Randomized routing strategies of fleets of CAVs may prove market efficient

Grzegorz Jamróz*[1], Łukasz Gorczyca[1], and Rafał Kucharski[1]

[1]Faculty of Mathematics and Computer Science, Jagiellonian University, Kraków, Poland

*Corresponding author, e-mail: grzegorz.jamroz@uj.edu.pl

Short summary

In future cities every driver may own a vehicle which could be either independently driven (HDV), or autonomously routed and piloted (CAV). The autonomous operations could be handled by a few competing companies. What is the market structure which would make this market aligned with city goals? In this paper we discuss a variant of the emerging market of collectively routed fleets of CAVs, where revenue for fleet operators is proportional to market share. We provide benchmark scenarios to compare the routing algorithms. We present several routing algorithms and demonstrate that, when the attitudes of human drivers towards CAVs exhibit significant diversity, randomised CAV routing, resulting in unpredictable travel times for HDVs, is more efficient than routing proportional to system optimum/user equilibrium. Based on this, we propose to improve the design of the market by augmenting the market-share objective with mean systemwide travel time in order to limit antisocial randomised strategies of fleet operators and drive the competition towards social welfare oriented cooperation.
Keywords: Autonomous driving and routing, choice modelling, market competition, randomized routing

1 Introduction

Design of a fair, efficient and competitive market of urban CAV operations is likely to be a challenge, like for other markets, compare Kominers et al. (2017). The unfair predatory pricing policies of companies backed by huge venture capital like Amazon or Uber, e.g. Khan (2016), resulting in long-term consumer harming mono/duopolies, even though common, are undesirable and preventable. As the market share of CAVs in major cities such as Los Angeles and Beijing is still too small, e.g. GoldmanSachs (2024), to significantly influence the urban routing market, we are in a position to design the market structure which would be most beneficial for cities and avoid learning from failures, see e.g. Roth (2018). Accordingly, in this paper we consider one variant of CAV autonomous-routing market where the revenue is derived only from market-share and not subscription/purchase cost. This potential future market is based on the following assumptions:

- Every driver owns a vehicle which can be either driven independently (HDV) or routed and navigated by one of several providers (CAV fleet operators).
- Autonomous driving is indistinguishable from human driving in traffic.
- Each driver can freely choose/switch between HDV and one of CAV fleets.
- Fleet operators are remunerated collectively from traffic taxes proportionally to the number of users (market-share maximization).

The above 'fee-free' design favours algorithmic routing and quality advancements over pure discount-based monopole building strategies which benefit from economy-of-scale effects, e.g. Cramer & Krueger (2016), and weak regulation, Collier et al. (2018). We study the dynamics of this market in a simplified case of one OD pair with two parallel routes. Fleet operators aim at maximum market share. Why would, however, a driver switch from independent driving (HDV) to autonomous driving and routing (CAV)? Three levels of incentives are presented in Fig. 1:

i) People's travel time cost spent in an AV is less than in an HDV. This is likely to be the case for most drivers, see de Almeida Correia et al. (2019). These customers are 'taken for free' by fleet operators by using user equilibrium (UE) routing.

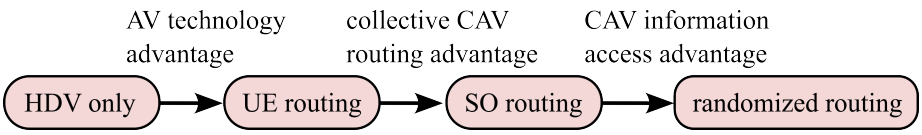


Figure 1: Incentives to switch to CAV. Compared to independent driving, AV technology offers better value of time for most drivers so that UE routing is enough to convince to switch to CAV. If this is not enough, leveraging collective market power via SO routing may convince some drivers. Finally, randomized routing which strips HDV drivers of the information which route will be fastest on a given day, may convince the most reluctant.

ii) Collective routing can improve average travel times compared to user equilibrium by routing according to System Optimum (SO) with turn-taking for equity, see Hoffmann et al. (2025) and references therein.

iii) SO routing is unstable as every driver may be inclined to defect and use the faster route. Moreover, drivers originally reluctant to use AVs will still not be convinced. Randomized routing addresses these two issues by deliberately introducing uncertainty and depriving independent drivers of the information which route is likely to be fastest on a given day.

To study the above mentioned mechanisms and design a reasonable market structure, we:

- Introduce a framework(benchmark) for comparing algorithms for routing fleets of CAVs aimed at maximizing market share, including both baseline algorithms and more sophisticated heuristics.
- Introduce dynamic randomized algorithms for routing and demonstrate experimentally that they are strictly superior to non-randomized algorithms; note that randomization in the static full market-share context was discussed in Jamróz et al. (2025).
- Show that including average system travel time minimization in the objective renders randomized algorithms less efficient and favors algorithms more aligned with city goals, while preserving competition.

2 Methodology

System dynamics

We present benchmark scenarios modelling the competition of fleet operators in systems consisting, for simplicity, of one OD pair and parallel routes $r = 1, \ldots, R$. We consider a fixed number N_D drivers, who every day decide whether to drive independently or join one of N_{FLEET} fleets of autonomously routed vehicles, based on travel times offered by each of the fleets and credibility of fleet offers. Consequently, the system can be considered a game with $N_D + N_{FLEET}$ players, who make decisions as follows. On a given day j, sequentially, see Fig. 2:

i) Each fleet operator $f \in \{0, 1 \ldots, N_{FLEET}-1\}$ creates an offer vector $O_i^{f,j}$ for $i = 0, 1, \ldots, N_D - 1$ based on some algorithm $\mathbf{Algo^f_{offer}}$ such that $O_i^{f,j}$ is the mean long-term travel time fleet f offers to driver i.

ii) Each driver $i = 0, 1, \ldots, N_D - 1$ decides, based on previous travel experience and offers from fleet operators whether to remain independent (HDV) or be part of one of the fleets, resulting in mode vector $M^j \in \{HDV, 0, 1, \ldots, N_{FLEET} - 1\}^{N_D}$, where M_i^j is the the mode chosen by driver i on day j.

iii) Once the drivers have committed to fleets or to staying independent HDVs, each fleet f creates, based on some algorithm $\mathbf{Algo^f_{routing}}$, a routing vector $[r^{f,j}]_i$ for drivers $i \in I^{f,j} := \{i : M_i^j = f\}$, i.e drivers who have chosen fleet f on day j. The value $r_i^{f,j} \in \{1, \ldots, R\}$ corresponds then to the route via which driver i will be routed on day j. The drivers i who have decided to remain independent, choose routes based on minimizing disutility, see below, resulting in vector $r_i^{HDV,j}$. The resulting routing vector on day j is given by $r_i^j = r_i^{M_i^j,j}$ such that r_i^j is the route via which driver i will go on day j.

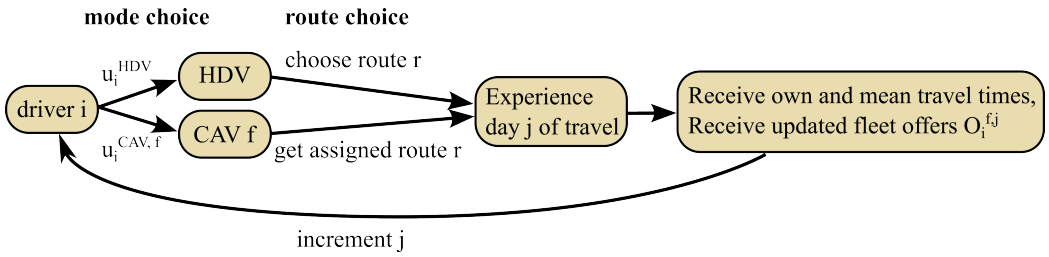


Figure 2: Dynamics of the system at a glance. Every day each driver i selects the mode (HDV or CAV f) and route based on utilities derived from experienced and mean travel times and fleet operators' travel time offers.

After all the routing decisions have been made, a day of travelling is simulated, resulting in route travel time t_r^j, for $r = 1, \ldots, R$. The travel times are assumed to be independent of driver identity, however they depend on congestion and in our considerations are given by the BPR function:

$$t^{BPR}(q) = (5min) * \left(1 + (q/(0.5 * N_D))^2\right) \tag{1}$$

where q is the number of vehicles driving via the route on a given day.

Human behaviour

We assume that each driver i has, for each fleet f, a discount factor $\gamma_i^{F,f} > 0$, which models the attitude towards using fleet f. This attitude can be decomposed into a part corresponding to the general liking of using fleets, γ_i^F (drawn from a given known distribution, in our simulations assumed to be gaussian, clipped at 0, centred at 0.7 with standard deviation 0.2, compare de Almeida Correia et al. (2019)) and a factor specific to a given fleet $\gamma_i^{F,specf}$, drawn from $N(0, (0.15)^2)$ (gaussian centred at 0 and standard deviation 0.15), resulting in $\gamma_i^{F,f} = \gamma_i^F + \gamma_i^{F,specf}$. The utility of driving independently u_i^{HDV} is given by the mean travel times over the past j_{avg} days, sampled uniformly from $\{1, 2, \ldots, 9\}$, weighted with logit probabilities. Contrariwise, the utility $u_i^{CAV,f}$ of using fleet f is proportional to the travel time offered multiplied by the discount factor γ_i^f and divided by credibility . Consequently, on day j, for $\overline{t_r} = \sum_{\iota=j-j_{avg}}^{j-1} t_r^\iota$, disutilities of using different modes are:

$$u_i^{HDV} = \sum_{r \in R} e^{-\beta \overline{t_r}} \overline{t_r} / \sum_{r \in R} e^{-\beta \overline{t_r}} \tag{2}$$

$$u_i^{CAV,f} = \gamma_i^{F,f} * O_i^{f,j} / Cred_i^f, \tag{3}$$

where $O_i^{f,j}$ is the travel time offered by fleet f to driver i and $Cred_i^f > 0$ is the credibility of fleet f to driver i. The initial credibility of fleets is assumed to be 1 and it is updated by

$$Cred_i^f \leftarrow (1 - \alpha) * Cred_i^f + \alpha * O_i^{f,j} / t_{r_i^j}^j \tag{4}$$

for some update rate α (which in this paper is assumed to equal 0.2), if driver i chose fleet f on the previous day j and was routed via r_i^j. This means that human drivers are inclined to try out different fleets initially, provided the (honest or dishonest) travel time offer $O_i^{f,j}$ is attractive enough. However, if the offers are later not confirmed in the real driving, the credibility of a given fleet decreases. The gradual update of credibility is important as the fleets are supposed to offer long-term average travel times and not exact travel times on the following day.
Finally, driver i is assumed to choose the mode (HDV or one of the fleets f) with lowest disutility. If HDV is chosen, the route choice is based on logit choice of routes, i.e. probability of choosing route r is given by $e^{-\beta \overline{t_r}} / \sum_{r \in R} e^{-\beta \overline{t_r}}$, where we assume $\beta = 0.2$. The rationale for choosing this model is that in the future even HDVs are likely to use some apps to propose routes, and a simple way to choose roughly optimal routes on a population level with small error is to use logit probabilities with small β, while mode choice, which would perhaps follow some bounded rationality model with near-deterministic logit probabilities can be approximated by deterministic mode choice.

Algorithms for CAV fleet behaviour

Algorithm 1: Proportional SO router

This baseline algorithm offers routing aimed at achieving System Optimum. To reach this goal it first determines the system optimal flows $\mathbf{q^{SO}} = (q_1^{SO}, \ldots, q_r^{SO})$, which is straightforward in the case of two parallel routes as well as the average travel time at System Optimum, t_{avg}^{SO}. $\mathbf{Algo^f_{offer}}$ is based on offering every driver average system optimal travel time, i.e.

$$O_i^{f,j} = t_{avg}^{SO}$$

for every day j and driver i. $\mathbf{Algo^f_{routing}}$ is the following. Given the set of drivers $I^{f,j}$ that have chosen fleet f on day j the drivers are split randomly with probabilities proportional to $\mathbf{q^{SO}}$. In the case of two available routes, it consists of selecting a uniformly random sample I^{rand} of $\left\lfloor |I^{f,j}| * q_1^{SO}/|\mathbf{q^{SO}}| \right\rfloor$ agent indices from $I^{f,j}$ and setting

$$r_i^{f,j} = \begin{cases} 0 & \text{for } i \in I^{rand}, \\ 1 & \text{for } i \in I^{f,j} \backslash I^{rand}. \end{cases}$$

The rationale behind this algorithm is that if all the drivers choose this fleet, routing according to system optimal assignment is the most efficient way to minimize average travel time, which is desirable. The downside is that this routing treats all the drivers in the same way and disregards different attitudes towards the fleet, making the routing potentially inefficient. Furthermore, in the case of non-full market penetration, the offered travel times cannot in fact be realized, which decreseas the credibility of the fleet. A variant of the algorithm consists in offering travel times lower than system optimum, i.e.

$$O_i^{f,j} = \kappa t_{avg}^{SO} \tag{5}$$

for some $\kappa \in (0, 1]$ in order to win over customers at the beginning.

Algorithm 2: Proportional UE router

This algorithm is analogous to Proportional SO routes, with $\mathbf{q^{SO}}$ replaced by $\mathbf{q^{UE}}$, i.e. user-equilibrium flows and setting $O_i^{f,j} = t^{UE}$, where t^{UE} is the travel time via both routes at User Equilibrium. We note that in the case of two equivalent routes, SO and UE coincide and there is no distinction between the two algorithms.

Algorithm 3: Randomized algorithm RFlexV

The standard offer algorithm $\mathbf{Algo^f_{offer}}$ is to offer the scaled average system optimal travel time (5) with $\kappa = 0.5$ (by default). The routing is given by the following randomized $\mathbf{Algo^f_{routing}}$ algorithm. Let $|I^{f,j}|$ be the number of drivers who have chosen fleet f on day j. For two equivalent routes $0, 1$ the fleet operator proposes a symmetric randomized strategy such that

- all the current members of the fleet are happy (i.e. will not want to leave the fleet to become an HDV on the next day),
- the strategy is as randomized as possible in order to make independent routing difficult for other drivers and coax them to switch to the fleet.

To achieve this, we assume that the remaining drivers will split 50-50 (as the routing is symmetric) and choose

$$n^{faster*} \in \arg\min \left\{ N_{unhappy}(n_{faster}) : n_{faster} = 1, 2, \ldots, \lfloor |I^{f,j}|/2 \rfloor \right\}, \tag{6}$$

where $N_{unhappy}(n_{faster})$ stands for the number of drivers who would rather switch to HDV than stay with the fleet when the fleet puts n_{faster} agents with highest discount factors on the faster route and is given by

$$N_{unhappy}(n_{faster}) = \left| \left\{ i \in I^{f,j} : \gamma_i^{F,f} t^{sim}_{r_i^{f,j}} > \overline{t_r^{sim}} \right\} \right|,$$

where $\overline{t_r^{sim}} := \frac{1}{N_R} \sum_{r=1}^{N_R} t_r^{sim}$ and t_r^{sim} is the simulated travel time of route r provided $n_{faster} < |I^{f,j}|/2$ fleet members are routed, without loss of generality, via route 0 (which will be the faster

route) and $|I^{f,j}| - n_{faster}$ drivers are put on the other (slower) route 1. We obtain:

$$\begin{aligned} t_0^{sim} &= t_0\left(n_{faster} + \lfloor (N_D - |I^{f,j}|)/2 \rfloor\right) \\ t_1^{sim} &= t_1\left(|I^{f,j}| - n_{faster} + \lceil (N_D - |I^{f,j}|)/2 \rceil\right) \end{aligned}$$

and $t_0 = t_1$ are the known functional dependences of travel time on flow, see (1).
Now, let $n^{faster*}$ be the smallest such that equation (6) is satisfied. The algorithm assigns $n^{faster*}$ drivers with highest discount factors to a random route and $|I^{f,j}| - n^{faster*}$ drivers to the alternative route, i.e. for r^* - a random number from $\{0, 1\}$ we have

$$r_i^{f,j} = \begin{cases} r^* & \text{if } s(i) < n^{faster*}, \\ 1 - r^* & \text{otherwise,} \end{cases}$$

where

$$s : \{0, 1, \ldots, |I^{f,j}| - 1\} \to I^{f,j} \tag{7}$$

is a sorting function such that $\gamma^{F,f}_{s(0)}, \ldots, \gamma^{F,f}_{s(|I^{f,j}|)}$ is a non-decreasing sequence (note s depends on $I^{f,j}$). The algorithm tries to maintain its current customer base while randomizing the assignment so that it is less convenient to remain an independent driver for discount factors $\gamma_i^{F,f} > 1$. This, however, (even in the case of one fleet) may be suboptimal, as a driver i with $\gamma_i^{F,f} > 1$ does not have to be routed via the faster route every day, but for instance only on 80% of days, see Jamróz et al. (2025).

Algorithm 4: RFlex

RFlex is an extension of RFlexV, optimizing the choice of drivers to be routed via the faster route so that drivers who do not need to be routed via the faster route are not routed via it. To this end, it computes the least necessary share sh_i of routing driver i via the faster route to keep i happy as:

$$sh_i^* = \min\left\{ s : \gamma_i^{F,f}(s t_0^{sim} + (1 - s) t_1^{sim}) < \overline{t_r^{sim}} \right\}.$$

Solving the above equation we obtain:

$$sh_i^* = \frac{t_1^{sim} - \overline{t_r^{sim}}/\gamma_i^{F,f}}{(t_1^{sim} - t_0^{sim})}$$

with the caveat that if $sh_i^* < 0$ then driver i will always be happy (we put $sh_i^* = 0$) and when $sh_i^* > 1$ then driver i cannot be made happy (we put $sh_i^* = \infty$, because this driver cannot be convinced anyway for the given n_{faster}). Let $s(i)$ be the sorting function from (7). The theoretical maximum number of drivers that can be made happy is given by

$$n_{maxhappy} = max\left\{ n : \frac{1}{n} \sum_{i=0}^{n-1} sh_{s(i)}^* < \frac{n_{faster}}{n} \right\}.$$

The routing should theoretically now be given by randomly assigning the drivers to routes by $r_i^{f,j}$ such that

$$\mathbb{E} r_{s(i)}^{f,j} < sh_{s(i)}^* \tag{8}$$

for every $i = 0, \ldots, n-1$, where, without loss of generality 0 is the faster route and 1 is the slower route. The overall routing could consist in finding n^{faster} such that $n_{maxhappy}$ is the highest and routing such that (8) is satisfied. This would work if human drivers fully trusted the fleet operator to deliver the offered travel times. However, in practice, the drivers might be discouraged if they are not offered good travel times initially after joining the fleet and, moreover, the number of participants of the fleet changes making the theoretical delivery of offered travel times tricky. In view of the above, we opt for a heuristic $\mathbf{Algo}^{\mathbf{f}}_{\mathbf{routing}}$ given by setting the target ratio of routing every driver via the faster route as $sh_i^{target} = sh_i^*/\sigma$, where $\sigma \in (0, 1]$ is a parameter to optimize (we use $\sigma = 0.4$). We obtain:

$$r_i^{f,j} = \begin{cases} r^* & \text{if } \frac{d_i^{faster}}{d_i^{faster} + d_i^{slower} + 1} > sh_i^{target}, \\ 1 - r^* & \text{otherwise,} \end{cases}$$

where d_i^{faster} and d_i^{slower} is the number of days driver i was routed via the faster (slower) alternative since joining the fleet and r^* is a random number from $\{0, 1\}$.

Estimation of human drivers' discount factors

Across most scenarios, we assume that the discount factors are known to the fleets.

3 Results and discussion

In this paper we consider the simplest benchmark scenario of two equivalent routes and discount factors known to fleet operators, leaving extensions with more realistic details to further work. The parameters used are listed in Table 1.

Table 1: Simulation parameters

Parameter	Symbol	Value
Number of drivers	N_D	200
Number of days	J	300
Route 0 delay function	$t_0(q_0)$	t^{BPR}
Route 1 delay function	$t_1(q_1)$	t^{BPR}
Initial fleets' credibility	$Cred_{init}$	1
Fleet discount factor distribution	γ^{FD}	$N(0.7, (0.2)^2)$
Fleet specific discount factor term	$\gamma^{F,specfD}$	$N(0, (0.15)^2)$
Discount factors known to Fleets?	$DiscF_{known}$	True
Discount factor distribution known?	$DiscD_{known}$	True
HDV logit parameter	β	0.2

The execution of the benchmark scenarios proceeds as follows.

1. A list of N_D human drivers is created. At the beginning, each driver $i \in \{0, \ldots, N_D - 1\}$ has the initial credibility towards the fleets equal $Cred_{init}$ and a couple of discount factors $\gamma_i^{F,0}, \gamma_i^{F,1}$, where $\gamma_i^{F,f} = \xi_i + \xi_{i,f}$ for $f = 0, 1$ where ξ_i a number drawn from γ^{FD} and $\xi_{i,f}$ drawn independently from $\gamma^{F,specfD}$ for $f = 1, 2$. If any of the resulting factors is ≤ 0 then it is set to ε, where $\varepsilon = 0.0001$ is a very small positive number. In this way, the attitude of a human towards fleet f is broken down into the general attitude towards using fleets of AVs and fleet specific components.
2. Fleets f for $f \in \{0, 1\}$ are created. Every fleet uses one of the algorithms specified in Table 2 and described in detail in Section *Methodology - Algorithms for CAV fleet behaviour.*
3. The dynamics of the system are simulated in line with Section *Methodology - System Dynamics.* Various parameters are recorded.
4. The parameters like market share of a fleet are plotted and compared to other algorithms.

Table 2: Algorithms used

Algorithm	Offer creation $\mathbf{Algo^f_{offer}}$	Assignment method $\mathbf{Algo^f_{routing}}$
SO	Avg SO travel time	Proportional to SO + random turn-taking
SO-	0.8∗ Avg SO travel time	Proportional to SO + random turn-taking
UE	Avg UE travel time	Proportional to UE
UE-	0.5∗ Avg UE travel time	Proportional to UE
RFlexV	Avg SO travel time	highest γ^F to (randomized) faster route
RFlexV-	0.5∗ Avg SO travel time	highest γ^F to (randomized) faster route
RFlex	Avg SO travel time	Adaptive randomized routing to make happy
RFlex-	0.5∗ Avg SO travel time	Adaptive randomized routing to make happy
Infty	Infinite	No assignment, used for one-fleet scenarios
Your Algorithm	?	?

Experiment 1: Offering realistic travel times vs. unrealistically short travel times

Fig. 3 shows the performance of different algorithms in the variants with best realistically possible mean travel times offered (without '-') and with offered travel times much faster than achievable (with '-'). For SO the advantage for '-' only persistis temporarily whereas for RFlex and RFlexV it is long-term. Moreover the logit parameter of human assignment equal 0.2 offers better, stable travel times, which is preferred over every driver striving to choose the faster route every time which results in high variance (P = 1.0). Therefore, for further experiments we choose to set the logit parameter to 0.2 and study only algorithms offering unrealistically short travel times.

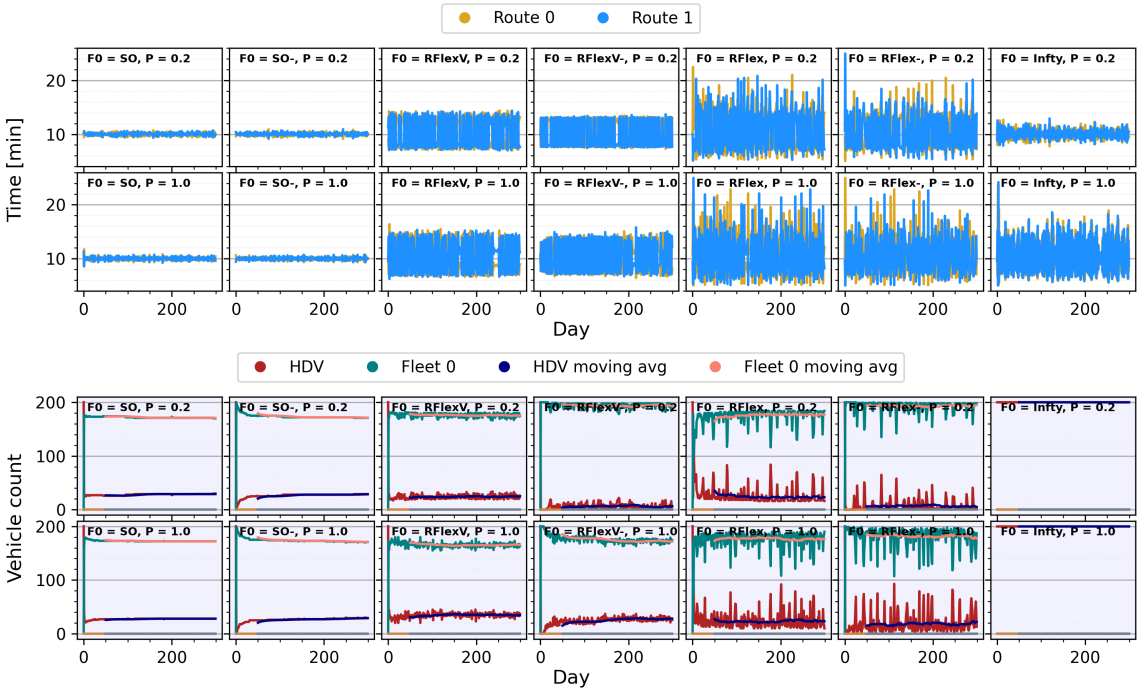


Figure 3: Comparison of properties of the system for different logit parameters of human driver assignment (P) and routing algorithms (F0) with travel time offers either unrelistically low (with '-') or more realistic (without '-'). Upper panel: Logit parameter 0.2 is a preferred choice for the human population as it results in relatively stable assignment when Fleet is absent (F0 = infty). Lower panel: The performance of '-' algorithms is superior to those offering more realistic travel times.

Experiment 2: Randomized algorithms vs. baseline algorithms

In Fig. 4 we pitch the algorithms against each other. The outcomes reveal that RFlexV is superior to SO and RFlex is superior to RFlexV when there is one fleet (last row/last column in bottom panel). However, when there are two fleet operators, the results are mixed, with RFlex/RFlexV performing slightly better than SO against SO (first row/first column, bottom panel), however not necessarily so against either RFlexV or RFlex (four central subplots). The upper panel demonstrates that randomized strategies entail large variations of travel times and the increase of mean travel time in the system. Fig. 5 shows which vehicles belong to which fleet on a sample day 150 plotted against discount factors towards both fleets. The baseline natural split is for both fleets using SO algorithms - drivers join the fleet with better discount factor provided it is < 1. This changes when more advanced algorithms are used. E.g. in the last row it can be seen how RFlexV/RFlex can convince even drivers who, without randomization would not join the fleet. Moreover, if a fleet uses a randomized algorithm then the other fleet sticking to the baseline SO algorithm fails to be competitive (first row and left column) and has to switch to a different algorithm. The four central panels illustrate continual competition for customers. The sample of discount factors is the same for every subplot.

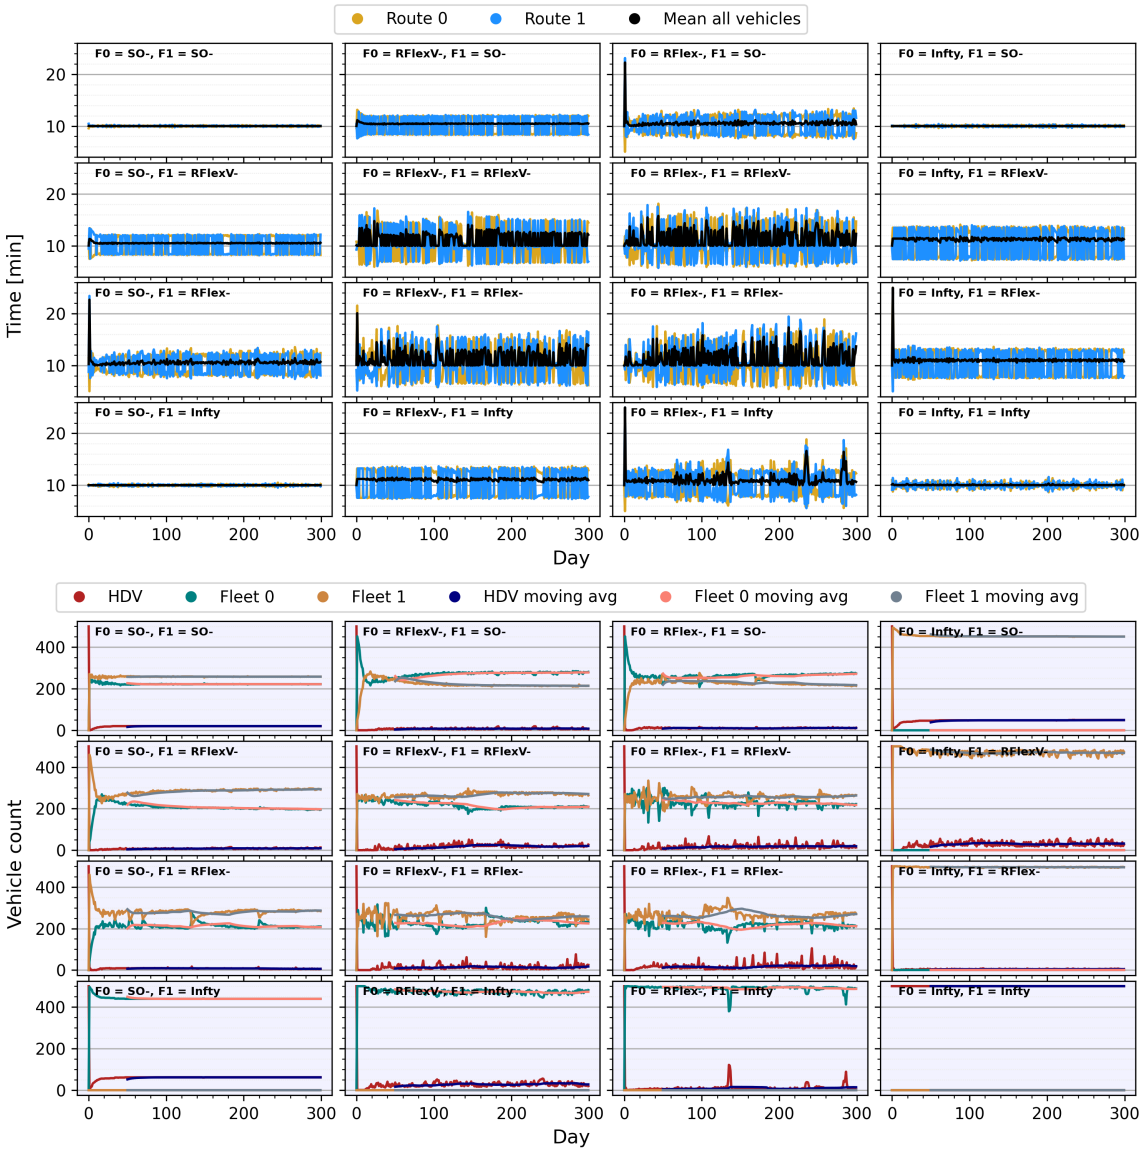


Figure 4: Competition of fleet operators in a benchmark scenario. Randomized algorithms RFlexV- and RFlex- perform better in one-fleet scenario (against F1 = Infty), while in two-fleet scenarios the results are mixed. The travel times when one or two fleets use randomized algorithms are highly variable.

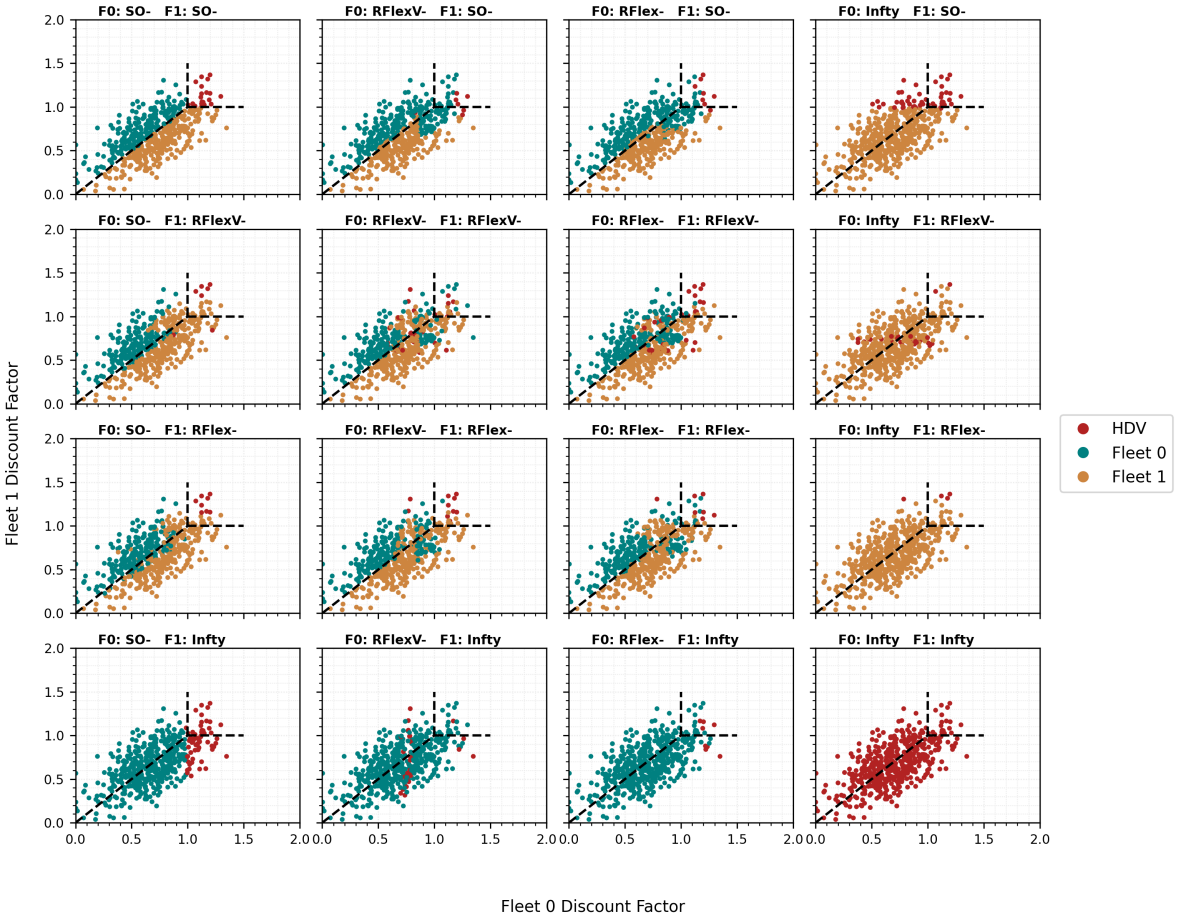


Figure 5: Usage of HDV or Fleet 0 or Fleet 1 on day 150 in simulations. Every dot in a panel corresponds to one vehicle and is plotted at coordinates corresponding to respective discount factors of fleets. Dashed lines separate the natural basins of belonging - as in the upper left panel. More efficient algorithms are able to attract and retain customers from outside of their basins. The HDV sample of discount factors is the same in every panel.

Experiment 3: SO component in objective

Experiment 2 shows that randomized algorithms are efficient in reaching higher market share, however for the price of increased and highly varying average travel times. Here, we demonstrate that a combination of share maximization with minimization of average travel time could be an objective which would be more socially acceptable (less oscillations), while preserving competition and discouriging randomization. To this end let

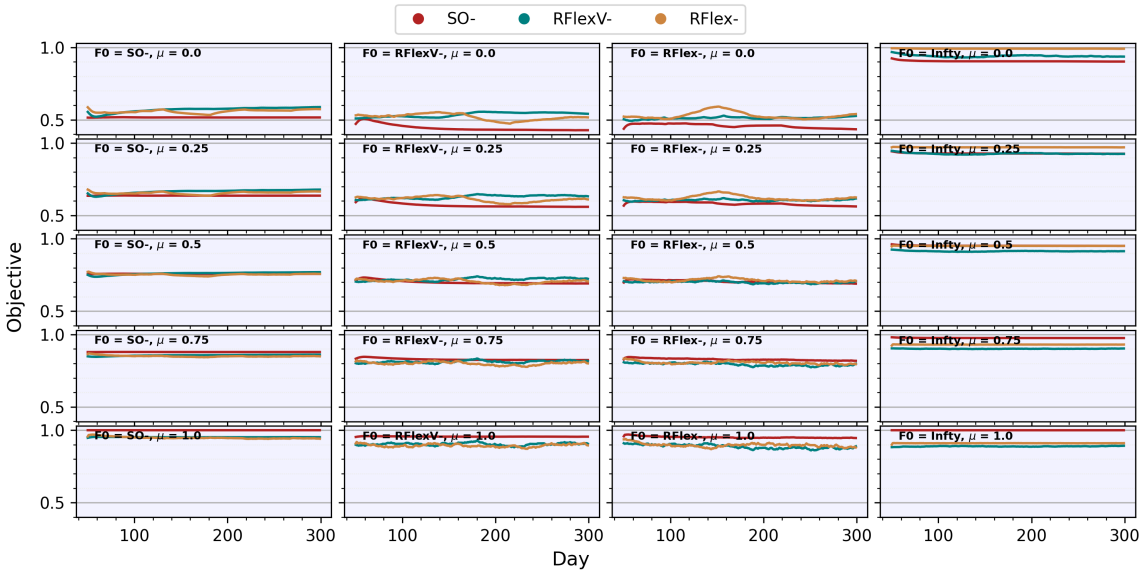


Figure 6: Smoothed-out objective (50-day moving average) for different values of μ and algorithm used by the fleet $F0$ against the fleet $F1$ is competing. Red/green/orange line correspond to $F1$ using SO-/RFlexV-/RFlex-, respectively. With growing μ the algorithm aiming at SO, ie. SO- becomes more and more competitive and eventually superior to randomized algorithms aiming solely at share maximization.

$$t^j_{avg} = \frac{1}{N_D} \sum_{r \in R} q^j_r t^j_r$$

where q^j_r is the flow (number) of vehicles on route r and day j. Let $\tau^j_{avg} := t^{SO}_{avg}/t^j_{avg}$ be the normalized average travel time and let $n^{f,j} = |I^{f,j}|/N_D$ be the share of fleet f on day j. Fig. 6 shows the objectives we propose (to be maximized):

$$Obj^\mu = (1-\mu)n^{f,j} + \mu\tau^j_{avg}$$

for different values of parameter μ. We note that $\mu = 0$ corresponds to market share maximization while $\mu = 1$ corresponds to minimizing average travel time. As already noticed before, for $\mu = 0$ the randomized algorithms $RFlex-$ and $RFlexV-$ are superior to SO oriented $SO-$. This changes with increasing μ and for $\mu \approx 0.5$ the plain algorithm $SO-$ becomes competitive enough and randomization is discouraged; indeed using randomization resulting in increased average travel times reduces the objective, i.e. payoff for the fleet operator.

4 Conclusions

In the paper we introduced a benchmark for comparing routing algorithms competing for customers in the scenario with diverse drivers. We provided baseline algorithms (SO) as well as more advanced heuristic randomized algorithms (RFlexV, RFlex). Running the benchmark scenario we found that:

- Dynamic routing algorithms using randomized strategies, which result in highly variable travel times seem to be more efficient than baselines when the sole objective is maximization of market share (and attitudes of drivers towards fleets are known).

- If one fleet uses a randomized algorithm, the other cannot stick to the baseline SO routing to stay competitive.
- The city could shape the market by setting the remuneration scheme based on combining the market share objective and deviation from System Optimum objective.
- There are indications that cities might be able to preserve competition while discouraging unwanted market-share optimizing behaviours such as randomized routing.

Future work will encompass:

- Including benchmark scenarios with more routes than 2 and more fleets than 2.
- Developing even more efficient routing algorithms which directly address competition against other fleet providers as well as take into account the objective set by city authorities.
- Developing ML-based routing algorithms in the competition case.
- Considering scenarios where fleet discount factors are not known to fleet operators and have to be inferred.
- Simulations in realistic city-scale simulators such as open-source SUMO, Lopez et al. (2018).

Acknowledgements

This work was financed by the European Union within the Horizon Europe Framework Programme (ERC Starting Grant COeXISTENCE no. 101075838). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

References

Collier, R. B., Dubal, V. B., & Carter, C. L. (2018). Disrupting regulation, regulating disruption: The politics of uber in the united states. *Perspectives on Politics*, *16*(4), 919–937.

Cramer, J., & Krueger, A. B. (2016). Disruptive change in the taxi business: The case of uber. *American Economic Review*, *106*(5), 177–182.

de Almeida Correia, G. H., Looff, E., Van Cranenburgh, S., Snelder, M., & Van Arem, B. (2019). On the impact of vehicle automation on the value of travel time while performing work and leisure activities in a car: Theoretical insights and results from a stated preference survey. *Transportation Research Part A: Policy and Practice*, *119*, 359–382.

GoldmanSachs. (2024). *Autonomous vehicle market forecast to grow ridesharing presence.* Retrieved from `https://www.goldmansachs.com/insights/articles/autonomous-vehicle-market-forecast-to-grow-ridesharing-presence` (Accessed: 2026-02-21)

Hoffmann, M., Bujak, M., Jamróz, G., & Kucharski, R. (2025). Wardropian cycles make traffic assignment both optimal and fair by eliminating price-of-anarchy with cyclical user equilibrium for compliant connected autonomous vehicles. *arXiv preprint arXiv:2507.19675*.

Jamróz, G., Kucharski, R., & Watling, D. (2025). Market share maximizing strategies of cav fleet operators may cause chaos in our cities. *arXiv preprint arXiv:2512.03524*.

Khan, L. M. (2016). Amazon's antitrust paradox. *Yale lJ*, *126*, 710.

Kominers, S. D., Teytelboym, A., & Crawford, V. P. (2017). An invitation to market design. *Oxford Review of Economic Policy*, *33*(4), 541–571.

Lopez, P. A., Behrisch, M., Bieker-Walz, L., Erdmann, J., Flötteröd, Y.-P., Hilbrich, R., ... Wiessner, E. (2018). Microscopic traffic simulation using sumo. In *2018 21st international conference on intelligent transportation systems (itsc)* (p. 2575-2582). doi: 10.1109/ITSC.2018.8569938

Roth, A. E. (2018). Marketplaces, markets, and market design. *American Economic Review*, *108*(7), 1609–1658.

A Appendix: SUMO-based simulations



Figure 7: Screenshot of SUMO-based simulation experiment. The two equivalent routes with stop signs exhibit BPR-like dependence of mean travel time on flow.

While the experiments in the main text clearly demonstrated the advanteges of randomized algorithms, they still used simplified flow-delay relations based on BPR functions. This experiment demonstrates the first step towards more realistic scenarios which is conducted as a microsimulation in SUMO, Lopez et al. (2018). To do this, we designed a system of two routes with properties similar to the one used in the main experiments. We assumed that:

- There are two equivalent routes with stop signs, see Fig. 7.
- There are 150 vehicles, departing within 600 seconds.
- Each vehicle (be it HDV or CAV) chooses/is assigned route, however it cannot choose the exact departure time.
- Vehicles arrive at start of routes with uniform probability within 10 min.
- The travel time depends on congestion on each route, and is stochastic: a driver may get stuck at STOP for long if unlucky or may pass quickly.
- A day of travel is an independent SUMO run yielding actual private travel times for every driver and public means on routes.
- HDV utilities are given by (2) with $\overline{t_r} = \sum_{\iota=j-j_{avg}}^{j-1} t_r^{SUMO}(q_r(\iota))$, where $q_r(\iota)$ is the number of vehicles that chose route r on day ι and t_r^{SUMO} is the mean dependence of travel time on flow from Fig. 8.
- CAV utilities are based on (3) with CAV fleet credibilities based on (4) with $t^j_{r_i^j}$ replaced by $t_i^{j,SUMO}$, where $t_i^{j,SUMO}$ is the private travel time of driver i in the SUMO simulation of day j.
- CAVs use the (known) dependece in Fig. 8 for their algorithms.

The experimental results shown in Fig. 9 demonstrate that:

- The randomized algorithms RFlexV- manages to achieve higher market share than baseline SO-.
- The results presented in this paper paper are likely to generalize to real-world scenarios.

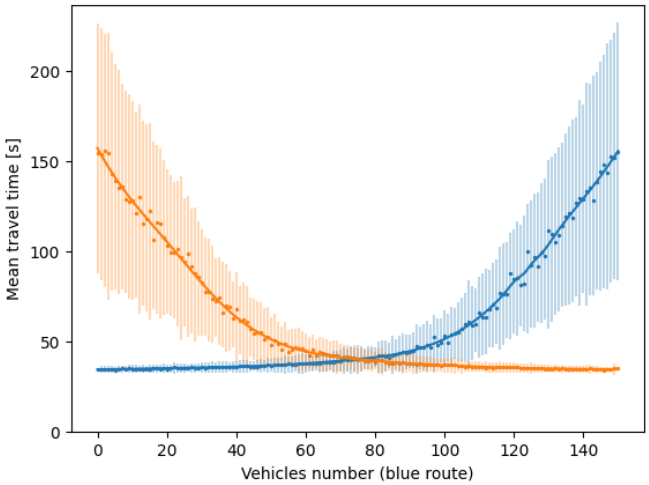


Figure 8: Dependence of average travel time on routes (blue - upper route, orange - lower route) on vehicle number in the SUMO-based system. The number of vehicles on orange route equals 150 minus the number of vehicles on blue route. Error bars correspond to standard deviation. Data from 10 indepent simulations.

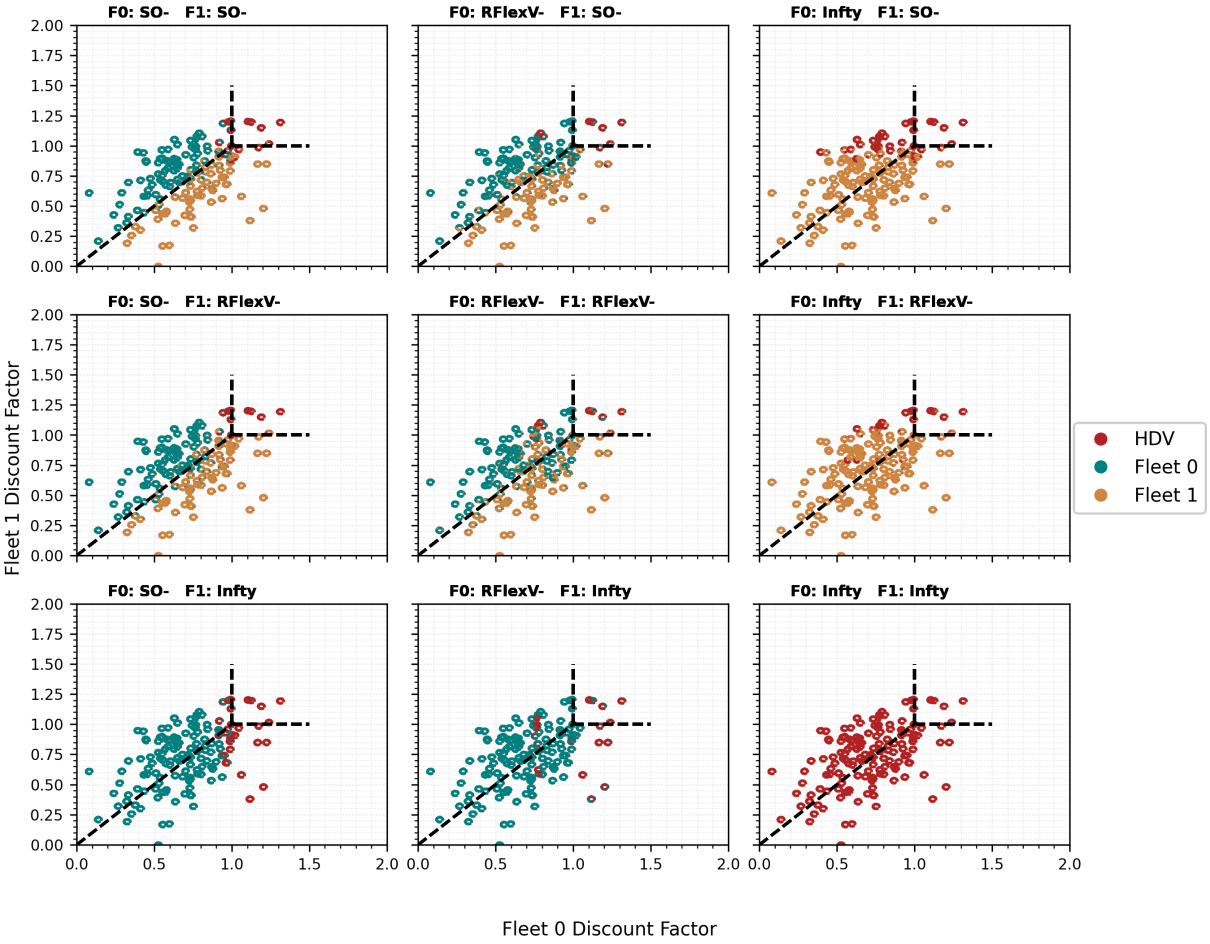


Figure 9: Modal split on days $1 - 200$ in the SUMO-based experiment. Each circle sector depicts the chosen modes of a driver (angle 0 - day 1, angle $3\pi/2$ - day 200). The advantage of randomised algorithms over SO- is clearly present in this setting as well.