \text{Var}(v)/SMS$, i.e., the gap between the two measures is entirely determined by the variance of vehicle speeds $\text{Var}(v)$. Similarly, Eq. (1) implies that $\mathbb{E}[c^t | \ell^t] = \bar{c} + \text{Var}(\ell^t)/16\bar{c}$, where $\bar{c} = 1/(4N)$ is the cost under uniformly spaced drivers. Therefore, in both settings, heterogeneity generates a systematic bias away from the ideal benchmark. That is, similar to the difference between TMS and SMS, the resulting cost is greater than if every driver would have the same probability of being chosen.

4.1. Some theoretical results on optimal assignments, and a comparison with the greedy rule

So far, we have assumed a greedy rule when assigning users to drivers. In this subsection, we argue why this is a reasonable matching mechanism, and more specifically, why a better one would not prevent the IGD from happening⁴. As we are discussing the (sub)optimality of matching mechanisms, it is worth noting that our system can be described as a Markov Decision Process (MDP), where we are searching for a *policy*, i.e., a function $\rho : [0, 1]^{N+1} \rightarrow \{1, \dots, N\}$ that takes as inputs the current position $(s_1, \dots, s_N, u)$ of the drivers and the user, and returns which driver is assigned to the user. The greedy rule is one possible such ρ . The question is which is the optimal ρ , i.e., the one minimising the limit of the average cost when $t \rightarrow \infty$. In the *continuous* case, where drivers and users can appear in any position, an optimal policy might not exist. In the *discrete* case, where the circle is subdivided into M subintervals of length $1/M$ each, there is a finite number of policies, so one of them must be optimal. We now describe our results.

4.1.1. General results about the optimal policies

Theorem 3. *If there are two drivers, greedy is optimal. This is valid in the discrete and in the continuous case.*

Proof. We note that when a user appears, the system's status after the assigned driver disappears is equivalent regardless of which driver was assigned. Namely, it has one driver located somewhere in the circle, and another driver will appear randomly. This implies that the future costs of the system are independent of the current assignment, and therefore a greedy assignment is optimal. $\square$

The previous argument does not hold when the number of drivers exceeds two. Even with just three drivers, if two are close to each other but the third is not, under certain conditions it may be preferable to incur a slightly higher cost now by assigning one of the two drivers that are clustered together. However, there is a remarkable result that is valid for any N , namely, that an optimal assignment rule should always match to one of the two *neighbouring* drivers.

In order to formalise this, let us consider a user u and a vector of drivers s . We define $R(u, s)$ as the first driver to the right of u , and $L(u, s)$ as the first driver to the left. If there is a driver exactly at the same position as u (which might happen in the discrete case), we have $R(u, s) = L(u, s) = u$. We call $R(u, s)$ and $L(u, s)$ the “neighbours” of u .

Theorem 4. *Consider the discrete case, where the circle is partitioned into M sub-intervals of length $1/M$ each. Then any optimal policy will always assign a user u to either $R(u, s)$ or $L(u, s)$. As a consequence, at least half of the actions of any optimal policy are greedy. These properties hold regardless of whether the arrival probabilities of users and drivers are uniform or not.*

Proof. Consider a policy ρ that does not always assign to a neighbour. We will show that ρ can be improved, implying that it is not optimal. We will construct a *non-Markovian* policy, meaning one that depends not only on the current state but also on past states. This does not pose a problem, since in any finite MDP, the expected value achieved by any non-Markovian policy can also be attained by a Markovian policy (Puterman, 2014). In other words, it suffices to show that there exists a non-Markovian policy that is better than ρ .

This non-Markovian policy ρ' operates as follows. It has a passive/active label that begins as passive. It assigns exactly the same as ρ for the initial iterations (potentially none) until the first time ρ does not assign to a neighbour. Let us denote by s_a the driver that was assigned by ρ , and by s_n the neighbour that is *between* s_a and u . That is, s_a is to the left of u if the distance between s_a and u is achieved by u moving to the left, and vice-versa; if the distance is exactly $1/2$, we can take s_n to be any of the two neighbours. In this iteration, ρ' assigns to s_n instead of s_a . This is shown in Fig. 6. We now say that ρ' is active.

Note that after this assignment, ρ and ρ' do not have the same set of drivers, but they differ by one: s_n would be available if using ρ , but we have s_a available instead. In the next iterations (potentially none), ρ' copies what ρ would do if having s_n . This is stopped

⁴ The literature reports related results in similar settings. Let us denote by $c(N)$ the limit cost for $\mathbb{E}[c^t]$ when we have N drivers. In the continuous case, Kanoria (2025) proves that there exists K such that no policy can achieve $c(N) \leq K \log(N)/N$. Note that if drivers would remain uniformly distributed, then $c(N) = 0.25/N$. Numerically, we find that greedy performs better than $\log^2(N)/N$ (more specifically, that $\frac{c(N)}{\log^2(N)/N} \rightarrow 0$ when $N \rightarrow \infty$), i.e., greedy is not worse than $\log(N)$ times the uniform result, which is obviously better than what can be achieved with an optimal policy. This resonates with the theoretical results for the greedy algorithm by Balkanski et al. (2023) in a similar setting, but when all drivers are available from the beginning and disappear one by one, and where they prove that the ratio between greedy and the optimal solution is not worse than a constant.

when ρ would assign to s_n , in which case ρ' assigns to s_a . That is:

$$\rho'(s, u) = \begin{cases} \rho(s \cup \{s_n\} \setminus \{s_a\}, u) & \text{if } \rho(s \cup \{s_n\} \setminus \{s_a\}, u) \neq s_n \\ s_a & \text{Otherwise} \end{cases} \quad (4)$$

When the latter happens, ρ' becomes passive and the cycle starts again. Note that in a whole cycle, i.e., since ρ' becomes passive (or the first iteration) until it becomes passive again, ρ and ρ' have the same cost in every iteration but two. Denoting by u_1 and u_2 the users involved in these two iterations, we have that the cost using ρ would be $d(u_1, s_a) + d(u_2, s_n)$. And then

$$d(u_1, s_a) + d(u_2, s_n) = d(u_1, s_n) + d(s_n, s_a) + d(u_2, s_n) \geq d(u_1, s_n) + d(u_2, s_a) \quad (5)$$

where the first equality in Eq. (5) holds because $s_n \in (u_1, s_n)$, and the inequality is just the triangular inequality. Note that the right hand side $d(u_1, s_n) + d(u_2, s_a)$ is exactly the cost yielded by ρ' . In other words, ρ' yields a cost that is lower or equal than ρ in every cycle. Moreover, there is a positive probability that the inequality is strict: for instance, if $u_2 = s_a$ the inequality would be strict, and this has a positive probability. It suffices that all new users and drivers appear in the same position as s_a from the moment ρ' becomes active until it becomes passive. This completes the proof that ρ' is better than ρ , and as argued above, we conclude that any optimal policy must always assign to a neighbour.

We now show why at least half of the actions are greedy. To do this, note that if d_1 and d_2 are two consecutive drivers, then we can divide the interval $[d_1, d_2]$ into two subintervals $[d_1, z]$, $[z, d_2]$ for some z , such that if a user appears in $[d_1, z]$ it would be assigned to d_1 , and a user in $[z, d_2]$ would be assigned to d_2 (to ease the notation, we are using $[a, b]$ to describe all the discrete slots that belong in that interval). To see why such a z exists, suppose instead there exist users $u_1 < u_2$ with u_1 assigned to d_2 and u_2 assigned to d_1 . Then, swapping these assignments would keep the probabilistic distribution for the next iteration unchanged, and diminish the expected costs in this iteration, which cannot happen in an optimal policy. This division into two subintervals directly implies that at least half of the users' positions would be assigned greedily. $\square$

Let us remark that a similar result holds in the continuous case. While there might be no optimal policy, any policy that sometimes does not assign to a neighbour can be improved doing exactly the same as we did in the proof of Theorem 4, i.e., building ρ' . However, this change would require admitting Non-Markovian policies, as the result mentioned above from Puterman (2014) is not necessarily valid when the state space is uncountable.

A natural question is whether a similar Theorem 4 holds in other domains? The very notion of “neighbours” only makes sense in one dimension. On the other hand, during the proof, the fact that we were considering neighbours was only used in Eq. (5), specifically in the first equality, as we knew that the shortest path from the assigned driver s_a to the user u_1 passed through the neighbour s_n . This means that in other spatial domains, whenever the same situation happens, we can discard s_a as an assignable driver. In Euclidean regions in two or more dimensions, this would typically have probability zero, as it requires all the points to lie on a one-dimensional line. But in graphs, this can help. All together, this implies the following corollary:

Corollary 1. Consider $X = (V, E)$ a graph (directed or undirected) where every arc has an associated distance. Take an optimal policy ρ . If the user is in $u \in V$, and there are vehicles in $v_1, v_2 \in V$, with v_1 lying within a shortest path from v_2 to u , then ρ will not assign v_2 to u .

While Corollary 1 seems too specific, there are some types of graphs where it can be useful to reduce the number of potential actions. Two notable examples are:

1. Take G an undirected grid, and u a node there. If v_1 and v_2 are both to the north and east of u , but v_2 is further in both dimensions, then v_2 can be discarded. Obviously, the same happens if replacing “north” with “south”, and/or “east” with “west”.
2. Take G a tree, and define its root to be where the user is located. If v_1 is an ancestor of v_2 , then v_2 can be discarded when deciding which vehicle to assign.

Theorem 4 has several implications. The first one is about reinforcing the intuition that greedy is reasonable. Indeed, while we should not always assign greedily, we should not do something too different either. We need to always select one of the two neighbours, and greedy selects the closest one.

There is also a direct relationship between Theorem 4 and the IGD. If we use the optimal policy, and a user arrives in a long gap, this long gap will increase because we will remove one of the gap's borders. This can be seen using Fig. 3 again: the optimal policy would assign either s_i or s_{i+1} to the new user, and in any case the long interval $[s_i, s_{i+1}]$ would get merged with another interval.

On the other hand, Theorem 4 significantly narrows the policy space by showing that, in each state, at most two assignments need to be considered instead of all N possibilities. This reduction makes it feasible to explore optimal policies directly. In the next subsection, we leverage this result to compute the optimal policy for small values of M and N , compare its performance with the greedy policy, and confirm that the IGD phenomenon still emerges, even under optimal decision-making.

4.1.2. Explicit optimal policies for small instances of the discrete case

In any finite MDP, the optimal policy can be derived through a Linear Program (Puterman, 2014). This LP has as many variables as the number of possible states, and as many restrictions as the number of states multiplied by the number of actions for each state, which in our case is lower or equal than two (thanks to Theorem 4). In practice, we can solve this problem for some combinations of $M \leq 15$, $N \leq 7$.

Let us first fix $M = 8$, and summarise the results in Table 1. For $N = 3, \dots, 7$, we show two results:

Table 1

Comparison of greedy and the optimal policy for $M = 8$ and different number of drivers.

Drivers	3	4	5	6	7
% increased costs	2.13	4.07	6.05	7.9	9.47
% non-greedy actions	0	1.8	2	2.1	2.2

Table 2

Comparison of greedy and the optimal policy for $N = 4$ and different values of M .

M	5	6	7	8	9	10
% increased costs	6.91	5.3	4.66	4.2	3.74	3.4
% non-greedy actions	0	0	0.95	1.82	2.02	1.96




Fig. 7. Comparison of the evolution of the system using the optimal policy versus greedy.

- In the top row, we have the extra cost yielded by greedy in percentages (computed through Monte Carlo simulations), compared to the optimal policy. While these values increase, they do it concavely, and moreover, the percentage is never over 10%. This suggests that greedy is indeed competitive.
- We then show which percentage of the states require a strictly non-greedy decision under the optimal policy, i.e., assigning to the furthest neighbour. While we know theoretically that this percentage is below 50%, we observe that the actual numbers are very low. We remark that having 0% in the case of 3 drivers occurs only because $M = 8$, as for greater values of M we do find strictly positive percentages.

To further understand these numbers, we now show in Table 2 the results for $N = 4$ and varying M . Crucially, we observe that greedy becomes more competitive as M increases.

Tables 1 and 2 reinforce that greedy is not so far from the optimum. We now show that even if we use the optimal policy, the IGD also appears. Let us first consider $N = 5$, $M = 9$, and analyse the temporal evolution of two indices, starting with drivers randomly located: (i) the cost and (ii) the variance of the intervals. We know that in the case of greedy these two quantities are correlated (Eq. 1), but this is not necessarily true for the optimal policy. Nevertheless, the variance still acts as a proxy for the IGD: If the IGD exists, then the longer intervals will become even longer, raising the variance. This is shown in Fig. 7, where we show the average of 10,000 repetitions of the experiment, and we consider 40 iterations with the same criteria (using moving averages) as done for Fig. 5. We observe that:

- While greedy is indeed worse than the optimal, both curves are also worse than the drivers being randomly distributed, represented by the starting point of both curves. That is, even when using an optimal policy the **IGD equilibrium is worse than a fully random distribution**.
- The reason for this increased costs stems from the variance. Even in the optimal policy, the variance increases significantly before converging.

Table 3

Expected next max-gap given current max-gap a , for optimal and greedy policies, using $M = 8$ (and omitting the division by 8).

a	2	3	4	5	6	7	8
W_a Optimal	3.5	3.77	4.18	4.71	5.2	5.65	6
W_a Greedy	3.5	3.84	4.34	4.79	5.25	5.66	6

- We can observe that the shape of the two red curves are almost identical, which is expected as the greedy costs depend directly on the variance. This is no longer true for the optimal policy, where there is some correlation but the curves are different.
- In the optimal policy, there is a clear drop at the beginning, after which the costs start to increase. This suggests that there is some smart way to respond to the initial random distribution of the drivers, but this smarter response does not persist for the subsequent iterations.

We now study the appearance of the IGD in a more direct way. As we have few drivers, there cannot be many long intervals. So we focus on the longest one. To be precise, we take the following steps:

1. We consider a policy ρ that can be optimal or greedy.
2. For this policy, we compute the transition matrix between the positions of the drivers, i.e., $\Pi_{j,i} = P(s_1^t = i_1/M, \dots, s_N^t = i_N/M | s_1^{t-1} = j_1/M, \dots, s_N^{t-1} = j_N/M)$. Note that this can be done trivially because we know the assignment rule ρ .
3. We compute the stationary probability π of the matrix Π , and define $W_{a,b} = P(I(s^t) = b | I(s^{t-1}) = a)$, where $I(s)$ is the longest gap between two consecutive drivers, when the drivers are described by s . We then compute $W_a = E_b(W_{a,b})$, where the expected value is taken over b using the probability distribution π .

In plain words, W_a measures the expected longest gap in the next iteration, when the current longest gap is of length a . In Table 3 we report the results for $M = 8$, $N = 4$, where the ideal distribution is having a longest interval of $2/8$. To keep the Table simple we do not divide by 8. The main observations from this Table are:

- The optimal policy is indeed better than greedy by keeping a better drivers' distribution. In every row, the optimal policy achieves a strictly lower value than greedy, with the exception of the two extreme cases where the assignment is trivial so the values are the same for both policies.
- However, the differences are quite mild. There is not much room from improvement compared to using the greedy assignment policy.
- The IGD can be observed when $W_a \geq a$, i.e., when the longest gap tends to further increase. Obviously, this cannot always hold, as in the extreme case where all drivers are in the same position, the longest gap will reduce. However, (i) This condition holds for exactly the same values of a for both policies, and (ii) It even holds when $a = 4$. That is, if all the drivers are concentrated in one half of the interval, the odds are that this gap will increase further in the next iteration, which constitutes a prototypical example of the IGD: it is just too probable that the user will arrive within this gap, and it will be assigned to one of the gap's borders regardless of the policy. This is not yet compensated by the arrival of the next driver.

In all, Table 3 shows that although the optimal policy does a better job than greedy in preparing for the future, this improvement is very minor and, more importantly, cannot prevent the IGD.

4.2. Non-uniform arrivals

So far, most of this section has assumed that passengers' and drivers' appearances follow a uniform distribution in $[0, 1]$. However, in real transport systems there are usually hotspots where users and drivers tend to concentrate. In this subsection we discuss how this generalisation would affect our analysis and results.

We first note that the overall dynamic remains almost the same. Instead of analysing the length of the intervals, we now need to focus on their probability. That is, an interval (gap) with a greater probability would still be more likely to attract the new users, making the interval to grow. Moreover, we know that Theorem 4 remains valid. That is, even if we used an optimal policy, we would still assign one of the neighbours so that the gap increases.

However, there is one relevant difference: the expected cost for a user arriving in an interval is determined by its length, whereas the probability of a user arriving there is governed by the interval's probability mass. In the uniform case, both quantities are proportional to the length, but this is not true when the probabilities are not uniform (i.e., Eq. 1 is no longer valid in this general case). This introduces a random effect on the IGD: consider a situation where all the intervals have the same probability but some begin merging due to the IGD; if these were low-demand intervals, thus lengthy ones, the impact of the IGD is even worse, but the contrary would happen if merging short intervals.

In all, the IGD still operates, but whether its impact is worse than when probabilities are uniform is unclear. We study this through simulations. We consider three different probability distributions. The first two give more probability to the interval $[0, 1/2]$ (the high-demand zone) than to $[1/2, 1]$ (the low-demand zone). The third one has a traditional triangular pdf. For each distribution, we run Monte Carlo simulations with 10 drivers and 10,000 repetitions. The number of passengers (211) follows the same convergence




Fig. 8. Effects of the IGD dynamics when the arrival of users and drivers is not uniform in space.

criterion used for Fig. 5, namely, the point at which the moving average has stabilised for all three distributions. The initial set of drivers is distributed randomly following the same corresponding probability distribution, so we can compare the final equilibrium with the random one.


The trends in Fig. 8 clearly show that the IGD equilibrium is worse than the original setup with drivers randomly located. This is true for all three curves. Moreover, the comparison between the two skewed cases reveal that the more concentrated the drivers, the worse the impact of the IGD, as the relative increase in costs is significantly greater.

In all, the conclusions remain similar to what was discussed in the previous subsections. As a complement, let us remark that in Section 6.2 we run simulations using a real-world dataset from Manhattan, where the origins and destinations of the users are not uniform.

5. Visualisations on the unit square

Let us consider now the (continuous) unit square in two dimensions. Conceptually, the system remains quite simple. Moreover, as this is a compact subset of $\mathbb{R}^2$, we know that Lemma 1 and Theorem 1 apply; that is, the probability distribution of the drivers' position converge, as well as the expected costs faced by the users. On the other hand, we now explain why an analytical treatment seems beyond reach. Note that what we described as the dominance zone in Section 1 corresponds exactly to the Voronoi diagram induced by the vehicles (De Berg, 2000). This means that in the unit square the dominance zone of every vehicle is a convex polygon, and their geometry now plays a role. Crucially, the probability of assigning a given car is given by the area of its corresponding dominance zone, which is exactly the same as in the single-dimensional case (using the length instead of the area). But the expected cost is now quite different: if the area is regular (e.g. an hexagon or a square), the expected cost will be proportional to the polygon's side, i.e., to the square root of its area. But if the polygon is more irregular, its area becomes less informative.

The complexity described in the previous paragraph hinders the development of a theoretical analysis. But on the other hand, the square is an ideal environment for visualisations, which can be very helpful for transport analysis (Andrienko et al., 2017). Let us first remark that the overall dynamics remain quite similar. This can be explained by three key observations. First, it is still desirable to maintain dominance zones of comparable size. Second, larger zones have a higher probability of attracting users. Third, although a zone's area is no longer directly proportional to the expected user cost (as in the one-dimensional case), it does remain positively correlated, since larger zones tend to include points that are farther from the assigned vehicle. This is the version of the IGD dynamics in the square.



(a) Random distribution





(b) Initial distribution ($t = 0$)(c) $t = 100$ (d) $t = 200$ (e) $t = 300$

Fig. 9. Evolution of the Markovian system in the unit square with 100 vehicles. Triangles represent the vehicles.

In order to develop the visualisations, we consider 100 vehicles that begin at the centres of a perfect partition of the square into 100 subsquares. We then run 300 iterations. The resulting process is shown in the video attached to this submission⁵. We further

⁵ The video can be accessed at youtu.be/2caWJjQyvXg




Fig. 10. Evolution of the variance of the dominance zones in the unit square.

visualise the results in Fig. 9, where vehicles are represented as triangles. First, Fig. 9a shows a benchmark of sorts, namely one realisation of the Voronoi diagram with 100 points randomly distributed in the square. Figs. 9b–9e show the evolution of the system, from the perfect initial to the distribution to the final one, showing the state of the system every 100 iterations. The main observations from Fig. 9 are:

- The system's distribution keeps worsening over time. While already at iteration 100 the heterogeneity between the zones' areas is evident, this situation is significantly sharper in $t = 200$ and $t = 300$. Crucially, some very large zones start to appear.
- While the random distribution is also quite heterogeneous, its zones are clearly less distinct than those yielded by the IGD. Fig. 9a shows no huge zones and relatively few very small ones, compared to Fig. 9e.

This visual analysis can be quantified. To do this, in every iteration we compute the area of every dominance zone, and then calculate the variance of the resulting vector. The results are presented in Fig. 10, where we benchmark against the expected variance if the drivers are distributed randomly (computed through a Monte Carlo simulation). It becomes evident how the variance increases, and how it rapidly exceeds the random one. We note that other simulations might reach lower numbers, but in all of them we observe the same increasing trend and a much greater variance than the random case.

It is worth noting that the sub-optimality of greedy assignments has been studied in a two-dimensional context by Hyland and Mahmassani (2018), who compare it to a matching method in which multiple users are assigned simultaneously (this is known as “batch assignment”, further discussed and tested in Section 6.1). They found that while greedy can be significantly outperformed, the performance difference narrows substantially as demand increases.

6. Simulation of a ride-hailing system and the emergence of the WGC

The theoretical model permits a clear analysis, but how transferable is this analysis when going to more realistic scenarios? Can we still observe the dynamic described above, i.e., that the heterogeneity of the dominant areas creates a negative cycle? Moreover, when the number of idle vehicles is not constant, but depends on the actual service times, can we observe a relationship between the IGD and the WGC?

In order to study these questions, we run simulations of a ride-hailing system in two different settings. First, we use a ring-radial network (Section 6.1), where we can control and ensure some regularity in the network, demand, and supply, enabling us to analyse the IGD and WGC in better isolation. Next, we use a real-world dataset (Section 6.2) to observe how these dynamics emerge when combined with the many additional aspects and complexities of real cities. Finally (Section 6.3), we come back to the circular setting but now extend it to consider a variable number of drivers, where we can compare the limit situations with and without IGD.

6.1. Ring-radial network

We consider a ring radial network composed by 45 zones, each of them with 5×7 nodes. This network has been previously utilised for the analysis of on-demand mobility, specifically to study scale economies in ride-pooling (Fielbaum et al., 2023). The network, although with fewer zones, is illustrated in Fig. 11. All arcs are bidirectional and are traversed in half a minute. Requests are created randomly but keeping the distance between origin and destination within a narrow interval (between 5.8 and 6.2 minutes), in order to limit the relevance of randomness. We have one request every 30 seconds, and each request is assigned immediately upon appearing. Note that it is no longer true that every time a vehicle is assigned, another one becomes available, as in the theoretical model. Instead,




Fig. 11. An illustration of the network, and the starting position of the vehicles, used for the simulations.




Fig. 12. Simulation results on the circular grid.

a vehicle is assigned to a request, and it becomes again available after the dropoff. We consider 120 minutes of simulation, with part of the vehicles beginning with a user onboard as the result of a 30 minutes warm-up phase (that starts with one vehicle at the centre of each zone), and the remaining perfectly distributed in the network. The experiment is repeated 100 times and we report the average results in Fig. 12.

Fig. 12 shows three panels. On the left, we show the average waiting time of the passengers, i.e., the exact cost function we have been considering. At the centre, we show the number of idle vehicles. Previous models that do not consider the evolution of the position of the vehicles have frequently assumed an inverted relationship between the two, e.g., $Waiting = Idle^{-\alpha}$ for some $\alpha > 0$ (Yan et al., 2020). Remarkably, at the beginning of our simulations (after a few minutes of adjustment), the waiting time goes up even though the number of idle vehicles also increases. The reason for this apparent contradiction lies precisely on the non-uniformity of the idle vehicles' position (i.e., the IGD), which is reported at the right and measured as follows: for every idle vehicle v , we compute the number of nodes that belong to the dominance zone of v , we calculate the variance of the resulting vector, and normalise.

We have divided the second and third panels of Fig. 12 into three distinct phases:

1. In the first phase, the number of idle vehicles increases, yet the waiting time also increases, indicating that the IGD is the dominant mechanism.
2. In the second phase, the number of idle vehicles decreases while spatial heterogeneity continues to increase. This means that both the IGD and the WGC are acting simultaneously.
3. In the final phase, the system converges to a steady state and the observed metrics stabilise.

Let us now quantify both the equilibrium impact of the IGD, as well as the extra waiting times experienced during the convergence process. To do so, we apply the same convergence criterion described for Fig. 5 (comparing consecutive moving averages), except that the tolerance was increased from 0.5% to 5%. This adjustment was necessary because the circular-grid simulations are computationally more demanding and therefore rely on fewer repetitions, resulting in noisier curves. Using this criterion, convergence was reached after approximately 63 passengers. The average waiting times increased from 1.99 minutes for the first passenger to a long-run value of 3.62 minutes, representing an increase of roughly 82%. During the transient phase, passengers experienced an average excess cost of 0.98 minutes relative to the first passenger. These results indicate that a substantial share of the welfare loss generated by the IGD is already incurred during the convergence process itself, rather than only after the system reaches its stationary regime.

These results remain the same if we use *batch assignment* (Yan et al., 2020; Ramezani and Valadkhani, 2023). That is, instead of assigning a passenger each time it appears, we wait for some time to collect more requests, and then assign them all together. In this case, we assign every six minutes. The idea is that by collecting more information, the assignment decision will be better (less greedy). As our dynamic is related to assignments and how current decisions affect the future state of the system, it is worth analysing if assigning by batches might eliminate (or mitigate) the IGD.

As shown in Fig. 13, the trends are consistent with those observed in Fig. 12. The numbers do change compared to the case where passengers are assigned right away: Waiting times increase, as some passengers have to wait before being assigned, and the number of idle vehicles also increases, reflecting that the distance between users and their assigned drivers does decrease. However, the main feature from Fig. 12 also appears here: at the beginning of the simulations, the waiting time increases (Fig. 13 left) even though the number of idle vehicles increases as well (Fig. 13 centre), which is explained because the distribution of the vehicles gets worse (Fig. 13 right).

Intuitively, batch assignment can improve the quality of individual matches by considering several requests simultaneously, thereby reducing the average distance between passengers and their assigned drivers within a given iteration. However, similar to the discussion about optimal policies in Section 4.1, batch assignments do not consider the created gaps when making decisions. Consequently, assigned drivers still leave behind gaps where future requests are more likely to appear, causing those gaps to grow over time. That is, batch assignment improves trip-level efficiency but leaves the feedback mechanism underlying the IGD essentially unchanged, and the system continues to evolve towards a more heterogeneous spatial distribution.




Fig. 13. Results of the simulations on the circular grid when the assignment is batch-based.

6.2. Simulations on a real-world dataset from manhattan

We now analyse the emergence of the IGD in a realistic context, namely the widely used dataset from Manhattan. We consider a subnetwork of the Manhattan roadmap, first proposed by Zhang et al. (2025), to work over a more homogeneous context by removing the areas with very low demand. This network is composed by 1966 nodes and 4235 edges, as depicted in Fig. 14.




Fig. 14. The subnetwork of Manhattan used in the simulations. Source: [Zhang et al. \(2025\)](#).

We consider one hour of data from the morning peak, which comprises 14,213 requests, and a fleet of 1500 vehicles⁶ Requests are assigned as soon as they emerge. We implement a warm-up phase of 15 minutes, after which the idle vehicles begin their journey “uniformly” distributed: we use a k-medoids clustering method and situate the vehicles in the resulting centres⁷

Crucially, the demand is no longer spatially homogeneous, as some nodes might generate or attract more trips than others. Moreover, the length of every trip might differ. These changes hinder an analytical treatment, but test whether our theoretical results remain valid under this data-informed configuration. The fleet size was chosen to be large enough to ensure that all requests can be served, as otherwise, selecting which ones to reject could bias the results.

We compute the same set of results as for the previous experiments, shown in [Fig. 15](#). In this case, the waiting times present a very large variation; moreover, it frequently reaches zero due to the high density of the vehicles in the network, that implies that sometimes a request would appear in a node that already has an idle vehicle. This is why we include a moving average, computed every 100 consecutive requests, to provide a cleaner interpretation. Crucially, we observe exactly the same trends as in the circular grid: the waiting time tends to increase ([Fig. 15 left](#)) until converging. This is true even at the beginning of the simulations, where the number of idle vehicles increases as well ([Fig. 15](#), disregarding the initial drop due to the inertia from the warm-up phase). The explanation lies in the IGD, as revealed by the increase in the spatial heterogeneity of the idle drivers ([Fig. 15 right](#)).

As mentioned above, the Manhattan experiments differ from the previous scenarios in that demand is now substantially more heterogeneous in space. This makes the similarity between the results obtained here and those from [Section 6.1](#) particularly noteworthy, as it suggests that the observed dynamics are robust to the underlying spatial distribution of requests. One could even argue that spatial heterogeneity strengthens the relevance of the IGD. Indeed, when trip origins are concentrated in specific areas, a gap created by the assignment of a nearby driver is more likely to attract future requests, causing the gap to grow further. At the same time, because destinations are no longer uniformly distributed, drivers tend to reappear in concentrated areas rather than across the entire space, reducing their ability to compensate for existing gaps. As a result, both the demand and supply dynamics may reinforce the spatial imbalances generated by the IGD.

This intuition is also reflected when quantifying the observed effects. Using the same convergence criterion as above, with a 5% tolerance to account for the greater variability of the simulations, convergence was reached after approximately 3625 requests. The average waiting time increased from 0.02 minutes at the beginning of the simulation to a long-run value of 0.27 minutes, representing a more than tenfold increase. During the transient phase, passengers experienced an average excess waiting time of 0.10 minutes

⁶ Formally, we redistribute these requests within 5 hours. This is done because we need a fleet that is large enough to ensure no rejections. This would require approximately 5000 vehicles if not redistributing the requests, implying that almost every node would have a vehicle, and some several, which would hinder the analysis of the dominance zones.

⁷ This clustering method is very appropriate for this context, as it can consider arbitrary distances and returns centres that are nodes in the network, as opposed to k-means that is based on the Euclidean distance and finds clusters but not explicit centres ([Kaufman and Rousseeuw, 2009](#)).




Fig. 15. Results of the simulations on Manhattan.

relative to the initial state. While these waiting times remain small in absolute terms due to the large fleet size and high vehicle density, the relative increase is substantially larger than in the circular-grid experiments.

6.3. An extension of the one-dimensional model

The two previous subsections show a relationship between the increase in the spatial heterogeneity and the appearance of the WGC. However, the IGD and WGC are very entangled in those simulations. To isolate their effects, let us go back to the one-dimensional circular model studied in Section 4, but now considering an extension where the number of available vehicles is no longer constant. The idea is to capture the relationship between the number of idle vehicles and resulting user-driver distance. To do this, let us start with a quick recap of the equations describing the WGC by Castillo et al. (2024).

Consider a fleet of F vehicles that are serving q passengers per time unit. The average pickup time, i.e., the distance between a car and its assigned user, is denoted by t_p . The average in-vehicle time, i.e., the distance between a user origin and destination, is denoted by t_i . By Little's law, we have that

$$F = q(t_p + t_i) + I \quad (6)$$

where I is the number of idle drivers in the system. As discussed in Section 6.1, it is assumed that t_p is a decreasing, and thus invertible, function of I (e.g., $t_p = I^{-\alpha}$). This implies that

$$q = \frac{F - I(t_p)}{t_p + t_i} \quad (7)$$

The crucial observation is that as t_p increases from some lower bound, q first increases as well but then it starts to decrease. This backand behaviour shows that a given demand q can be served with the same fleet F offering two different waiting times.

As thoroughly discussed in the paper, these equations do not consider the spatial heterogeneity of the drivers and its evolution in time. To model a system with both IGD and WGC in the circle discretised into M points, we now introduce a tweak to our previous model. Instead of having a new driver s appearing immediately once the former one is removed, we consider a lag equal to the distance between s and its assigned user (multiplied by M so it belongs in $\{0, 1, \dots, M\}$). We note that we are not including in-vehicle time ($t_i = 0$), because doing so would simply add an average of $t_i q$ busy vehicles to the system that do not affect the rest of the dynamics⁸. As one user emerges at each iteration, we have $q = 1$. If all I drivers were randomly distributed and users are assigned greedily to the closest vehicle, it is easy to see that $t_p(I) = M/4I$, so Eq. (6) would yield $F = I + M/4I$.

To study this system, we run simulations with random demand and with a starting set of I idle vehicles randomly distributed. For the first $M/4I$ iterations, we include one random driver after assignment, representing the drivers that were on their way to pickup a passenger at the beginning of the simulation - a warmup phase of sorts.

⁸ Analogous simplifications are usual in transport analysis. See, for example, Vickrey's bottleneck model as described by Boyles et al. (2020). Having $t_i = 0$ is also a natural representation for the specific case of curbside parking of delivery vehicles described in Section 1, where the lots are either empty (idle), or with a vehicle (waiting for the deliverer to complete the delivery). Complementary experiments without this assumption can be found in the Appendix, where the same trends persist.

We consider two different assignment rules: (i) *Greedy*, i.e., assigning an emerging user to the closest vehicle, and (ii) *No IGD*, where the vehicle is chosen randomly. Crucially, the latter is called “No IGD” because, by assigning randomly, we are removing the very cause of the IGD: with this rule, when a gap becomes longer, there is no feedback cycle that pushes it to increase even further. In both cases the number of idle drivers is no longer fully predictable. In particular, it might happen that for some iterations there are none, and in those cases we assume the passenger is rejected and leaves the system.




Fig. 16. Results of the simulations when assigning with the No IGD rule, starting with 10 drivers (left) or 3 drivers (right).

With the users assigned through No IGD, if we consider $t_p(I) = M/4I$ and have a total fleet of $F = N + M/4N$ (for some exogenous N), the two solutions for Eq. (7) are $I = N$ or $I = M/4N$. Clearly, the greater of these values represents the good equilibrium, and the lower represents the WGC. We run simulations using $M = 120$ and $N = 10$ (thus $M/4N = 3$), and show the average of 70 repetitions in Fig. 16. The left panel starts with 10 drivers and the right panel with 3. Remarkably, we observe that in both cases the number of idle drivers converges to a bit more than 10, i.e., to the good equilibrium. The conclusion is very clear: in the case of the circle, when the IGD is not present, the WGC is not triggered. It is noteworthy that we repeated the experiment with numerous other values for M and N obtaining equivalent results.




Fig. 17. Results of the simulations when assigning with the greedy rule, starting with 10 drivers (left) or 3 drivers (right).




Fig. 18. Summary of the spatial dynamics in ride-hailing and similar systems.

When we use the greedy rule (i.e., the more realistic one), the solutions for Eq. (7) are unclear, because the relationship between I and t_p depends on the limit distribution of the drivers. Nevertheless, if we replicate the experiments underlying Fig. 16 but now apply the greedy rule, we obtain the results in Fig. 17, where the key observation is that the number of idle drivers converges to an average just slightly greater than zero, which is obviously a WGC state. Crucially, this result does not hold for every N and M : for example, if N is increased to $N = 15$, the system converges to having approximately 4 drivers on average. Although this value seems low, we lack theoretical guarantees that 4 idle drivers would correspond to a WGC state. In short, under the greedy dynamics with IGD, we cannot claim that the system *always* converges to a WGC, but we do show that it *sometimes* does.

Developing a theoretical model for this extended framework, in which the IGD and WGC interact endogenously, is a natural direction for further research. Two aspects seem particularly important.

1. It remains to characterise the conditions under which the WGC emerges. In the example above, when $M = 120$ and $N = 15$, the system converges to approximately four idle drivers on average. Although this is clearly a low-idle-vehicle regime, the present simulations do not allow us to determine whether it should be formally interpreted as a WGC state. A theoretical characterisation would make it possible to identify whether there is a threshold in M , N , or in the initial number of idle drivers below which the system converges to the WGC.
2. The No IGD rule should not be interpreted as a realistic assignment policy. Rather, its purpose is purely diagnostic: by removing the feedback mechanism that gives rise to the IGD, it demonstrates that, in this circular model, the WGC does not emerge in its absence. In practice, an effective assignment policy would ideally match passengers to nearby drivers (though not necessarily the nearest one) while simultaneously preventing the IGD and the resulting WGC. Whether such a policy can be designed in this extended setting, and more generally, what the optimal assignment rule would look like, remains an important open question.

In all, this section reveals that, in the circular model considered here, the WGC would not be triggered at all if the IGD was avoided, but the same is not true when the IGD is present.

7. Summary of the spatial dynamics under online matching

Let us go back to the general spatial model, combining our theoretical findings with previous studies and the insights from the simulations. How do the IGD and the WGC relate and affect the spatial distribution of the servers? How does the distribution of the servers at equilibrium look like?

All of our analysis depict the following dynamics, illustrated in Fig. 18. There are two feedback cycles, the IGD and the WGC. When present, both cycles are first triggered by the IGD. To be more specific, and assuming that the servers begin perfectly distributed:

1. The random arrival of users and servers perturbs the original uniform distribution so that idle servers become unevenly distributed, which directly increases the average distance between users and servers.
2. As explained by the IGD, the large dominance zones are more likely to attract users, so the servers become even less balanced - namely, because every user that appears in an already large dominance zone increases this dominance zone further.
3. The increased distance between users and servers implies that drivers take a longer time to become idle again, which reduces the number of idle servers, as described by Castillo et al. (2024).
4. Having fewer idle servers obviously further increases the distance between a user and their assigned server.

These feedback cycles degrade the system's quality of service until a new probabilistic equilibrium is reached. This equilibrium is not only worse than a servers' perfect distribution, but significantly worse than a fully random distribution.

As discussed in Section 4.1 (where we discuss the suboptimality of the greedy policy), and reinforced in our experiments using batch assignment (Section 6.1), there is no improvement to the matching algorithm that could substantially modify these dynamics. On the other hand, a proactive rebalancing of the idle servers could help maintain a better distribution. This is particularly relevant when drivers' behaviour is accounted for, as they might relocate to undersupplied areas (the gaps). However, two caveats must be noted: first, under high-demand circumstances, the impact of rebalancing might be mild because the percentage of idle servers remain low⁹ (Fielbaum et al. (2022)); second, rebalancing is relatively simple in ride-hailing, but in other instances of spatial matching this is no longer the case, such as the examples given in the introduction, i.e., dockless bike-sharing and on-street parking when delivering.

⁹ Additional results are presented in the Appendix, that extend the Manhattan simulations by including a simple rebalancing step. The trends are not changed at all compared to Fig. 15.

Let us recall that WGC models (e.g. [Castillo et al. \(2024\)](#), [Li and Ge \(2025\)](#)) have typically identified that for a given fleet and demand rate, there are two possible equilibria, the “good one” with low pickup times and the WGC equilibrium with large pickup times. However, these models do not take into account how regions become spatially heterogeneous, i.e., the IGD. What the dynamics summarised in [Fig. 18](#) reveal is that, because of the IGD, pickup times will tend to increase from the ideal situation. This does not imply that the probabilistic limit of the process corresponds to the WGC equilibrium. Nevertheless, whenever it does converge to a WGC, it is induced by the IGD.

8. Conclusions

Numerous emerging transport systems rely on some form of online matching between a user and the server (vehicle) it will utilise to travel. Some traditional ones, including some examples of parking, are not done online but still follow a similar process. This *spatial matching* is subject to particular dynamics that explain the temporal evolution of the spatial distribution of the servers. In this paper, we have identified and described the *Increasing Gap Dynamics* (IGD), namely, that once the servers become unevenly distributed, those covering greater gaps are more likely to be assigned, further increasing those gaps. Through a number of analytical and experimental results, we have shown that the servers’ distribution converges to a probabilistic equilibrium that is worse not only than a perfect distribution, but also than a fully random one.

Most of our theoretical analysis has been done over a circle and assuming a greedy assignment rule. Therein, we formally prove that a random distribution tends to become worse and that this happens exactly because of the IGD. Moreover, we have also proved that the optimal assignment rule will always match a user to one of its neighbouring servers, which suffices for the IGD to appear. We have also provided very evident visualisations of the phenomenon in a unit square. Simulations of a proper ride-hailing system using real-world data have shown not only that the IGD occurs, but also that it may trigger the well-known *Wild Goose Chase* (WGC).

As this is an emerging topic, there are numerous directions for further research. At the theoretical level, a formal bound on the relationship between the optimal and greedy algorithms on the circle remains an open problem. Similarly, extending our theoretical results to obtain further analytical insights about the IGD limit distribution on the circle would be very valuable. A related research direction is a formal analysis of the system where the number of drivers is not constant (as in [Section 6.3](#)), when drivers rebalance or act strategically, or when there is a given correlation between the pickup point and the place where the driver reappears (i.e., the dropoff point). Moreover, studying other policies in more complex environments such as a graph or the unit square (e.g., the hierarchical greedy proposed by [Kanoria \(2025\)](#)) could provide relevant insights into the evolution of the IGD in real-life scenarios. From an application perspective, while our results show that the IGD cannot be fully prevented, a crucial question emerges, namely, what remedial measures can be implemented to improve the efficiency and quality of service of the numerous systems that can be represented by this general spatial model. Finally, the models have been done for *Assigned* modes, i.e., where users and vehicles are matched online ([Fielbaum \(2025\)](#)): would the IGD still appear for *Onsite* modes, such as taxis, where the servers and users move randomly till finding each other?

CRedit authorship contribution statement

Andres Fielbaum: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization; **Roberto Cominetti:** Writing – review & editing, Validation, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization; **Jose Correa:** Writing – review & editing, Validation, Methodology, Investigation, Formal analysis, Conceptualization.

Acknowledgments

Dr Andres Fielbaum has been partially funded by the [Australian Research Council](#) (ARC) Discovery Early Career Researcher Award (DECRA) [DE250100417](#). The research of Roberto Cominetti was partially supported by FONDECYT [1241805](#). The research of Jose Correa was partially supported by FONDECYT [1260036](#) and by ANID grant [FB210005](#). The idea of this paper was originally discussed at the Dagstuhl Seminar 24281 on Dynamic Traffic Models in Transportation Science. The authors thank the comment of the three anonymous reviewers which greatly helped improved the quality of this paper.

Data availability

Data will be made available on request.

Appendix A. Lemma 1 and Theorem 1 in the continuous case

In this section we establish the convergence of the greedy assignment for the stochastic dynamics in the case where drivers and users are distributed over a sufficiently regular continuous space X . Specifically, we consider X to be a compact subset of $\mathbb{R}^d$ with nonempty interior, endowed with its d -dimensional Lebesgue measure denoted by μ (i.e., the usual length if $d = 1$, area if $d = 2$, volume if $d = 3$, etc.). We assume that X is regular in the sense that for all $x \in X$ and $r > 0$ we have $\mu(X \cap B(x, r)) > 0$. Since this measure varies continuously with x over the compact set X , for every fixed $r > 0$ we can find $\alpha(r) > 0$ such that $\mu(X \cap B(x, r)) \geq \alpha(r)$ uniformly for all $x \in X$. We also suppose that the probabilities P and Q governing the appearance of users and drivers over the

ground set X have densities with respect to μ , and that there exists $\rho > 0$ such that for every Borel set $A \in \beta(X)$ we have $P(A) \geq \rho\mu(A)$ and $Q(A) \geq \rho\mu(A)$ (in particular this holds when the densities are continuous functions because X is compact). We denote by $P_N(\cdot)$ and $Q_N(\cdot)$ the corresponding N -fold product probability measures on the space X^N , and by $P(s, C) = \mathbb{P}(S^{t+1} \in C | S^t = s)$ the one-step transition kernel of the chain for each state $s \in X^N$ and every Borel set $C \in \beta(X^N)$. Further, we denote $D = \{s \in X^N : \exists i \neq j, s_i = s_j\}$ the exceptional set of N -tuples where two or more drivers coincide.

In order to extend [Lemma 1](#) and [Theorem 1](#) to this continuous setting, we use Theorem 13.3.3 in [Meyn and Tweedie \(2012\)](#)¹⁰, which provides sufficient conditions ensuring that the chain has a unique invariant probability measure $\pi(\cdot)$ and that for every initial state $s \in X^N$ the t -step distribution $\mathbb{P}(S^t \in C | S^1 = s)$ converges in total variation towards $\pi(C)$. Namely, this is guaranteed as long as the chain is:

- a) *φ -irreducible*: that is, there exists a non-zero measure φ on $\beta(X^N)$ such that for every state $s \in X^N$ and each Borel set $C \in \beta(X^N)$ with $\varphi(C) > 0$ we have $\mathbb{P}(S^t \in C | S^1 = s) > 0$ for some $t \geq 1$ (see Proposition 4.2.1 (ii) in [Meyn and Tweedie \(2012\)](#));
- b) *Harris recurrent*: that is, every Borel set $C \in \beta(X^N)$ with $\psi(C) > 0$ satisfies $P_x(\Gamma_C = \infty) = 1$, where Γ_C is the number of times C is visited and ψ is the maximal measure described in Proposition 4.2.2 in [Meyn and Tweedie \(2012\)](#);
- c) *Aperiodic*: the maximum m for which an m -cycle exists is $m = 1$. An m -cycle is a partition of X^N into m subsets $C_1, \dots, C_m$, such that for all $s \in C_i$ we have $P(s, C_{i+1}) = 1$; and
- d) *Positive*: there exists an invariant probability measure.

In order to show that these conditions hold in our setting, we will exploit the following technical fact.

Lemma 3. *Let $C \in \beta(X^N)$ be any Borel set with positive Lebesgue measure $\mu_N(C) > 0$. Then, there exists $\varepsilon > 0$, which depends only on C , such that $\mathbb{P}(S^{2N+1} \in C | S^1 = s) \geq \varepsilon$ uniformly for $s \in X^N$.*

Proof. We first argue that it suffices to consider the case where C is a product $A_1 \times \dots \times A_N$ of disjoint closed subsets $A_1, \dots, A_N$ of X . Indeed, by Lebesgue's density theorem almost every point in $C \setminus D$ has density 1 so that taking such a point $x \in C \setminus D$ and any $\kappa \in (0, 1)$ we can find a small radius $r > 0$ such that the $\|\cdot\|_\infty$ -ball $R = \{s \in X^N : \|s - x\|_\infty \leq r\}$ satisfies $\mu_N(C \cap R) \geq \kappa \mu_N(R)$. Observe that R is a product of rectangles $R = R_1 \times \dots \times R_N$ and since $x \notin D$ we can choose r small enough so that the R_k 's are pairwise disjoint. Setting $A_i = R_i \cap X$ it follows that the product set $\tilde{C} \triangleq A_1 \times \dots \times A_N$ is a Borel subset of $X^N \setminus D$, while the regularity of X implies that $\mu_N(\tilde{C}) = \prod_{i=1}^N \mu(A_i) > 0$. Moreover, since P_N and Q_N have a strictly positive density it follows that there exists some $\gamma > 0$ such that $\mathbb{P}(S^{2N+1} \in C | S^1 = s) \geq \gamma \mathbb{P}(S^{2N+1} \in \tilde{C} | S^1 = s)$. Hence, it suffices to find a uniform lower bound for the right hand side, that is, for sets of product form $A_1 \times \dots \times A_N$ for pairwise disjoint A_k 's with $\mu(A_k) > 0$ and therefore $P(A_k) > 0$ as well as $Q(A_k) > 0$.

Now, considering a potentially smaller radius $r > 0$, we can further assume that there is a second family of pairwise disjoint closed subsets $B_1, \dots, B_N$ included in the interior of X , which are disjoint from all the A_k 's and such that $\mu(B_k) > 0$ (so that $P(B_k) > 0$ and $Q(B_k) > 0$). Without loss of generality we may assume that the sets B_k have a small diameter $\delta > 0$ with δ smaller than the distance between any two of these sets B_k and also smaller than the distance between these sets and the boundary of X . With this reduction we can assume from now on that C is of the form $C = A_1 \times \dots \times A_N$ with all these properties. The rest of the proof consists in two steps: first, we show that starting from any initial state $s \in X^N$, after N periods there is a uniformly positive probability that the chain reaches an intermediate state S^{N+1} in which every B_k contains exactly one driver. Then after N additional periods we show that with a uniformly positive probability the chain reaches a state $S^{2N+1} \in A_1 \times \dots \times A_N$.

First step. We claim that there exists $\varepsilon_1 > 0$, independent of the initial state $S^1 = s$, such that with probability at least ε_1 , after N steps the drivers are relocated so that each B_k contains exactly one driver.

Let $s \in X^N$ be an arbitrary initial state and let η_k^1 be the number of drivers in B_k . Denote $\eta_{N+1}^1 = N - \sum_{k=1}^N \eta_k^1$ the drivers located in $X \setminus \cup_{k=1}^N B_k$. The superindex represents the time step $t = 1$, as these values will evolve. We will show that there exists $\varepsilon'_1 = \varepsilon'_1(\eta^1)$ that only depends on η^1 such that with probability at least ε'_1 after N steps every set B_k contains exactly one driver. Since there is a finite number of possible configurations for η^1 , by taking ε_1 as the minimum of these ε'_1 this property will hold uniformly for $s \in X^N$.

In order to find ε'_1 as required, we relocate the drivers one at a time until we have $\eta_k^{N+1} = 1$ for all $k = 1, \dots, N$ and $\eta_{N+1}^{N+1} = 0$. In every iteration t , we have some sets B_k that have exactly one driver. We call these the “achieved” sets with “achieved” drivers. When this happens for every $k = 1, \dots, N$ we have completed our task. If this is not yet the case, we take any j such that $\eta_j^t = 0$. We want to remove a driver that is not achieved, and replace it by a new driver located within B_j . We now prove that this occurs with a probability $\varepsilon''_1 > 0$ that only depends on the family of currently achieved sets B_k . Since there are finitely many combinations for such achieved sets, it then suffices to define ε'_1 as the product of all such ε''_1 .

Let $p = \min_k P(B_k)$ so that the probability that the new driver appears in any B_j is at least p . Notice that this p is strictly positive and does not depend on s nor η^1 . What about the probability of selecting a driver that is not achieved? It suffices to show that there exists $v_1 > 0$ such that the Lebesgue measure of the Voronoi regions covered by non-achieved drivers is larger than v_1 , since then the probability of the user being assigned to a non-achieved driver is at least ρv_1 and we may then take $\varepsilon''_1 = p \rho v_1$. In order to find an appropriate v_1 let us set $r = \delta/4$ and let v be the d -volume of a half ball of radius r . We claim that we can take $v_1 = \min\{\alpha(r), v\}$. In order to prove this, it suffices to consider the case where there is a single non-achieved driver s_o (if there are more, the joint volume of the Voronoi cells of non-achieved drivers would be larger). We distinguish two cases for s_o .

- **Case 1:** If for every achieved B_k we have $\text{dist}(s_o, B_k) \geq \delta/2$, then $B(s_o, r) \cap X$ is included in the Voronoi cell of s_o which then has a Lebesgue measure at least $\alpha(r) \geq v_1$.

¹⁰ The theorem can also be stated as a particular case of [Theorem 1](#), part (ii), in [Borovkov and Odell \(2007\)](#). However, we keep the explicit proof here as it provides some additional intuition of the underlying dynamics.

- **Case 2:** There is some achieved B_k such that $\text{dist}(s_o, B_k) < \delta/2$. Note that the choice of δ implies that this k is unique and also that $B(s_o, r) \subseteq X$. Let s_k be the position of the only driver in B_k . The Voronoi partition that considers just s_o and s_k is given by a hyperplane passing through the mid point of the segment connecting s_k and s_o . Now, at least half of the ball $B(s_o, r)$ is included in the half-space containing s_o , and since all the other achieved sets B_i are at a distance larger than $\delta/2$ from s_o , this half ball is included in the Voronoi cell of s_o which then has a Lebesgue measure of at least $v \geq v_1$.

Both cases combined show that regardless of the position of s_o , the described ε_1'' works. This concludes the first step: there is a lower bound ε_1 , independent of s , such that $P(\eta_1^{N+1} = 1, \dots, \eta_N^{N+1} = 1 | S^1 = s) \geq \varepsilon_1$.

Second step. We now prove that the probability of moving from an arbitrary state s having exactly one driver in each B_k to a state $\tilde{s} \in A_1 \times \dots \times A_N$, also has a uniform lower bound ε_2 independent of s . Let $q > 0$ be such that $Q(B_k) \geq q$ and $P(A_k) \geq q$ for all $k = 1, \dots, N$. Consider the driver at position s_1 and let B_j the set that contains s_1 . Then, with probability $Q(B_j) > q$ an arriving user will appear in this set B_j and will be assigned to the driver s_1 , while with probability $P(A_1) \geq q$ the new driver will appear in a position $\tilde{s}_1 \in A_1$. Hence, with probability q^2 the state s will transition to a new state with the first driver located in A_1 . Proceeding sequentially by replacing the drivers at $s_2, s_3, \dots$ with new drivers at $\tilde{s}_2 \in A_2, \tilde{s}_3 \in A_3, \dots$ it follows that after N periods the state s will transition to a state in $A_1 \times \dots \times A_N$, with probability at least q^{2N} .

Combining the first and second steps it follows that after $2N$ periods the state S^{2N+1} has a probability at least $\varepsilon = \varepsilon_1 \varepsilon_2$ to belong to the target set $C = A_1 \times \dots \times A_N$, uniformly in the initial state $S^1 = s$. $\square$

Corollary 2. Let $C \in \beta(X^N)$ be any Borel set with positive Lebesgue measure $\mu_N(C) > 0$. Then, there exists $\varepsilon > 0$, which depends only on C , such that $\mathbb{P}(S^t \in C) \geq \varepsilon$ for all $t \geq 2N + 1$, regardless of the starting state S^1 .

Proof. By conditioning on the state S^{t-2N} visited at stage $t - 2N$ we have $\mathbb{P}(S^t \in C) = \mathbb{E}[\mathbb{P}(S^t \in C | S^{t-2N})]$. Since the chain is homogeneous, this is equal to $\mathbb{E}[\mathbb{P}(S^{2N+1} \in C | S^1)]$ and we conclude by Lemma 3. $\square$

Using these results we may now proceed to establish the conditions that ensure the existence of a unique invariant measure and the convergence in total variation of the t -step distribution of the chain.

Condition a). In order to check φ -irreducibility we may just take $\varphi = \mu_N$ since Lemma 3 ensures that for every initial point $s \in X^N$ and any Borel set C with strictly positive Lebesgue measure $\mu_N(C) > 0$ the chain has a strictly positive probability of reaching C at stage $t = 2N + 1$.

Condition b). Lemma 3 implies that $\mathbb{P}(\Gamma_C = \infty) = 1$ for all C with positive Lebesgue measure $\mu_N(C) > 0$. However, Harris recurrence needs this to be true for any set with $\psi(C) > 0$ where ψ is a maximal measure as defined in Proposition 4.2.2 in Meyn and Tweedie (2012). Now, since the chain is also irreducible with respect to ψ (as stated in the Proposition 4.2.2 just cited), every set C with $\psi(C) > 0$ satisfies $\mu_N(C) > 0$. Indeed, if $\mu_N(C) = 0$, and since the probabilities P_N and Q_N have densities with respect to μ_N , the probability for the drivers to reach C would be zero for every t , implying that $\psi(C) = 0$. Hence, Lemma 3 is indeed sufficient to ensure that the chain is Harris recurrent.

Condition c). In any m -cycle $C_1, \dots, C_m$ the sets C_i are only visited at regular intervals of length m . Now, since P_N and Q_N are absolutely continuous with respect to the Lebesgue measure we have $\mu_N(C_i) > 0$ and then Corollary 2 shows that C_i can be reached with positive probability at every stage $t \geq 2N + 1$. This necessarily implies that $m = 1$.

Condition d). To prove the existence of an invariant probability we use Theorem 3.1 in Lasserre (2000). Since the set $D = \{s \in X^N : \exists i \neq j, s_i = s_j\}$ is such that $\mu_N(D) = 0$, and P as well as Q have a density with respect to the Lebesgue measure, the probability that the chain jumps in one step from a state $s \in X^N \setminus D$ towards D is zero, while Lemma 3 gives $\mathbb{P}(S^{2N+1} \notin D | S^1 = s) > 0$ for all $s \in D$. On the other hand, we claim that the weak Feller property holds at every point $s \notin D$, namely, for every continuous function $f \in C(X^N)$ and each $s \in X^N$, by considering the Voronoi cells $V^i(s) = \{x \in X : \|x - s_i\|_2 \leq \|x - s_j\|_2 \text{ for all } j \neq i\}$ for the Euclidean norm $\|\cdot\|_2$, we have

$$\mathbb{P}f(s) \triangleq \int_{X^N} f(x) P(s, dx) = \sum_{i=1}^N Q(V^i(s)) \int_X f(s_1, \dots, s_{i-1}, x_i, s_{i+1}, \dots, s_N) dP(x_i).$$

Below we prove that for each $i = 1, \dots, N$ the function $s \mapsto Q(V^i(s))$ is continuous at every $s \in X^N \setminus D$, so that the same holds for the map $s \mapsto \mathbb{P}f(s)$. Finally, take $f_0 : X^N \rightarrow [0, 1]$ a continuous function that vanishes exactly on D , that is $\{s \in X^N : f_0(s) = 0\} = D$. Take C a compact subset of $X^N \setminus D$ with $\mu_N(C) > 0$, and let $m = \min_{s \in C} f_0(s)$ so that $m > 0$. For an arbitrary initial state $s_0 \in X^N$ we have $\mathbb{E}_{s_0}[f_0(S^t)] \geq m \mathbb{P}(S^t \in C)$ and then Corollary 2 implies

$$\limsup_{n \rightarrow \infty} \mathbb{E}_{s_0} \left[\frac{1}{n} \sum_{t=0}^{n-1} f_0(S^t) \right] \geq m \varepsilon > 0.$$

These conditions allow us to invoke Theorem 3.1 in Lasserre (2000) which ensures that the chain has an invariant probability measure π , establishing condition d).

Continuity of $s \mapsto Q(V_i(s))$. Consider two N -tuples of drivers $s, r \in X^N \setminus D$ and let $\delta = \|s - r\|$. Consider a point $x \in X$ such that $x \in V_i(s) \setminus V_i(r)$, that is, a user at x would be assigned to the i -th driver when the drivers are located according to s but to a different driver $j \neq i$ under r , so that $\|x - r_j\| \leq \|x - r_i\|$ while

$$\|x - r_i\| \leq \|x - s_i\| + \|s_i - r_i\| \leq \|x - s_j\| + \delta \leq \|x - r_j\| + \|r_j - s_j\| + \delta \leq \|x - r_j\| + 2\delta.$$

It follows that $V_i(s) \setminus V_i(r) \subseteq \cup_{j \neq i} \Delta_{ij}^\delta(r)$ where $\Delta_{ij}^\delta(r) \triangleq \{x \in X : 0 \leq \|x - r_i\| - \|x - r_j\| \leq 2\delta\}$. These regions decrease with δ and, because $r_j \neq r_i$, for $\delta = 0$ they become the intersection of X with a hyperplane through the midpoint between r_i and r_j . Since X is bounded, it follows that $\mu(\Delta_{ij}^\delta(r)) \rightarrow 0$ as $\delta \rightarrow 0$. Then, the Lebesgue measure of the symmetric difference $V_i(s) \Delta V_i(r)$ of the Voronoi cells, and *a fortiori* the probabilities $Q(V_i(s) \Delta V_i(r))$, tend to 0 when δ tends to 0. Since $|Q(V_i(r)) - Q(V_i(s))| \leq Q(V_i(s) \Delta V_i(r))$, we conclude that for any fixed tuple $s \in X^N \setminus D$ and an arbitrary sequence $r_n \rightarrow s$, we have $Q(V_i(r_n)) \rightarrow Q(V_i(s))$.

Conclusion. Having proved that the conditions a)-d) hold for any continuous ground set X and probabilities P and Q that satisfy the stated regularity assumptions, Theorem 13.3.3 in [Meyn and Tweedie \(2012\)](#) implies the convergence of the chain in total variation towards the unique invariant measure π . It follows that [Lemma 1](#) and [Theorem 1](#) in [Section 3](#) remain valid under these conditions. $\square$

Appendix B. Proof of [Lemma 2](#)

Proof. We compute the expected value $\Delta V(\ell)$ by conditioning on which intervals are merged and where the new driver falls. Namely, denote M_k the event in which the new user appears in a position that produces the merging of I_k and I_{k+1} , which has probability $P(M_k) = (\ell_k + \ell_{k+1})/2$, and let D_j the event where the new driver appears in the interval I_j so that $P(D_j) = \ell_j$. We distinguish the cases $j \in \{k, k+1\}$ and $j \notin \{k, k+1\}$, by considering separately the events $\Omega_k = M_k \cap (D_k \cup D_{k+1})$ and $\Omega_{kj} = M_k \cap D_j$ for $j \notin \{k, k+1\}$. Then, using the law of total expectation we have

$$\Delta V(\ell) = \sum_{k=1}^N \mathbb{E}(\Delta V(\ell) | \Omega_k) \cdot P(\Omega_k) + \sum_{k=1}^N \sum_{j \notin \{k, k+1\}} \mathbb{E}(\Delta V(\ell) | \Omega_{kj}) \cdot P(\Omega_{kj}). \quad (\text{B.1})$$

Since the arrivals of users and drivers are independent, we have $P(\Omega_k) = \frac{(\ell_k + \ell_{k+1})}{2}(\ell_k + \ell_{k+1})$ and $P(\Omega_{kj}) = \frac{(\ell_k + \ell_{k+1})}{2}\ell_j$. Hence, it remains to compute the conditional expectations appearing in this formula. To this end, let us suppose that we are in event M_k with I_k, I_{k+1} the two intervals that were merged, and let $I_M = I_k \cup I_{k+1}$ the merged interval, with respective lengths $\ell_k, \ell_{k+1}, \ell_k + \ell_{k+1}$. Consider the two cases.

Case 1: the new driver appears within the merged interval (event Ω_k). Before we had two intervals of length ℓ_k, ℓ_{k+1} , and now they changed to $z, \ell_k + \ell_{k+1} - z$ for some $z \in [0, \ell_k + \ell_{k+1}]$. Since the remaining ℓ_j 's are unchanged, the change in V only depends on the difference these two terms, which is

$$\Delta V_1(\ell_k, \ell_{k+1}, z) = z^2 + (\ell_k + \ell_{k+1} - z)^2 - [\ell_k^2 + \ell_{k+1}^2] = 2[z^2 + \ell_k \ell_{k+1} - z(\ell_k + \ell_{k+1})]. \quad (\text{B.2})$$

From this equation we can see that:

- If $\ell_k = 0$ or $\ell_{k+1} = 0$, the change is zero or negative. This is sensible, as $\ell_k = 0$ means that two drivers were in the same place, and now they have been separated (or there was no change).
- If $z = \frac{\ell_k + \ell_{k+1}}{2}$, the change is zero or negative. This is sensible because the merged interval is now perfectly divided.
- If $\ell_k = \ell_{k+1}$, the change is zero or positive. The reason is the same as in the previous bullet point: the interval used to be perfectly divided.

Given that users arrivals are uniform on $[0, 1]$, we have that z is also uniform in $[0, \ell_k + \ell_{k+1}]$. Hence, the expected value of the change conditional on M_k and the new driver appearing inside $I_k \cup I_{k+1}$ is

$$\mathbb{E}(\Delta V(\ell) | \Omega_k) = \frac{1}{\ell_k + \ell_{k+1}} \int_0^{\ell_k + \ell_{k+1}} 2[z^2 + \ell_k \ell_{k+1} - z(\ell_k + \ell_{k+1})] dz = \frac{1}{3} [4\ell_k \ell_{k+1} - \ell_k^2 - \ell_{k+1}^2]. \quad (\text{B.3})$$

Case 2: the new driver appears outside the merged interval (events Ω_{kj}). Let us denote by ℓ_j the length of the interval divided by the new driver into two shorter intervals of length $z, \ell_j - z$ respectively. The change in V is now

$$\Delta V_2(\ell_k, \ell_{k+1}, \ell_j, z) = (\ell_k + \ell_{k+1})^2 + (l_j - z)^2 + z^2 - [\ell_k^2 + \ell_{k+1}^2 + \ell_j^2] = 2[z^2 - z\ell_j + \ell_k \ell_{k+1}]. \quad (\text{B.4})$$

The analysis of [Eq. \(B.4\)](#) reveals that:

- If $\ell_j \approx 0$ (implying that $z \approx 0$ as well), the change is positive, which is sensible because we are dividing an interval that was already small.
- If $z \approx 0$, the change is positive, which happens because the new driver appeared very close to a previous one, contribution almost nothing to the coverage of the segment.
- If $\ell_k \approx 0$ or $\ell_{k+1} \approx 0$, the change is negative. This occurs because there were two drivers very close to each other and one of them was replaced by someone in a different position, improving the overall situation.

The conditional expected change in this case is given by

$$\mathbb{E}(\Delta V(\ell) | \Omega_{kj}) = \frac{1}{\ell_j} \int_0^{\ell_j} 2[z^2 - z\ell_j + \ell_k \ell_{k+1}] dz = 2 \left[\ell_k \ell_{k+1} - \frac{\ell_j^2}{6} \right] \quad (\text{B.5})$$

The result now follows directly by substituting [\(B.3\)](#) and [\(B.5\)](#) and the corresponding probabilities in [\(B.1\)](#). $\square$

Appendix C. Proof of Theorem 2

Let us start analysing (*1) from Eq. (2). To do this, we will show that actually the expected value of Eq. (B.3) is equal to zero, which suffices because then the first term in Eq. (B.1) nullifies. We omit the multiplying $1/3$ as it does not affect the zero result.

$$E_{\ell}([4\ell_k\ell_{k+1} - \ell_k^2 - \ell_{k+1}^2]) = \int_0^1 E([4\ell_k\ell_{k+1} - \ell_k^2 - \ell_{k+1}^2]|\ell_k + \ell_{k+1} = L) \cdot f_{\ell_k + \ell_{k+1}}(L) dL$$

where f_X stands for the pdf of X . Note that this proof assumes the continuous case, while the discrete case is analogous replacing the integral by a sum. We now show that every term in this integral is zero. In fact,

$$\begin{aligned} E([4\ell_k\ell_{k+1} - \ell_k^2 - \ell_{k+1}^2]|\ell_k + \ell_{k+1} = L) &= E([4\ell_k(L - \ell_k) - \ell_k^2 - (L - \ell_k)^2]|\ell_k + \ell_{k+1} = L) \\ &= 6LE(\ell_k|\ell_k + \ell_{k+1} = L) - L^2 - 6E(\ell_k^2|\ell_k + \ell_{k+1} = L). \end{aligned}$$

Given $\ell_k + \ell_{k+1} = L$, we have that ℓ_k is uniformly distributed in $[0, L]$. Therefore, $E(\ell_k|\ell_k + \ell_{k+1} = L) = L/2$, and $E(\ell_k^2|\ell_k + \ell_{k+1} = L) = L^2/3$, so putting everything together, the term inside the integral becomes

$$6 \frac{L^2}{2} - L^2 - 6 \frac{L^2}{3} = 0.$$

We now show (*2) is strictly positive. Note that we can shift the vector without affecting the result, so we assume that the first driver is located at zero. We then have $n - 1$ uniform variables in $[0, 1]$, so the interval vector $(\ell_1, \dots, \ell_N)$ follows a Dirichlet distribution¹¹ with parameter $(1, \dots, 1)$, where the vector has length N . In that case, it is also an established fact that

$$E(\prod_{i=1}^N \ell_i^{\beta_i}) = \frac{\Gamma(N)}{\Gamma(\sum_i (1 + \beta_i))} \cdot \prod_{i=1}^N \Gamma(1 + \beta_i).$$

Note that in the expression above β is general, but for integers, we have that $\Gamma(k) = (k - 1)!$. So in our case we have to compute $E(\ell_k^2\ell_{k+1}\ell_j - \ell_k\ell_j^3/6)$. Using the general formula above, we get

$$\begin{aligned} E(\ell_k^2\ell_{k+1}\ell_j) &= \frac{\Gamma(N)}{(N + 3)!} \cdot 2! \\ E(\ell_k\ell_j^3/6) &= \frac{\Gamma(N)}{(N + 3)!} \cdot 3! \cdot \frac{1}{6}. \end{aligned}$$

As $\frac{3!}{6} = 1 < 2!$, we conclude that $E(\ell_k^2\ell_{k+1}\ell_j - \ell_k\ell_j^3/6) > 0$, i.e., (*2) is positive completing the proof. $\square$

Appendix D. Results in Manhattan with rebalancing

As discussed in Section 7, here we show and discuss the results in Manhattan when some simple rebalancing step is included. To be specific, after each assignment, we instruct up to 50 idle vehicles to relocate towards the origins of the most recent 50 requests. The vehicles are selected by minimising the total relocating distance. This approach is similar to the one by [Alonso-Mora et al. \(2017\)](#), [Fielbaum et al. \(2022\)](#). It is noteworthy that these vehicles will start moving as instructed, but will be considered available in the next iteration, so they might never reach the relocation objective. The idea is that these partial movements will help have more vehicles in the areas where a vehicle has just been removed. The crucial observation is that the overall trends in [Fig. D.19](#) are the same as in [Fig. 15](#), with changes that only appear as random. This is a preliminary comparison, and a proper model that includes rebalancing is regarded as a relevant direction for further research (as discussed in [Section 8](#)).

¹¹ We use here well-known facts about the distribution of the intervals that follows a random partition of the unit interval. The original source is [Rényi \(1953\)](#), but there are numerous lecture notes and links that describe them as well.




Fig. D.19. Results of the simulations on Manhattan with rebalancing.

Appendix E. Results in the extended one-dimensional model with dropoff time

The results in Section 6.3 were obtained under the simplifying assumption that the dropoff time is negligible, i.e., $t_i = 0$. We now consider a more realistic setting in which vehicles remain unavailable not only while travelling to pick up a passenger, but also while transporting that passenger to their destination. To account for vehicles already engaged in trips when the simulation starts, the warmup phase is extended from $M/4I$ iterations to $M/4 + M/4I$ iterations.

Fig. E.20 shows the evolution of the average number of available drivers for $M = 120$, considering $N = 10$ (top) and $N = 15$ (bottom). The results are qualitatively very similar to those reported in Section 6.3. In the absence of IGD, the system stabilises around the initial number of available drivers. By contrast, when IGD is present, the system converges to a much lower level of availability: for $N = 10$, the average number of available drivers remains between one and two, while for $N = 15$ it stabilises at approximately eight drivers. Therefore, incorporating non-negligible dropoff times does not alter the main conclusions of the paper, and the impact of IGD on fleet availability remains substantial.




Fig. E.20. Evolution of the average number of available drivers under the IGD and no-IGD dynamics for different initial numbers of drivers. Top: $N = 10$. Bottom: $N = 15$.

References

- Alonso-González, M.J., van Oort, N., Cats, O., Hoogendoorn-Lanser, S., Hoogendoorn, S., 2020. Value of time and reliability for urban pooled on-demand services. *Transp. Res. C: Emerg. Technol.* 115, 102621.
- Alonso-Mora, J., Samaranayake, S., Wallar, A., Frazzoli, E., Rus, D., et al., 2017. On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment. *Proc. Natl. Acad. Sci.* 114 (3), 462–467. <https://doi.org/10.1073/pnas.1611675114>
- Alonso-Mora, J., Wallar, A., Rus, D., 2017. Predictive routing for autonomous mobility-on-demand systems with ride-sharing. In: 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, pp. 3583–3590.
- Andrienko, G., Andrienko, N., Chen, W., Maciejewski, R., Zhao, Y., 2017. Visual analytics of mobility and transportation: state of the art and further research directions. *IEEE Trans. Intell. Transp. Syst.* 18 (8), 2232–2249.
- Ashkrof, P., Homem de Almeida Correia, G., Cats, O., Van Arem, B., 2024. On the relocation behavior of ride-sourcing drivers. *Transp. Lett.* 16 (4), 330–337.

- Ashkrof, P., de Almeida Correia, G.H., Cats, O., Van Arem, B., 2022. Ride acceptance behaviour of ride-sourcing drivers. *Transp. Res. C: Emerg. Technol.* 142, 103783.
- Balkanski, E., Faenza, Y., Périvier, N., 2023. The power of greedy for online minimum cost matching on the line. In: *Proceedings of the 24th ACM Conference on Economics and Computation*, pp. 185–205.
- Bansal, P., Liu, Y., Daziano, R., Samaranayake, S., 2020. Impact of discerning reliability preferences of riders on the demand for mobility-on-demand services. *Transp. Lett.* 12 (10), 677–681.
- Beauvoir, V., Moylan, E., 2020. Unreliability of delay caused by bike unavailability in bike share systems. *Transp. Res. Rec.* 2674 (5), 444–451.
- Bejone, C.V., Zhai, Y., Ouyang, Y., 2025. Analysis of ride-hailing systems' stability and dynamics around equilibrium points. Available SSRN 5801894.
- Berbeglia, G., Cordeau, J.-F., Laporte, G., et al., 2010. Dynamic pickup and delivery problems. *Eur. J. Oper. Res.* 202 (1), 8–15. <https://doi.org/10.1016/j.ejor.2009.04.024>
- Besbes, O., Castro, F., Lobel, I., 2021. Surge pricing and its spatial supply response. *Manag. Sci.* 67 (3), 1350–1367.
- Borovkov, K.A., Odell, D.A., 2007. On spatial thinning-replacement processes based on voronoi cells. *Adv. Appl. Probab.* 39 (2), 293–306.
- Boyles, S.D., Lownes, N.E., Unnikrishnan, A., 2020. *Transportation network analysis. Vol. I: Static Dyn. Traffic Assign.*
- Castillo, J.C., Knoepfle, D., Weyl, E.G., 2024. Matching and pricing in ride hailing: wild goose chases and how to solve them. *Manag. Sci.* In press. <https://doi.org/10.1287/mnsc.2022.00096>
- De Berg, M., 2000. *Computational geometry: algorithms and applications*. Springer Science & Business Media.
- Emami, E., Ramezani, M., 2024. Integrated operator and user-based rebalancing and recharging in dockless shared e-micromobility systems. *Commun. Transp. Res.* 4, 100155.
- Fielbaum, A., 2025. Coordination costs in spatial matching: assigned versus onsite transport modes. *Transp. Res. a: Policy Pract.* 199, 104556.
- Fielbaum, A., Alonso-Mora, J., 2020. Unreliability in ridesharing systems: measuring changes in users' times due to new requests. *Transp. Res. C: Emerg. Technol.* 121, 102831.
- Fielbaum, A., Kronmueller, M., Alonso-Mora, J., 2022. Anticipatory routing methods for an on-demand ridepooling mobility system. *Transportation* 49 (6), 1921–1962.
- Fielbaum, A., Salas, D., Zhang, R., Castro, F., 2025. Idle wage as a tool to regulate the relationship between ride-hailing platforms and drivers. *Transp. Res. C: Emerg. Technol.* 174, 105113.
- Fielbaum, A., Tirachini, A., Alonso-Mora, J., 2023. Economies and diseconomies of scale in on-demand ridepooling systems. *Econ. Transp.* 34, 100313.
- Guo, X., Caros, N.S., Zhao, J., 2021. Robust matching-integrated vehicle rebalancing in ride-hailing system with uncertain demand. *Transp. Res. B: Methodol.* 150, 161–189.
- Hu, B., Hu, M., Zhu, H., 2022. Surge pricing and two-sided temporal responses in ride hailing. *Manuf. Serv. Oper. Manag.* 24 (1), 91–109.
- Hua, M., Yu, X., Chen, X., Chen, J., Cheng, L., 2025. Can bike sharing achieve self-balancing distribution? evidence from dockless and station-based cases. *Travel Behav. Soc.* 38, 100879.
- Hyland, M., Mahmassani, H.S., 2018. Dynamic autonomous vehicle fleet operations: optimization-based strategies to assign AVs to immediate traveler demand requests. *Transp. Res. C: Emerg. Technol.* 92, 278–297.
- Kanoria, Y., 2025. Dynamic spatial matching. *Ann. Appl. Probab.* 35 (5), 3086–3118.
- Kaufman, L., Rousseeuw, P.J., 2009. *Finding groups in data: an introduction to cluster analysis*. John Wiley & Sons.
- Knoop, V., Hoogendoorn, S.P., van Zuylen, H., 2009. Empirical differences between time mean speed and space mean speed. In: *Traffic and Granular Flow'07*. Springer, pp. 351–356.
- Lasserre, J.B., 2000. Quasi-Feller Markov chains. *J. Appl. Math. Stoch. Anal.* 13 (1), 15–24.
- Levin, D.A., Peres, Y., 2017. *Markov chains and mixing times. Vol. 107*. American Mathematical Soc.
- Li, X., Ge, J., 2025. Wild goose chase or not? equilibrium in a hybrid ride-hailing market. *Transp. Policy* 160, 73–88.
- Li, X., Ke, J., Yang, H., Wang, H., Zhou, Y., 2024. An aggregate matching and pick-up model for mobility-on-demand services. *Transp. Res. B: Methodol.* 190, 103070.
- Meyn, S.P., Tweedie, R.L., 2012. *Markov chains and stochastic stability*. Springer Science & Business Media.
- Mo, D., Chen, X.M., Zhang, J., 2022. Modeling and managing mixed on-demand ride services of human-driven vehicles and autonomous vehicles. *Transp. Res. B: Methodol.* 157, 80–119.
- Ouyang, Y., Yang, H., 2023. Measurement and mitigation of the “wild goose chase” phenomenon in taxi services. *Transp. Res. B: Methodol.* 167, 217–234.
- Pan, L., Cai, Q., Fang, Z., Tang, P., Huang, L., 2019. A deep reinforcement learning framework for rebalancing dockless bike sharing systems. In: *Proceedings of the AAAI Conference on Artificial Intelligence. Vol. 33*, pp. 1393–1400.
- Puterman, M.L., 2014. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.
- Ramezani, M., Valadkhani, A.H., 2023. Dynamic ride-sourcing systems for city-scale networks-part I: matching design and model formulation and validation. *Transp. Res. C: Emerg. Technol.* 152, 104158.
- Rényi, A., 1953. On the theory of order statistics. *Acta Math. Acad. Sci. Hung* 4 (2), 48–89.
- Stein, W.E., Dattero, R., 1985. Sampling bias and the inspection paradox. *Math. Mag.* 58 (2), 96–99.
- Xu, Z., Sun, X., 2024. Economic analysis of on-street parking with urban delivery. *Transp. Sci.* 58 (6), 1300–1318.
- Xu, Z., Yin, Y., Ye, J., 2020. On the supply curve of ride-hailing systems. *Transp. Res. B: Methodol.* 132, 29–43.
- Yan, C., Zhu, H., Korolko, N., Woodard, D., 2020. Dynamic pricing and matching in ride-hailing platforms. *Nav. Res. Logist.* 67 (8), 705–724.
- Yang, H., Qin, X., Ke, J., Ye, J., 2020. Optimizing matching time interval and matching radius in on-demand ride-sourcing markets. *Transp. Res. B: Methodol.* 131, 84–105.
- Yang, Y., Ramezani, M., 2022. A learning method for real-time repositioning in e-hailing services. *IEEE Trans. Intell. Transp. Syst.* 24 (2), 1644–1654.
- Zakharenko, R., 2023. Pricing shared vehicles. *Econ. Transp.* 33, 100296.
- Zha, L., Yin, Y., Xu, Z., 2018. Geometric matching and spatial pricing in ride-sourcing markets. *Transp. Res. C: Emerg. Technol.* 92, 58–75.
- Zha, L., Yin, Y., Yang, H., 2016. Economic analysis of ride-sourcing markets. *Transp. Res. C: Emerg. Technol.* 71, 249–266.
- Zhang, K., Alonso-Mora, J., Fielbaum, A., 2025. What do walking and e-hailing bring to scale economies in on-demand mobility? *Transp. Res. B: Methodol.* 192, 103156.
- Zhang, K., Chen, H., Yao, S., Xu, L., Ge, J., Liu, X., Nie, M., 2019. An efficiency paradox of uberization. Available at SSRN 3462912.
- Zhang, X., Zhong, S., Jia, N., Ling, S., Yao, W., Ma, S., 2024. A barrier to the promotion of app-based ridesplitting: travelers' ambiguity aversion in mode choice. *Transp. Res. a: Policy Pract.* 181, 103971.