



Sensitivity Analysis of the Spatial Parameters in Modelling the Evolutionary Interaction Between Autonomous Vehicles and Other Road Users

Isam Bitar¹ · David Watling¹ · Richard Romano¹

Received: 29 August 2022 / Accepted: 28 January 2023
© The Author(s) 2023

Abstract

The road user network is a dynamic, ever-evolving population in which road users interact to share and compete for road space. The advent of autonomous road vehicles (ARVs) will usher in numerous opportunities and challenges in road user dynamics. One of the challenges is whether an ARV population would be able to successfully enter the existing road user space. Previous work demonstrates that successful introduction of ARVs into the road network must consider the evolutionary dynamics of the existing population. This study examines the effect of different spatial parameters as starting conditions for the introduction of a small population of ARVs into a resident population of human-driven vehicles (HDV). The model utilises the concept of evolutionary game theory and uses a square lattice grid with a novel agent mobility approach. The results show that ARV success exhibits significant sensitivity to variations in initial cluster size, position, and travel range. ARVs seem to perform best in fewer, larger clusters with a shorter travel range. This suggests that the best form of early ARV introduction may take the shape of centralised, highly co-operative fleets of local passenger or freight transport.

Keywords Autonomous vehicles · Spatial evolutionary game theory · Evolutionarily stable strategies · Road user interaction · Hawk–Dove games

Introduction

Autonomous road vehicles (ARVs) are slowly reaching market maturity and will soon begin entering the road network. The short- and long-term success of ARVs will depend on their ability to interact effectively with human-driven vehicles (HDVs). Many researchers believe that human road users are likely to learn the nuances of ARV behaviour and so take advantage of them to force ARVs to yield at every interaction [1–3]. This would ultimately have a detrimental effect to the ARV population as a whole and may prevent or slow down any real uptake of the new technology. Experiments have shown that humans expect co-operative

behaviour from machines but are not generally willing to reciprocate it [4]. Indeed, there are fundamental differences in the behaviour, communication, and decision-making process between ARVs and HDVs [1, 2, 5–10], that the two groups of road users can be described as two distinct populations competing on a population level.

Competition amongst populations has long been the subject of study by evolutionary biologists. One framework within which population-level competition and co-operation can be studied and modelled is evolutionary game theory [11–14]. Evolutionary game theory revolves around the idea that individuals which inhabit overlapping habitats must often interact to compete for or share resources necessary for survival, such as food, mates, and territory. The outcomes of such interactions influence an individual's *fitness* (ability to survive and reproduce). Thus, strategies which do well against other strategies (and against copies of themselves) would grow in proportion within the population relative to other strategies. A strategy can be an evolutionarily stable strategy (ESS) if it meets one of the following two criteria [12]:

This article is part of the topical collection “Vehicle Technology and Intelligent Transport Systems” guest edited by Oleg Gusikhin and Markus Helfert.

✉ Isam Bitar
ts14isb@leeds.ac.uk

¹ Institute for Transport Studies, University of Leeds, 34-40 University Rd, Leeds LS2 9JT, UK

The subject strategy does better against itself than other strategies do.

If Strategy **S** exists, which does equally well against the subject strategy, the subject strategy does better against **S** than **S** does against itself.

In any population, resident strategies which are not evolutionarily stable are vulnerable to invasion by new strategies which may outperform the resident strategy. Eventually, the resident strategy may become a minority strategy in the population or be driven out of it entirely. On the other hand, invading strategies must themselves be evolutionarily stable if they are to outcompete a resident ESS. These principles would apply to “invading” ARV populations entering a resident—and likely evolutionarily stable—HDV population. Just as with natural populations, individuals (vehicles) on the road network must interact to compete for or share a resource (road space). The outcome of these interactions influences the *fitness* (delay, ride comfort, safety, etc.) of each vehicle. Just as living organisms reproduce, vehicles and driving styles on the road network use memetic “reproduction” through higher sales of “fitter” vehicle models and imitation of more successful driving styles [15].

Previous work has shown that evolutionary stability in a mixed ARV–HDV population can exist on the road network if certain design and introduction considerations are followed [16]. However, these findings have only been applied to theoretical, well-mixed populations. In reality, it is unlikely that the population of ARVs and HDVs will be well mixed, especially during initial introduction. Instead, it is expected that early ARV introduction will take place within geographically bounded, technologically advanced smart cities [17, 18], with few ARVs venturing outside such environments.

The road user network is a spatial one. Vehicles generally travel, park, and operate within proximity of their “home” geography. In fact, the US Federal Highway Administration’s 2020 National Origin–Destination report shows that over 78% of all trips generated in the US were less than ten miles in length [19]. This indicates a strong geographical element in the distribution of vehicles on the network. Spatial conditions can have a significant impact on the shape and evolution of dynamic populations. This is because the geographical element influences an individual’s probability of interaction with members of its own species and members of other species. In a spatially agnostic, well-mixed environment, an individual interacts with other members of the population with the same probability as all other individuals. Conversely, an individual in a spatially non-homogenous environment interacts with geographically nearer individuals with a higher probability than farther ones. For example, [20] found that Hawk populations fair worse in spatial simulations than well-mixed theory, because Doves tend to

cluster in impenetrable masses which ensure any one Dove’s disbenefit from interacting with a Hawk is counteracted by the benefit of interacting with its many Dove neighbours. Conversely, [21] saw that co-operation fared poorly in their snowdrift game simulations. The reason for this was once again a spatial one in that co-operator populations tended to cluster, whereas game rules dictated that an individual was better off being surrounded by members of the opposite type. These findings suggest a significant potential impact of the spatial elements of the road user network on the future *success* of ARVs and the shape of the evolving road user network. *Success* in this context refers to ARVs’ ability to grow and maintain a viable population within the road user network and the final, stable size of this population. Therefore, a good understanding of the spatial dynamics of the population-level interaction between ARVs and HDVs would serve to ensure proper and sustainable introduction of ARVs. To this end, this paper sets out to investigate the possible impact of a selection of spatial factors on the success of ARVs in a simulation setting and whether special attention ought to be paid to these spatial conditions.

Literature Review

The transport environment is a highly interactive one. As such, ARVs that enter the road user population are expected to be able to interact with other road users safely, efficiently and successfully. Road user interaction models abound in the literature and vary in scope, complexity and implementation. Notable recent works include [22] who outlined a three-tiered hierarchical model of road user interaction originally devised by John Michon [23]. Tiered interaction models such as this and [24] provide computationally efficient means of handling complex interaction algorithms in real time, which is vital for ARVs. Crucially, [24] employs a game-theoretic interaction protocol in its middle interaction layer. Game theory allows for the consideration of other road users as active agents with goals and incentives of their own, modelled as *payoffs*. Payoffs dictate how an agent would act given a set of circumstances and possible actions taken by opponents. The game-theoretic school of autonomous driving has so far largely focused on kinematic factors when formulating payoff functions and subsequent decisions [1, 7, 10, 25–32]. Many add an extra layer of complexity for smarter and/or more realistic decision-making, such as repeated games [10], hierarchical reasoning [33–35], receding horizon control [7, 36], Bayesian probability [30, 37] and proxemics [38, 39]. One common drawback of the mentioned models is that they are validated against virtual opponents playing by the same rules. This ignores the fundamental differences between ARVs and most existing road users with whom an ARV will interact. This in turn would

lead to a problem where ARVs are not equipped to deal with HDVs' ability to react to and exploit ARVs' decision-making. Thus, proper understanding of the evolutionary dynamics of the road user population is necessary to ensure ARV manufacturers and policymakers can introduce ARVs that can keep a meaningful and sustainable presence within the road user population.

Several studies have been conducted in the field of transport which employed evolutionary game theory as the main principle, primarily in the study and modelling of route and mode choice [40–44]. Other applications include predicting and building implementations of government subsidies in transport development that are effective and evolutionarily stable in the long term, such as in new-energy vehicles [45] and in public transport [46]. On the road user interaction level, evolutionary game theory has been used to predict driver attention, simulate driver co-operation and study social dilemmas in driving [47–50]. Previous work by the authors has demonstrated that evolutionary stability can be achieved in a well-mixed population of ARVs and HDVs using a combination of conditional strategies, effective V2V communication and external subsidies [16]. However, to the authors' knowledge, there has been no exploration of the spatial implementation of evolutionary game theory in the context of road user interaction within the road network.

Spatial evolutionary game theory has been the subject of many studies over the years outside the field of transport [20, 21, 51–54]. Studies range considerably in their spatial geometries, fitness parameters and replicator dynamics. However, literature generally agrees on the conclusion that the evolutionarily stable solutions for spatial games differ significantly, if not categorically, from their well-mixed, theoretical counterparts. For example, some simulations show that aggressive behaviour is less successful under spatial constraints than co-operative behaviour [20, 53]. Yet, others arrive at opposite conclusions [21]. In both cases, researchers cite spatial conditions for their findings.

An important aspect of spatial distribution is the way the members of one "species" (homogenous group) in the spatial population are distributed relative to each other. Different starting cluster sizes and numbers can greatly influence the performance of the species in the population as the frequency of each individual's interaction with members of its own and other species changes.

Many studies use grid-based lattice structures to represent the spatial distribution of agents [20, 21, 52, 53, 55, 56]. These structures space out agents into statically positioned coordinates on a grid and allow interaction between neighbours. Most are square lattices but some can be polygonal [21], which increases the number of neighbours each agent in the grid has. Others such as [54] employed an interaction protocol where agents create and terminate interaction links with selected neighbours based on benefit. This contrasts

with the typical, indiscriminate interaction with immediate neighbours employed in most lattice-based simulations and shows stable states closer to the theoretical standard thanks to the ability of agents to interact with other agents beyond the immediate vicinity. Other spatial distribution methods include diffusion dynamics as seen in [51]. These allow for population dispersal and variation in population densities and produce interesting results where free movement is applicable.

Method

Model Structure

The study aims to investigate the effect of several spatial characteristics on the success and viability of a small, introduced ARV population within a large, resident HDV population in a contained road user network. To model this, a conceptual visualisation of the spatial setup is demonstrated in Fig. 1. A square lattice grid is used as the basis for the spatial distribution. Each node in the grid represents a geographical region within a given road network. The node's colour reflects the dominant vehicle type in that region. The spatial grid consists mostly of HDV-dominated nodes, with one or more small clusters of ARV-dominated nodes. These clusters will vary in number, size, and position to reflect different initial conditions. Interaction will follow a novel approach better suited to model the nature of travel on the road network. Conventional interaction protocols where static agents are restricted to interaction with immediate neighbours would serve as poor representations of road

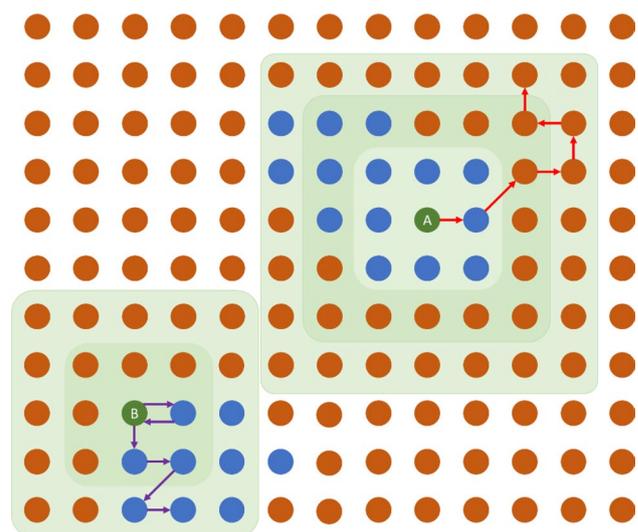


Fig. 1 Demonstrative spatial distribution of ARVs (blue) in a modelled road user network

network dynamics. Conversely, more dynamic approaches such as diffusion allow too many degrees of freedom for each agent compared to what is feasible on the road network. Instead, a novel approach to spatial interaction is adopted. At the start of each generation, a vehicle from each region (node) travels along a randomly generated path through neighbouring nodes. The extent of the travel is determined by the vehicle's *Range* property. Each vehicle then interacts once at each node with the dominant vehicle type of that node. At the end of the generation, vehicles return to their home nodes. Vehicle movement takes place in all directions in two-dimensional space (orthogonally and diagonally). Two-dimensional movement represents the degrees of freedom associated with the movement of ground vehicles; thus, conflicts and interactions are determined on a two-dimensional plane. Each vehicle will move (and interact) a total of six times per generation. This allows each vehicle multiple interactions per turn where the opponent is determined by the trajectory and range of the agent vehicle.

The maximum distance a vehicle travels at the end of each generation is determined by the vehicle's *Range* property. In this paper, *Range* (R) is an attribute that determines the distance required between a vehicle's start and end points at the end of a generation.

In Fig. 1, $R_A = 3$ and $R_B = 2$. A higher R means that the vehicle moves further away from its local cluster, and thus has an increased likelihood of encountering opponents outside of its immediate neighbours. Extended range has the obvious benefit of increased mobility, though can also potentially expose the vehicle to opponents outside of its cluster.

Each interaction represents the meet-up of two vehicles at a junction where priority *must* be negotiated (ignoring each other is not an option at that point). This negotiation is a two-strategy, non-co-operative game in which each vehicle has the option to *escalate* the interaction by accelerating to force priority or *facilitate* the interaction by conceding priority to the other vehicle.

As the road network is rarely a closed space, each edge of the grid is surrounded by a population of "ethereal HDVs". These HDVs do not travel and do not switch sides but will interact with any vehicles that travel out of bounds. These HDVs earn a consistent payoff equal to the standard payoff of HDV–HDV interactions (D) regardless of the type of vehicles travelling out of bounds. This payoff will then factor into the replication stage of each boundary vehicle. This ensures that ARVs do not cheat the game by latching onto a corner early on and eliminating competition from that side.

Each vehicle will accumulate a payoff score as it moves through the grid. At the end of each generation, each node will adopt the colour (vehicle type) of the node with the highest cumulative payoff amongst its immediate nine-node neighbourhood (itself included). This is akin to what's employed in much of the literature, e.g. [20, 21,

54] and is a useful replicator formula that provides a deterministic and relatively computationally efficient approach to a selection process which favours the fittest individuals. Payoffs do not reset between generations. This is to reflect a "lifetime fitness" property where a vehicle's performance in the current generation is added to its historic performance. This prevents an amnesiac form of replicator dynamics where nodes may repeatedly switch sides between generations (called blinkers) or a well-established node of one type is switched because of a single "bad" generation.

Experimental Design

To properly capture the effect of spatial characteristics and distribution on the evolution of ARVs within the model, several simulations are run whilst varying three sets of spatial parameters:

Spatial position of the starting ARV cluster. The three possible starting positions are the *centre*, *edge*, and *corner* of the grid. This tests whether adjacency to the "ethereal HDV" population which does not enter the grid nor change type has any effect on ARV evolution. *Centre* populations have no adjacent ethereal HDVs. *Edge* populations adjoin the ethereal from one side (approx. 25% of the ARV population's perimeter). *Corner* populations adjoin the ethereal from two sides (approx. 50% of the perimeter).

Starting number of ARV clusters, reflecting the level of centralisation in initial ARV introduction—*one*, *three* and *six* individual clusters. Varying the number of clusters whilst maintaining the same introductory ARV population size allows for the measurement of the effect of spatial fragmentation on the overall success of the population. It will also help establish whether a certain *critical mass* exists, below which a cluster cannot prevent being overrun by the surrounding HDVs. The number of ARV clusters bears no impact on the computational complexity of the model. As evolutionary game theory is by design an agent-centric concept, the formation of clusters is an emergent property. Computational load is only a function of the population size in this model.

ARV travel range, which represents different applications of early ARVs—ranging from localised application to longer-range travel. The three choices are *two*, *four* and *six* nodes of travel range per generation. All HDVs have a fixed travel range of four nodes per generation. Increased range means a higher probability to interact with individuals outside of one's nearest neighbours, thus reducing the impact of spatial conditions on the performance of the population. Investigating range will allow for a characterisation of the sensitivity of ARV performance with respect to overall population homogeneity.

The above parameters give a total of nine different starting conditions. Each of these nine is run in six different permutations based on the following:

Average expected payoff. Two scenarios are chosen to reflect two different dynamics which may exist between ARVs and HDVs on the road network. The two scenarios share the assumption that ARVs would face an immediate disadvantage upon introduction. This sets a conservative, “worst-case” representation of the possibility of HDV exploitation of ARVs. It also allows for a categorical evaluation of ARV success where it occurs. The two proposed scenarios reflect two different levels of evolutionary stability for the resident HDV population. They are as follows:

- A. HDVs do well against ARVs, but ARVs do not do as well against each other. This creates a profile where ARVs are thoroughly **dominated by HDVs** at every proportion of the mixed population, thus rendering ARV introduction theoretically impossible (Fig. 2a). HDVs in this scenario are an evolutionarily stable population, immune to ARV invasion
- B. A more efficient ARV–ARV interaction protocol sees ARVs do just as well against each other as HDVs do against ARVs, but still suffer a greater disadvantage against HDVs. This means that ARVs could steadily improve their average expected payoff (AEP) as a function of their proportion until the AEP of both populations is equal at near-100% ARV (Fig. 2b). HDVs here still form an evolutionarily stable population, albeit a weaker one, and so would theoretically still be able to resist invasion.

In both scenarios, the values constituting the average expected payoffs for each of the two vehicle types is determined by two elements. The first element is the payoff

functions which govern each vehicle type’s interaction. The main principles of interaction are designed as follows:

- ARVs interact with other ARVs in an efficient, coordinated manner, thus maximising payoff for both vehicles at a small operating cost (Payoff A in Table 1)
- HDVs which detect ARVs will force priority, relying on ARVs’ propensity to be risk averse and quicker to respond to threats, thus HDVs earn a much higher payoff (C) than ARVs (Payoff B) when the two interact
- HDVs interact with other HDVs in a manner so as, on average, all HDVs will do equally well against each other (Payoff D), though not as well as ARVs do against other ARVs (Payoff A) as HDVs do not engage in cooperative, electronic communication

The second element is the strategy profile for each vehicle type. Strategy profiles can be developed to reflect different driving styles or attitudes, such as increased aggressiveness towards ARVs or a more defensive driving culture. The current study will not investigate the effect of different strategy profiles on ARV success. We therefore use a single strategy profile in which vehicles randomise evenly between *escalation* and *facilitation* when interacting with other vehicles, regardless of the type of agent or opponent vehicle.

All interactions follow a simple payoff formula of *Reward – Cost*. The *Reward* element refers to whether priority is taken, and thus can take the value of 0 or 1 in any single interaction, denoting conceding or taking priority,

Table 1 Tabular representation of the average expected payoffs of interaction between two vehicles

	ARV	HDV
ARV	A	B
HDV	C	D

$$C > D > B, A > B$$

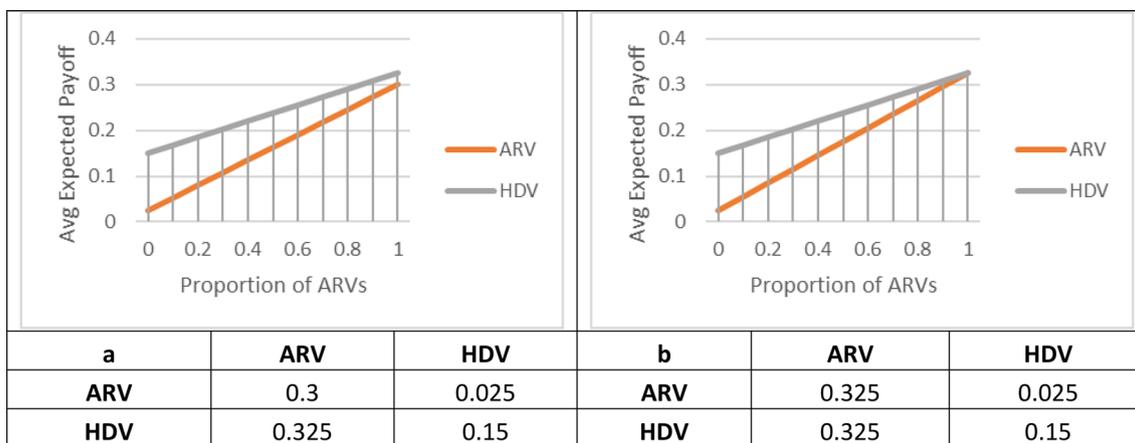


Fig. 2 Theoretical average expected payoffs and payoff profiles for Scenario A (a) and Scenario B (b)

Table 2 Outline of the reward and cost values used in the simulation

Payoff variable	Value	Justification
Reward: mutual <i>facilitation</i>	0.5	Assumes that, on average, each vehicle would receive priority half of the time
Reward: <i>facilitation</i> against <i>escalation</i>	0	The facilitating vehicle will concede priority to the escalating vehicle
Reward: <i>escalation</i> against <i>facilitation</i>	1	The escalating vehicle will force priority
Reward: mutual <i>escalation</i> (excl. HDV against ARV)	0.5	Assumes that, on average, each vehicle would succeed in forcing priority half of the time
Reward: mutual <i>escalation</i> (HDV against ARV)	1	HDVs consistently force priority against the more risk averse ARVs
Cost: <i>facilitation</i> (excl. ARV × ARV)	0.2	A low cost representing the effort taken to communicate and reach an agreement
Cost: <i>escalation</i> (against HDV only)	0.5	Higher cost associated with increased risk and less ideal driving
Cost: <i>escalation</i> (HDV against ARV)	0.4	A slightly lower cost to represent the relative ease in which HDVs force priority against ARVs
Cost: all interaction (ARV × ARV)—Scenario A	0.2	A universal cost representing the use of V2V communication
Cost: all interaction (ARV × ARV)—Scenario B	0.175	A lower cost representing more efficient use of V2V communication

Table 3 Normal-form games for the interaction between the different vehicle types

a ARV × HDV		
ARV \ HDV	Facilitate	Escalate
Facilitate	0.3, 0.3	-0.2, 0.6
Escalate	0.5, -0.2	-0.5, 0.6
b HDV × HDV		
HDV \ HDV	Facilitate	Escalate
Facilitate	0.3, 0.3	-0.2, 0.5
Escalate	0.5, -0.2	0, 0
c ARV × ARV (Scenario A)		
ARV \ ARV	Facilitate	Escalate
Facilitate	0.3, 0.3	-0.2, 0.8
Escalate	0.8, -0.2	0.3, 0.3
d ARV × ARV (Scenario B)		
ARV \ ARV	Facilitate	Escalate
Facilitate	0.325, 0.325	-0.175, 0.825
Escalate	0.825, -0.175	0.325, 0.325

respectively. The *Cost* element is dependent on vehicle type, opponent vehicle type and the action taken. Table 2 outlines the values which make up the costs and rewards for all interaction types.

Given the values in Table 2, Table 3 summarises the normal-form game matrices for each interaction type. Thus, the values shown in Fig. 2 represent the average expected payoffs for each vehicle type given the strategy profile of randomisation and the interaction payoff functions illustrated in Table 3.

The values used in Table 2 and therefore those devised in Table 3 are not validated against real-world data. However, the values themselves are of less importance than the *value difference* between the payoffs of the different vehicle types for the same action pairs. It is the value differences that generate the *selective pressure* which drives the growth or contraction of each population relative to the other. The use of conceptual values such as those devised above means that one cannot take the resulting population dynamics as individually accurate representations. However, one may still observe the differences between the different population dynamics that result from changes to the different spatial parameters. Since the focus of this study is on investigating the sensitivity of ARV evolutionary success to such changes, the use of the non-validated values in Tables 2 and 3 is justified.

Starting ARV population as a percentage of the total: three starting percentages are chosen, which in practice would require different levels of funding, coordination, and legislation. These are *one*, *five* and *ten* per cent. This tests whether there exists a minimum viable introductory population given a set of payoff profiles. It also allows for the study of the sensitivity of ARV success to the initial size of the population. No higher percentages are investigated as it is unlikely to see a real-world introduction of such size.

Finally, each of the resultant 54 simulations are repeated three times using three different seeds (the same three seeds are used for each of the 54 simulations) to normalise the effect of the random selection of vehicle trajectories at the beginning of each generation. All other model elements are deterministic.

Figure 3 below summarises the parameter choices which yield the 162 total simulation runs outlined previously,

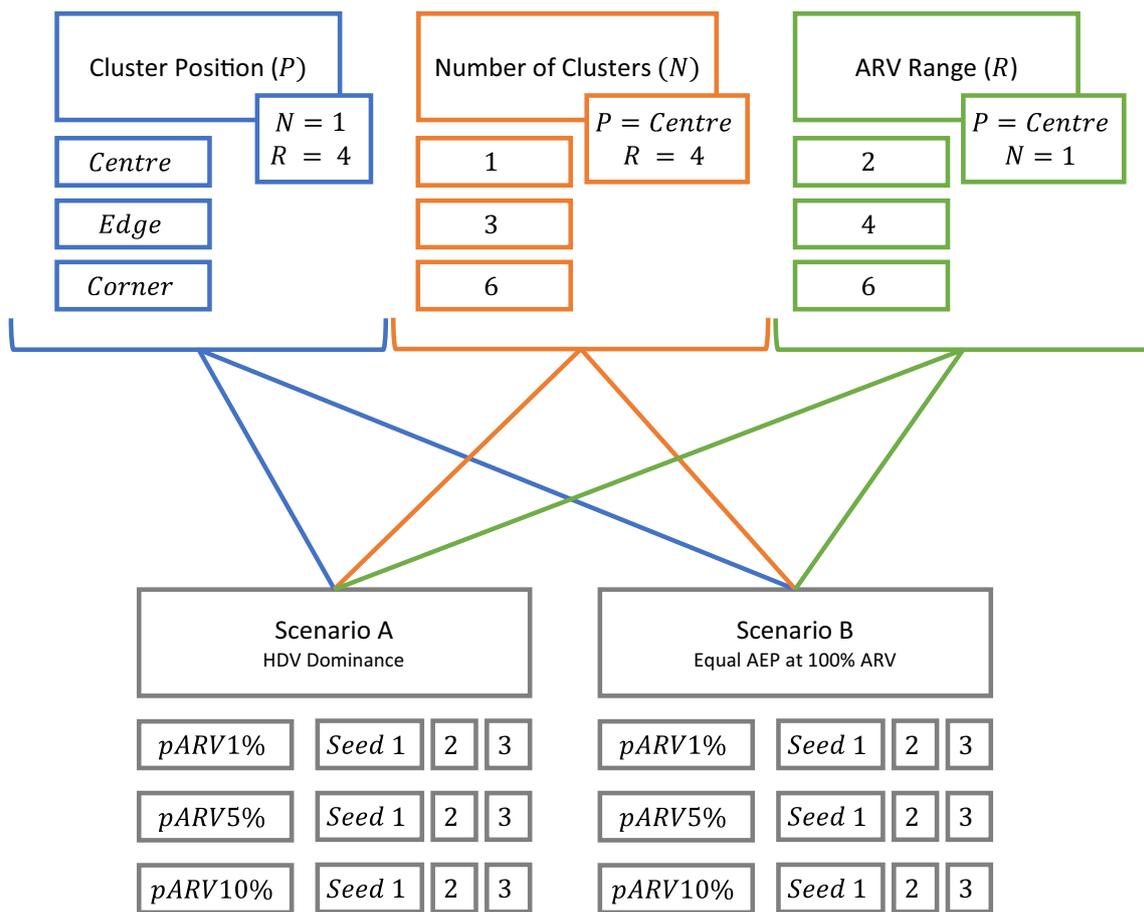


Fig. 3 Diagram outlining the 162 permutations of the chosen starting conditions

covering all permutations of the above parameters. Each simulation run comprises 2500 nodes in a 50-by-50 grid and run for 300 generations.

Figure 4 illustrates the steps of the model in flowchart form.

Hardware/Software Requirement

The simulations are carried out in a purpose-built simulator programmed in Python 3.10.4, in Visual Studio Code, and run on a Windows 10 Desktop PC housing a 2.9-GHz, six-core Intel Core i5-10,400 processor, 16 GB of DDR4 RAM at 3200 MHz data rate, and an Nvidia GeForce RTX 3060 Ti graphics processor.

Results

All 162 simulation runs were concluded successfully and the proportion of ARVs out of the total population in every generation was recorded and plotted. Figures 5, 6 and 7 outline the results of the runs.

All nine *pARV10% position* runs under *Scenario A* pay-offs successfully completed the entire 300-generation run with a stable ARV population. The *centre* starting position produced no other stable ARV populations. In contrast, the *edge* starting position produced one stable ARV population at *pARV5%* and the *corner* starting position produced three. None produced any viable ARV populations at *pARV1%* as all such runs concluded with ARV extinction within the first 15 generations.

ARVs fared better in the *Scenario B position* runs, with all *pARV5%* and *pARV10%* runs concluding with stable ARV populations. As with *Scenario A* runs, however, none of the *pARV1%* runs produced a viable population. ARVs seemed to perform marginally better in the *centre*, which

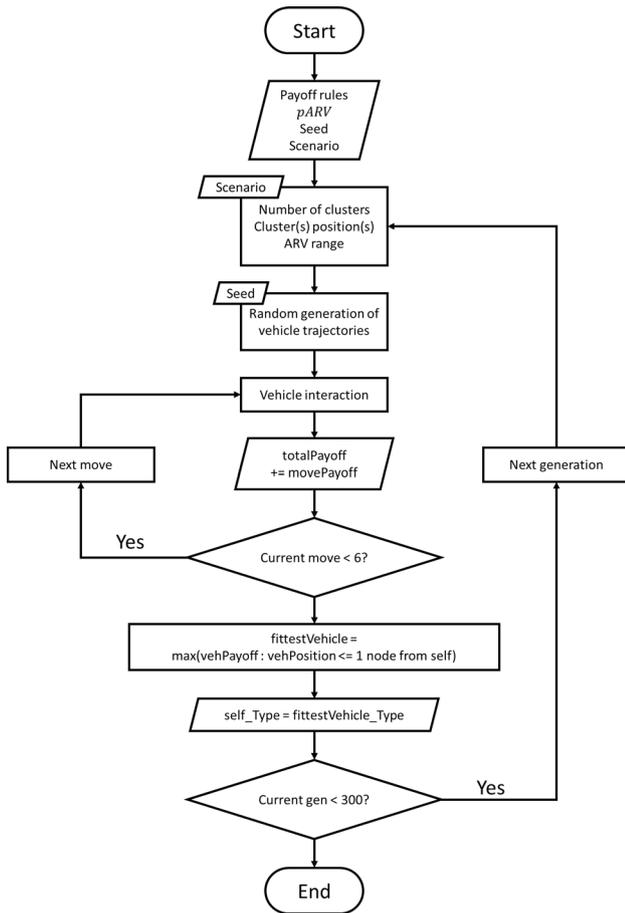


Fig. 4 Flowchart diagram outlining the simulation steps of each of the 162 runs

is in contrast with the observations made under *Scenario A* runs.

The *clustering* runs showed more varied results under *Scenario A* payoffs. Here, only three out of 27 runs produced stable ARV populations. These were the three *pARV10%* runs with the *one-cluster* start. The remainder all saw ARVs die out relatively quickly, though the failed *one-cluster* runs fared significantly better than both the failed *three-cluster* and *six-cluster* runs. On average, *one-cluster* runs had ARVs die out in 57 generations (max. 170), *three-cluster* runs in 18 generations (max. 57) and *six-cluster* runs in six generations (max. 11).

A relatively better picture once again materialises in the *Scenario B clustering* runs. Ten out of 27 runs produced stable ARV populations. Six of these were *one-cluster* runs (all *pARV5%* and *pARV10%*), and the remaining four were *three-cluster* runs (three *pARV10%*, one *pARV5%*). None of the *six-cluster* runs produced a viable ARV population. Whilst both the *one-cluster* and *three-cluster* scenarios produced an equal maximum viable population of 25% ARV (both *pARV10%*), the six successful *one-cluster* runs fared

better on average (16% ARV) than the four *three-cluster* runs did (13%).

The *range* runs showed significant variability with different ranges as well. Under *Scenario A* payoffs, ARVs fared best with a *range* of *two nodes* and worst with *six nodes*. Ten runs in total produced stable ARV populations—six *two-node* runs (all *pARV5%* and *pARV10%*), three *four-node* runs (all *pARV10%*) and one *six-node* run (*pARV10%*). A side-by-side comparison also shows that *pARV10%* produced higher stable ARV populations in *two-node* runs (9% average) than *four-node* runs (5% average) and the *six-node* run (4% average).

Finally, the *Scenario B range* runs were the only runs to produce a stable population of ARVs from a *pARV1%* start. Predictably, this result was obtained from a *two-node* run. 18 other runs (19 total) concluded with stable ARV populations. These were divided equally amongst the three *range* starting conditions, culminating in all *pARV5%* and *pARV10%* runs producing viable ARV populations. Side-by-side comparison shows a similar trend to the *Scenario A* runs in that *two-node* runs fared best, and *six-node* runs fared worst.

Overall, *Scenario B* runs fared better than the *Scenario A* ones in all nine different starting conditions in both the number of runs producing stable ARV populations and the final ARV population percentage in these runs.

Discussion

Several trends and observations can be established from the results discussed and shown.

Payoffs

The first observation of note is that unlike the theoretical profiles shown in Fig. 2, both payoff scenarios have yielded several evolutionarily stable ARV populations. This demonstrates the non-trivial role spatial distribution plays in the evolution of spatial populations. Such findings are echoed in the literature, e.g. [20, 21, 55, 57]. In this paper, we show that ARVs can under the right conditions form self-sufficient clusters where each individual ARV interacts with enough of its ARV neighbours to keep them sustainably viable, even if in a well-mixed population such viability would be impossible. This is reflected well in the positive correlation between individual cluster size and ARV success and the negative correlation between ARV range and ARV success. This is because both larger cluster sizes and lower ranges translate to a higher probability of interaction with other ARVs—the ideal opponent for an ARV.

Predictably, *Scenario B* payoffs produced better results for ARVs across the board. More importantly, however, the simulations show that even under significantly unfavourable

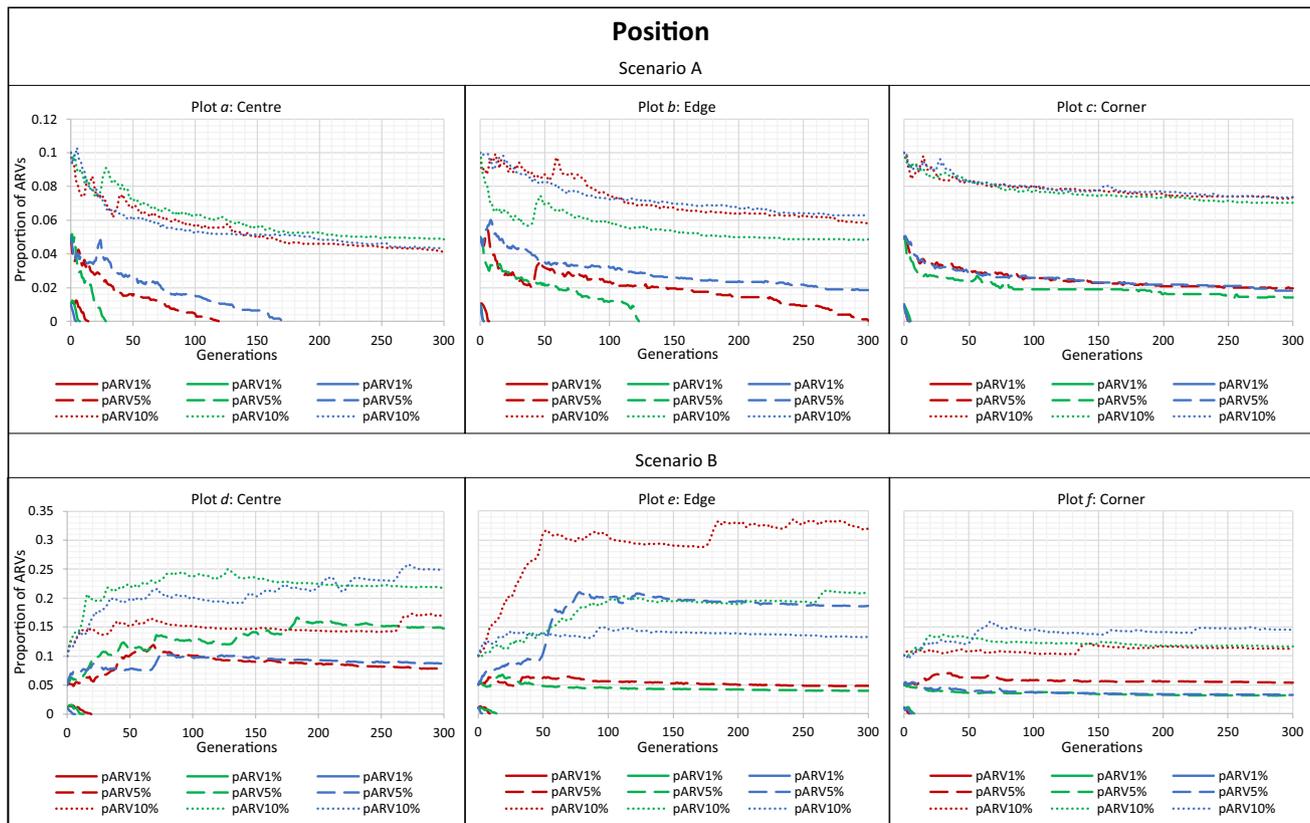


Fig. 5 Simulation results for the 54 runs examining the effect of cluster positioning on ARV success

conditions, as those afforded by the *Scenario A* payoffs, introducing, and maintaining viable ARV populations is not impossible.

Finally, lower ARV payoffs (as in *Scenario A*) are associated with significantly higher sensitivity to variation in the starting conditions. This is in line with [20]’s findings and is likely linked to the increased impact of such variations on the potential success of ARVs. Less favourable conditions would have their effect amplified by less forgiving payoff functions.

pARV

A higher introductory proportion of ARVs naturally translates to a better chance to develop a stable ARV population and a larger one at that. This is primarily due to the “strength in numbers” effect in play, as ARVs have a higher probability to interact with fellow ARVs and thus enjoy a higher payoff from these interactions. This, however, is likely to carry higher upfront implementation and coordination costs in the real world.

Position

The *position* conditions tested in this simulation relate to the “surface area” available for HDVs to interact with an ARV cluster. Vehicles from every node in the grid are free to travel to adjacent nodes and interact with the vehicles within them. This extends to interaction between HDVs entering an ARV cluster and an ARV from the destination node. The result from such an interaction is, of course, a high payoff for the HDV. By testing for the effect of “shielding” ARV clusters from one (*edge*) or two (*corner*) sides against HDV incursion, we can gain some insight on whether analogous arrangements in the real world would be beneficial for the introduction of ARVs. Such arrangements may take the form of “ARV-only” districts where HDV through-traffic is banned, thus allowing an introductory ARV population to mature under a far lighter pressure from HDVs.

The results show a significant improvement in ARV success along *edges* and at *corners* when payoff conditions are more unfavourable (*Scenario A*). This is in line with the principles described in the previous paragraph. In contrast, however, the effect seems to be reversed under *Scenario B* payoff conditions. This behaviour can be attributed to a side effect resulting from the spatial restriction of ARV

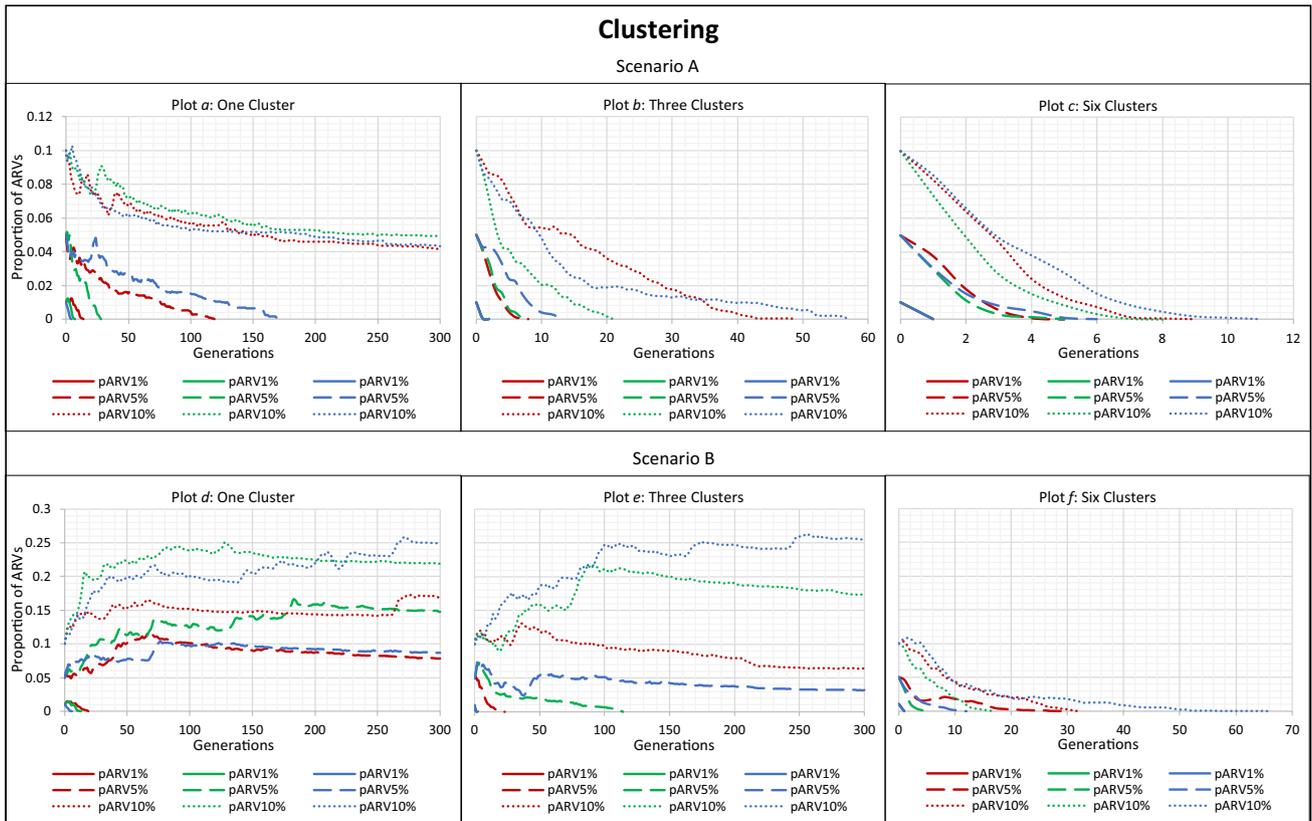


Fig. 6 Simulation results for the 54 runs examining the effect of the number of clusters on ARV success

growth beyond the grid boundaries. As explained earlier, this boundary helps protect ARVs when the conditions are unfavourable. However, under more favourable condition the boundary begins to restrict growth from the other side—the ARV side, which can hinder overall growth. *Centred* clusters would not suffer this limitation. This finding is of real-world relevance, though such boundaries would more likely be legislative rather than physical.

Clustering

A general inspection of the *clustering* graphs shows that multi-clustered populations generally perform worse than single-clustered ones. Interestingly, however, this inverse correlation has a notable exception. This is well illustrated by the *three-cluster* runs under *Scenario B* payoff conditions (Fig. 6, Plot *e*), where ARV success shows greater sensitivity to *pARV* than the *single-cluster* runs.

The reason for this is likely linked to the “strength in numbers” principle mentioned previously. At lower payoff values, dispersal of ARVs in multiple clusters forms a liability and exposes ARVs to a higher probability of interacting with non-ideal opponents (HDVs). Thus, coalescence into a single, geographically continuous unit makes better

evolutionary sense. When ARVs enjoy a relative improvement in payoff, however, multi-clustering can instead provide an avenue to multiply the number of simultaneously growing and self-sustaining ARV communities. Still, the dispersal does weigh lower-*pARV* populations down, as demonstrated by the much larger gap in Plot *e* between the *pARV10%* *three-cluster* runs and the single successful *pARV5%* *three-cluster* run compared to what is seen in the *single-cluster* counterparts in Plot *d*. These results suggest the existence of a *critical mass* below which ARVs simply cannot form a viable population under a given set of conditions. Our simulation results suggest that no fewer than 40 ARV nodes are needed in a single cluster to produce a viable ARV population under the standard *Scenario B*, *Range 4* conditions. This explains why none of the *pARV1%* runs discussed so far have produced a viable ARV population, as the maximum single-cluster size under such conditions would be 25—well below the observed *critical mass*.

Whilst the numerical values provided in this paper are mainly of demonstrative value, the principle of *critical mass* remains an important finding which should be considered in a real-world introduction.

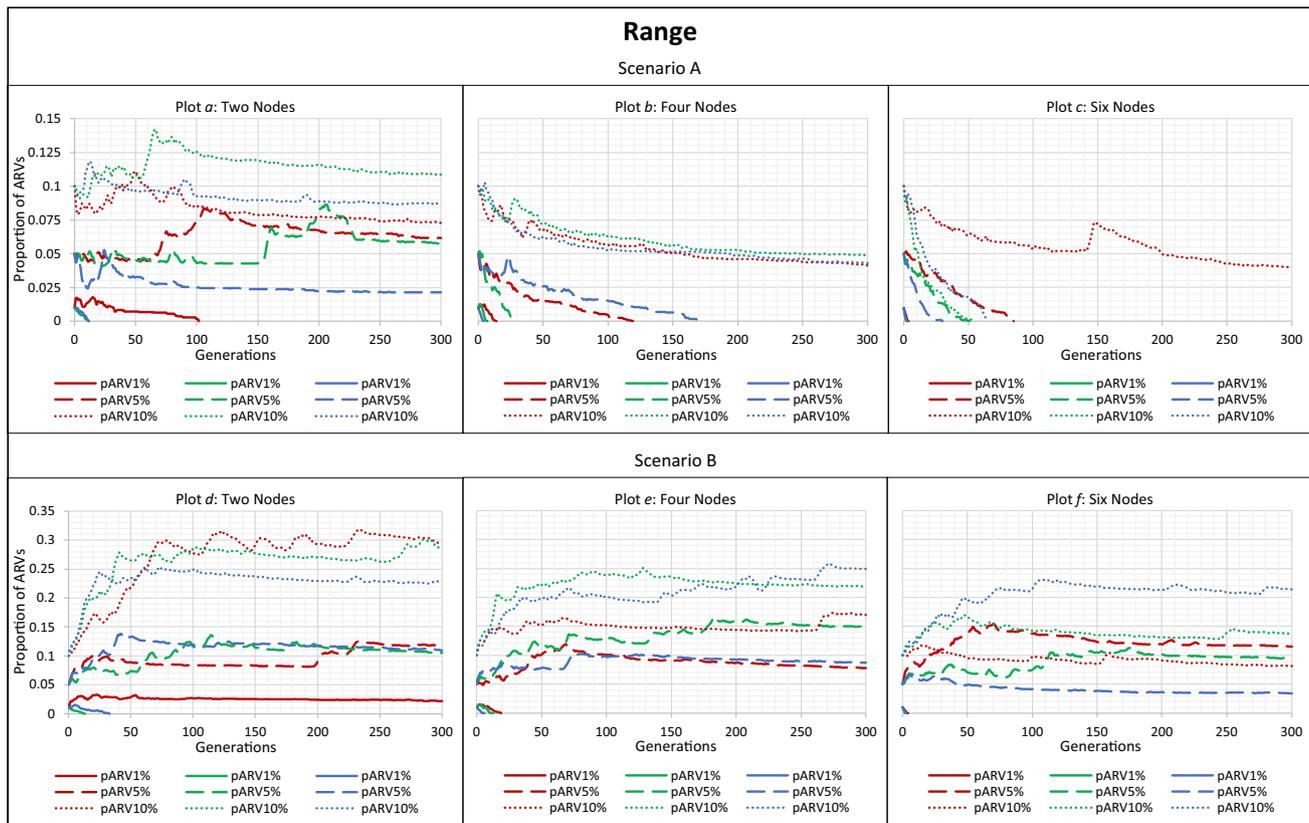


Fig. 7 Simulation results for the 54 runs examining the effect of ARV range on ARV success

Range

Range shows similar effects to *clustering* in that the higher the dispersal of ARVs, the lower the overall resultant fitness. This is especially pronounced under *Scenario A* pay-offs, where we go from six successful ARV runs at the *two-node* range (Fig. 7, Plot *a*) down to just one at *six nodes* (Plot *c*). The effect is less pronounced under *Scenario B* payoffs, though higher *pARV* rates show greater sensitivity to reduction in *range*. This is because the “strength in numbers” effect is compounded when larger clusters conduct their interactions closer to home. Therefore, we see largely comparable ARV success between the different *pARV* rates at *six-node* range (Plot *f*), whereas *pARV10%* runs show remarkable improvement over lesser *pARVs* at *four-* and *two-node* ranges (Plots *d* and *e*). These results make sense as longer ranges allow for higher probabilities of interaction with farther vehicles, thus diluting the effect of spatial heterogeneity. This means that longer ranges are closer in effect to the theoretical well-mixed populations than shorter ones.

These observations give evidence that the best approach to the introduction of ARVs may be a single, localised, short-range cluster where ARVs can maximise the “strength in numbers” effect observed and discussed throughout this

section. A practical application of this may take the form of a fleet of autonomous shuttles in a central business district or town centre, serving a circular line. Such application would ensure a large probability of ARV–ARV interactions. If coupled with a no-through-traffic restriction on HDVs, such introductions may prove effective and far more successful.

Conclusion

These conceptual experiments suggest a strong relationship between the spatial distribution, range, and size of an introductory population of ARVs and the shape of its evolution over time. The results show stronger performance for fewer, shorter-range clusters which suggests possible feasible applications of early ARVs in controlled, localised networks such as autonomous bus fleets in city centres. Dispersed, longer-range applications such as autonomous heavy goods vehicles may only be viable if deployed in larger numbers as closely co-operative platoons capable of generating collective efficiencies to counteract possible exploitation by human drivers. Future work may investigate a wider range of starting conditions, such as the shape of the clusters, the density of ARV-dominated nodes within each cluster, the effect of

population mixing, and whether external large-scale factors such as government subsidies could influence the resultant ARV populations.

The model parameters used in this study are demonstrative and further research will be required to characterise fitness parameters more accurately. However, there is value in testing the sensitivity of ARV success and the evolution of the road network against the variances in the different spatial parameters.

Funding This work was supported by the Engineering and Physical Sciences Research Council Doctoral Training Partnership, Grant No. EP/R513258/1.

Declarations

Conflict of Interest On behalf of all authors, the corresponding author states that there is no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Fox C, Camara F, Markkula G, Romano R, Madigan R, Merat N. When Should the Chicken Cross the Road? - Game Theory for Autonomous Vehicle—Human Interactions. In: Helfert M, Gusikhin O, editors. Proceedings of the 4th International Conference on Vehicle Technology and Intelligent Transport Systems: SciTePress; 2018. p. 431–9.
2. Millard-Ball A. Pedestrians, autonomous vehicles, and cities. *J Plan Educ Res*. 2018;38(1):6–12. <https://doi.org/10.1177/0739456x16675674>.
3. Cooper M, Lee JK, Beck J, Fishman JD, Gillett M, Papakipos Z, et al. Stackelberg Punishment and Bully-Proofing Autonomous Vehicles. In: Salichs MA, Ge SS, Barakova EI, Cabibihan J-J, Wagner AR, Castro-González Á, et al., editors. *Social Robotics*. Cham: Springer International Publishing; 2019. p. 368–77.
4. Karpus J, Krüger A, Verba JT, Bahrami B, Deroy O. Algorithm exploitation: Humans are keen to exploit benevolent AI. *iScience*. 2021;24(6):102679. doi: <https://doi.org/10.1016/j.isci.2021.102679>.
5. Elvik R. A review of game-theoretic models of road user behaviour. *Accid Anal Prev*. 2014;62:388–96. <https://doi.org/10.1016/j.aap.2013.06.016>.
6. van Loon RJ, Martens MH. Automated Driving and its Effect on the Safety Ecosystem: How do Compatibility Issues Affect the Transition Period? *Procedia Manufacturing*. 2015;3:3280–5. <https://doi.org/10.1016/j.promfg.2015.07.401>.
7. Meng F, Su J, Liu C, Chen W. Dynamic decision making in lane change: Game theory with receding horizon. 2016 UKACC 11th International Conference on Control (CONTROL)2016. p. 1–6.
8. Harris CM. Autonomous Vehicle Decision-Making: Should We Be Bio-inspired? In: Gao Y, Fallah S, Jin Y, Lekakou C, editors. *Towards Autonomous Robotic Systems*. Cham: Springer International Publishing; 2017. p. 315–24.
9. Liu C, Lin C, Shiraishi S, Tomizuka M. Improving Efficiency of Autonomous Vehicles by V2V Communication. 2018 Annual American Control Conference (ACC)2018. p. 4778–83.
10. Kang K, Rakha HA. A Repeated Game Freeway Lane Changing Model. *Sensors*. 2020;20(6):1554.
11. Smith JM, Price GR. The Logic of Animal Conflict. *Nature*. 1973;246(5427):15–8. <https://doi.org/10.1038/246015a0>.
12. Smith JM. *Evolution and the Theory of Games*. Cambridge: Cambridge University Press; 1982.
13. Wilkinson GS. Reciprocal food sharing in the vampire bat. *Nature*. 1984;308(5955):181–4. <https://doi.org/10.1038/308181a0>.
14. Bendor J, Swistak P. Types of evolutionary stability and the problem of cooperation. *Proc Natl Acad Sci USA*. 1995;92(8):3596–600. <https://doi.org/10.1073/pnas.92.8.3596>.
15. Dawkins R. *The selfish gene*. 4th ed. Oxford: Oxford University Press; 2016.
16. Bitar I, Watling D, Romano R. How Can Autonomous Road Vehicles Coexist with Human-Driven Vehicles? An Evolutionary-Game-Theoretic Perspective. Proceedings of the 8th International Conference on Vehicle Technology and Intelligent Transport Systems - VEHTS: SciTePress; 2022. p. 376–83.
17. Yaqoob I, Khan LU, Kazmi SMA, Imran M, Guizani N, Hong CS. Autonomous driving cars in smart cities: recent advances, requirements, and challenges. *IEEE Network*. 2020;34(1):174–81. <https://doi.org/10.1109/MNET.2019.1900120>.
18. Yvkoff L. The Success Of Autonomous Vehicles Hinges On Smart Cities. Inrix Is Making It Easier To Build Them. *Forbes Magazine*: Forbes; 2020.
19. FHWA FHA. 2020 NextGen NHTS National Passenger OD Data. Washington, DC.: U.S. Department of Transportation; 2020.
20. Killingback T, Doebeli M. Spatial evolutionary game theory: Hawks and Doves revisited. *Proc R Soc Lond B*. 1996;263(1374):1135–44. <https://doi.org/10.1098/rspb.1996.0166>.
21. Hauert C, Doebeli M. Spatial structure often inhibits the evolution of cooperation in the snowdrift game. *Nature*. 2004;428(6983):643–6. <https://doi.org/10.1038/nature02360>.
22. Thalya P, Kovaceva J, Knauss A, Lubbe N, Dozza M. Modeling driver behavior in interactions with other road users. 2020.
23. Michon JA. A Critical View of Driver Behavior Models: What Do We Know, What Should We Do? In: Evans L, Schwing RC, editors. *Human Behavior and Traffic Safety*. Boston, MA: Springer, US; 1985. p. 485–524.
24. Johora FT, Müller JP. Modeling Interactions of Multimodal Road Users in Shared Spaces. 2018 21st International Conference on Intelligent Transportation Systems (ITSC)2018. p. 3568–74.
25. Kita H. A merging–giveaway interaction model of cars in a merging section: a game theoretic analysis. *Transportation Research Part A: Policy and Practice*. 1999;33(3):305–12. [https://doi.org/10.1016/S0965-8564\(98\)00039-1](https://doi.org/10.1016/S0965-8564(98)00039-1).
26. Liu, Xin, Adam, Ban X. A game theoretical approach for modeling merging and yielding behavior at freeway on-ramp section. Proceedings of the 17th International Symposium on Transportation and Traffic Theory2007. p. 197–211.
27. Kim C, Langari R. Game theory based autonomous vehicles operation. *Int J Veh Des*. 2014;65(4):360. <https://doi.org/10.1504/ijvd.2014.063832>.
28. Kang K, Rakha HA. Game Theoretical Approach to Model Decision Making for Merging Maneuvers at Freeway On-Ramps.

- Transp Res Rec. 2017;2623(1):19–28. <https://doi.org/10.3141/2623-03>.
29. Camara F, Romano R, Markkula G, Madigan R, Merat N, Fox C. Empirical game theory of pedestrian interaction for autonomous vehicles. In: Grant R, Allen T, Spink A, Sullivan M, editors. Proceedings of Measuring Behavior 2018: 11th International Conference on Methods and Techniques in Behavioral Research: Manchester Metropolitan University; 2018. p. 238–44.
 30. Yu H, Tseng HE, Langari R. A human-like game theory-based controller for automatic lane changing. *Trans Res Part C*. 2018;88:140–58. <https://doi.org/10.1016/j.trc.2018.01.016>.
 31. Wu W, Liang Z, Luo Q, Ma F. Game theory modelling for vehicle U-Turn behavior and simulation based on cellular automata. *Discret Dyn Nat Soc*. 2018;2018:5972495. <https://doi.org/10.1155/2018/5972495>.
 32. Camara F, Dickinson P, Merat N, Fox CW. Towards game theoretic AV controllers: measuring pedestrian behaviour in Virtual Reality. Proceedings of TCVC2019: Towards Cognitive Vehicles: IROS; 2019. p. 7–10.
 33. Li N, Oyler D, Zhang M, Yildiz Y, Girard A, Kolmanovsky I. Hierarchical reasoning game theory based approach for evaluation and testing of autonomous vehicle control systems. 2016 IEEE 55th Conference on Decision and Control (CDC)2016. p. 727–33.
 34. Oyler DW, Yildiz Y, Girard AR, Li NI, Kolmanovsky IV. A game theoretical model of traffic with multiple interacting drivers for use in autonomous vehicle development. 2016 American Control Conference (ACC)2016. p. 1705–10.
 35. Fisac J, Bronstein E, Stefansson E, Sadigh D, Sastry S, Dragan A. Hierarchical Game-Theoretic Planning for Autonomous Vehicles. 2018.
 36. Wang M, Hoogendoorn SP, Daamen W, van Arem B, Happee R. Game theoretic approach for predictive lane-changing and car-following control. *Transportation Research Part C: Emerging Technologies*. 2015;58:73–92. <https://doi.org/10.1016/j.trc.2015.07.009>.
 37. Michieli U, Badia L. Game Theoretic Analysis of Road User Safety Scenarios Involving Autonomous Vehicles. 2018 IEEE 29th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)2018. p. 1377–81.
 38. Camara F, Fox C. Space Invaders: Pedestrian Proxemic Utility Functions and Trust Zones for Autonomous Vehicle Interactions. *Int J Soc Robot*. 2021;13(8):1929–49. <https://doi.org/10.1007/s12369-020-00717-x>.
 39. Heymann M, Degani A. Autonomous Vehicle Interactions with Other Road Users: Conflicts and Resolutions. 2019.
 40. Helbing D, SchÖnhof M, Stark H-U, Holyst J. How Individuals Learn to Take Turns: Emergence of Alternating Cooperation in a Congestion Game and the Prisoner's Dilemma. *Adv Complex Systems (ACS)*. 2005;08:87–116. <https://doi.org/10.1142/S0219525905000361>.
 41. Jiang X, Ji Y, Du M, Deng W. A Study of Driver's Route Choice Behavior Based on Evolutionary Game Theory. *Computational Intelligence and Neuroscience*. 2014;2014:124716. doi: <https://doi.org/10.1155/2014/124716>.
 42. Wu C, Pei Y, Gao J. Evolution Game Model of Travel Mode Choice in Metropolitan Beijing. *Discrete Dynamics in Nature and Society*. 2015;2015:638972. <https://doi.org/10.1155/2015/638972>.
 43. Alibabai H, Mahmassani HS. Foxes and sheep: effect of predictive logic in day-to-day dynamics of route choice behavior. *EURO Journal on Transportation and Logistics*. 2016;5(1):53–67. <https://doi.org/10.1007/s13676-015-0088-2>.
 44. Lei L, Gao S. Transportation network companies and drivers dilemma in China: an evolutionary game theoretic perspective. *Transport*. 2019;34:1–12. <https://doi.org/10.3846/transport.2019.11105>.
 45. Wang S, Fan J, Zhao D, Wu Y. The Impact of Government Subsidies or Penalties for New-energy Vehicles A Static and Evolutionary Game Model Analysis. *Journal of Transport Economics and Policy (JTEP)*. 2015;49(1):98–114.
 46. Zhang L, Long R, Huang Z, Li W, Wei J. Evolutionary game analysis on the implementation of subsidy policy for sustainable transportation development. *Journal of Cleaner Production*. 2020;267:122159. <https://doi.org/10.1016/j.jclepro.2020.122159>.
 47. Chatterjee I, Davis GA. Evolutionary game theoretic approach to rear-end events on congested freeway. *Transp Res Rec*. 2013;2386(1):121–7. <https://doi.org/10.3141/2386-14>.
 48. Cortés-Berruero LE, Gershenson C, Stephens CR. Traffic Games: Modeling Freeway Traffic with Game Theory. *PLOS ONE*. 2016;11(11):e0165381. <https://doi.org/10.1371/journal.pone.0165381>.
 49. Iwamura Y, Tanimoto J. Complex traffic flow that allows as well as hampers lane-changing intrinsically contains social-dilemma structures. *Journal of Statistical Mechanics: Theory and Experiment*. 2018;2018:023408. <https://doi.org/10.1088/1742-5468/aaa8ff>.
 50. Free C. How does Aggressive Driving Respond to Passenger Load and Type. 2018.
 51. Brown DB, Hansell RIC. Convergence to an evolutionarily stable strategy in the two-policy game. *Am Nat*. 1987;130(6):929–40.
 52. Nowak MA, May RM. Evolutionary games and spatial chaos. *Nature*. 1992;359(6398):826–9. <https://doi.org/10.1038/359826a0>.
 53. Yang H-X, Yang J. Cooperation percolation in spatial evolutionary games. *EPL (Europhysics Letters)*. 2019;124(6):60005. <https://doi.org/10.1209/0295-5075/124/60005>.
 54. Sakiyama T. A power law network in an evolutionary hawk–dove game. *Chaos, Solitons & Fractals*. 2021;146:110932. <https://doi.org/10.1016/j.chaos.2021.110932>.
 55. He J, Zhao Y, Cai H, Wang R. Spatial games and the maintenance of cooperation in an asymmetric Hawk-Dove game. *Chin Sci Bull*. 2013;58(18):2248–54. <https://doi.org/10.1007/s11434-013-5810-6>.
 56. Wang X-W, Nie S, Jiang L-L, Wang B-H, Chen S-M. Cooperation in spatial evolutionary games with historical payoffs. *Phys Lett A*. 2016;380(36):2819–22. <https://doi.org/10.1016/j.physleta.2016.06.026>.
 57. Roca CP, Cuesta JA, Sánchez A. Evolutionary game theory: temporal and spatial effects beyond replicator dynamics. *Phys Life Rev*. 2009;6(4):208–49. <https://doi.org/10.1016/j.plrev.2009.08.001>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.