Assessing expected ride-pooling performance with non-deterministic, heterogeneous travellers' behaviour.

Michal Bujak*[1, 2] and Rafal Kucharski[1]

[1]Department of Mathematics and Computer Science, Jagiellonian University, ul. Prof. Lojasiewicza 6, Krakow, 30-348, Poland
[2]Doctoral School of Exact and Natural Sciences, Jagiellonian University, Golebia 24, Krakow, 31-007, Poland

Short summary

Ride-pooling remains a promising emerging mode with a potential to contribute towards urban sustainability and emission reductions. However, recent studies revealed complexity and diversity among travellers' ride-pooling aptitudes. So far, ride-poling analyses assumed homogeneity and/or determinism of ride-pooling travellers. This, as we demonstrate, leads to a false assessment of ride-pooling system performance. We experiment with an actual NYC demand from 2016 and classify travellers into four groups of various ride-pooling behaviour (value of time and penalty for sharing), as reported in the recent SP study. We replicate their random behaviour to obtain meaningful distributions. Heterogeneity assumption proves to have a significant impact on the system. The performance indicators are shifted compared to the deterministic scenario. Albeit the high variability of travellers' preferences, system-wide results remain within reasonably narrow confidence intervals.
Keywords: ride-pooling, behavioural heterogeneity, shareability graph.

1 Introduction

Ride-pooling is a ride-hailing service which enables travellers to share a ride. Two or more trip requests are submitted to the platform and pooled into a single vehicle. Shared ride is typically longer (due to pick-up delay and detour) and yields a sharing discomfort, which has to be compensated by a lower ride fare.

Most ride-pooling algorithms (Alonso-Mora et al. (2017), Ke et al. (2021), Shah et al. (2020), Bilali et al. (2020), Wang et al. (2021)) operate on based on fixed time windows constraints. In reality, travellers are rational utility maximisers, who opt for the most attractive alternative. In this spirit, we proposed our previous ExMAS algorithm (Kucharski & Cats (2020)), which puts the traveller in the centre. Recent studies (Lavieri & Bhat (2019), Chavis & Gayah (2017), Alonso-González et al. (2020)) show that preferences towards pooling vary across population. Here, we introduce the behavioural heterogeneity (varying value of time, perceived discomfort of sharing) into the algorithm. While we are not aware of the individual traveller's preferences, we assume to know the distribution in the population. Applying results from Alonso-González et al. (2020) study on the actual taxi requests from NYC, we replicate the ride-pooling experiment 1000 times to obtain a meaningful estimation of the results.

We find that introduction of the heterogeneous preferences shifts performance indicators compared to the homogeneous scenario. While expected mileage reduction and profitability are worse than in the deterministic benchmark, traveller satisfaction with the service improves.

Our method allows the provider for more meaningful estimation of the system performance and its distribution. Results, apart from the expected values, reveal confidence intervals and tails. Estimates of performance (e.g. vehicle kilometres travelled, fares collected) and perceived attractiveness for travellers are now more meaningful and robust.

2 Methodology

To reach the objectives of this study, we extend the previously proposed off-line utility-based ride-pooling algorithm ExMAS. We introduce non-deterministic utility formulas to integrate them

with the previous deterministic approach and solve the ride-pooling problem for the heterogeneous population.
The original, traveller-oriented ExMAS is a utility-based algorithm. Traveller is assumed to share a ride only if it is more attractive than the alternatives (in our case, private ride). We measure attractiveness via utility, which has always a negative sign as it represents the perceived cost of a trip. The utility of a shared ride (for passenger i participating in ride r_k denoted as U^s_{i,r_k}) and of the non-shared ride (denoted U^{ns}_i) are as follows:

$$\begin{aligned} U^{ns}_i &= -\rho l_i - \beta_t t_i \\ U^s_{i,r_k} &= -(1-\lambda)\rho l_i - \beta_t \beta_s (\hat{t}_i + \beta_d(\hat{t}^p_i)) + \epsilon, \end{aligned} \tag{1}$$

where ρ stands for per-kilometre fare, while λ denotes discount for sharing a ride. Both are controlled by the operator. β^t, β^s β^d are the behavioural parameters: value of time, penalty for sharing and delay sensitivity, respectively. t_i and $\hat{t}_i$ stand for travel time of non-shared and shared ride, respectively, $\hat{t}^p_i$ is updated for each evaluated shared-ride candidate and is typically greater than t_i due to both pooling detour (to pick-up and drop of other travellers) and delay (to wait for others). ϵ is a random term. By conducting the hierarchical search, while preserving exhaustiveness, the algorithm prevents the search space explosion and yields an optimal solution. Details are described in Kucharski & Cats (2020).
In this study, we extend the above approach to account for the heterogeneous behavioural characteristics of individuals. Namely, any of behavioural parameters (β's in eq. 1) instead of being constants, can be now a random variable. We assume that preferences of an individual are unknown, while we know the distributions within the population. Moreover, to account for additional uncertainty, we introduce a panel noise ϵ_i (traveller-specific) and idiosyncratic noise $\epsilon_{i,r}$ (ride-specific for specific traveller), leading to the following formulas:

$$U^{ns}_i = \rho l_i + \mathring{\beta}_t t_i \tag{2}$$

$$U^s_{i,r_k} = (1-\lambda)\rho l_i + \mathring{\beta}_t \mathring{\beta}_s (\hat{t}_i + \mathring{\beta}_d(\hat{t}^p_i)) + \epsilon_i + \epsilon_{i,r}, \tag{3}$$

where all notions introduced in eq. 1 still apply, yet $\mathring{\beta}$s are now random variables. Consistently with a discrete choice theory, we assume that the probability that a traveller i finds a shared ride r_k attractive is expressed with:

$$P^s_{i,r_k} = \Pr(U^s_{i,r_k} > U^{ns}_i), \tag{4}$$

which, depending on the random variables distributions, may become e.g. a logit or a probit model. While in general any distribution of random variable can be applied in the method, in the experiment we use multiple classes $\mathcal{C}$ and assume the random variables to follow a multimodal normal distribution (unimodal within the class):

$$\mathring{\beta}_j = \sum_{C \in \mathcal{C}} p_C X(\bar{\beta}_{j,C}, \sigma_{j,C}) \text{ for } j \in \{c, t, s, d\}, \tag{5}$$

where p_C is the probability of belonging to the class C, $X(\bar{\beta}_{t,C}, \sigma_{j,C})$ is a random variable following normal distribution with mean of $\bar{\beta}_{t,C}$ and a standard deviation of $\sigma_{j,C}$ (coefficient- and class-specific).
We measure the performance of ride-pooling solution with the four following indicators. For the environment and city perspective we look at the vehicle-kilometres saved due to pooling $\mathcal{D}$. For travellers, we observe on one hand the pooling costs (relative detours $\mathcal{T}$) and on the other benefits (relative improvement in utility $\mathcal{U}$). For the platform, we look at the profitability of ride-pooling service $\mathcal{P}$. All values are calculated as relative to the value of the private ride scenario. Furthermore, indicators are calculated not only system-wide, but on the levels of a single traveller (utility, travel time extension), ride (mileage reduction, profitability).

3 Results and discussion

Experiment settings

We illustrate how the proposed method enhances assessment of ride-pooling services with the case-study of New York City. We reproduce the likely case when the service provider can predict the demand and its behavioural structure, yet the actual traits of individual travellers remain

Figure 1: Demand dataset for experiments: ca. 150 trip requests from Jan 2016 in Manhattan. Green dots are origins, orange destinations.

latent. To obtain reliable estimates, we fix the demand (requests) and replicate the experiment (assuming various ride-pooling behaviour of individual travellers) 1000 times. Results are available to reproduce on the public repository[1]

We experiment with the trips actually requested in the New York City on January 2016. We use the set 147 trips requested in the 30 minute batch (which is beyond the critical mass needed to induce pooling). Each trip request links origin with destination at a given time (as illustrated in Figure 1). We used the fare λ_p of 1.5 €/km (converted from \$, consistent with NYC Taxi & Limousine Commission (2022)), sharing discount λ is 30% (within the range suggested by Shaheen & Cohen (2019)). We assume $\beta_d = 1$ while $\mathring{\beta}_t$ and $\mathring{\beta}_s$ are derived from study by Alonso-González et al. (2020).

Value of time (VoT) found by Alonso-González et al. (2020) was explicitly applicable to our study. However, penalty for sharing (PfS) needed to be adjusted. ExMAS assumes a penalty for sharing irrespective of the number of co-travellers. Moreover, the penalty is scaled with trip length, not fixed. Hence, we took value for 3 additional passengers, which was found for one class. We scaled the results to obtain values for the remaining three classes (based on proportions in the fixed penalty for sharing with one additional passenger). According to formulas by Seltman (2012), we retrieved the variance of the variables (which was missing in the original study).

Eventually, we successfully obtained: class membership probabilities (p_C), mean values of value-of-time and penalty for sharing for four different classes ($\beta_{s,C}$, $\beta_{t,C}$) and class-respective variances ($\sigma_{s,C}$, $\sigma_{t,C}$) as reported presented in Table 1.

Table 1: Ride-pooling behavioural parameterization from Alonso-González et al. (2020). Mean values and standard deviations of value of time (VoT) and penalty for sharing (PfS) for four classes. *for class 4 we arbitrarily clipped the otherwise too wide st. dev. of VoT

parameter		class name	C1 "It's my ride"	C2 "Sharing is "saving"	C3 "Time is gold"	C4 "Cheap and half empty"
VoT	$\beta_{t,C}$	mean	16.98	14.02	26.25	7.78
		st.dev.	0.318	0.201	5.777	1*
PfS	$\beta_{s,C}$	mean	1.22	1.135	1.049	1.18
		st.dev.	0.082	0.071	0.06	0.076
Share	p_C		29%	28%	24%	19%

Expected ride-pooling performance with non-deterministic travelers

We dive into details how the behavioural heteroscedasticity impacts the ride-pooling performance. We assess the ride-pooling system's performance with the four selected indicators. Calculated first for consecutive replications (realisations of random variables) and then accumulated over all 1000 replications (one value corresponds to one realisation). Results are presented in Figure 2. Baseline for comparison is the deterministic ExMAS (with VoT and PfS being weighted average of means introduced in Table 1) which yields mileage reduced by 30%, 9.8% detour, utility increased by 4.5% and profitability of 1.097.

[1]Script to reproduce results is on the branch *probabilistic_topological* of original ExMAS (direct link: `https://github.com/RafalKucharskiPK/ExMAS/blob/probabilistic_topological/Utils/Probabilistic_ExMAS_wrapper.py`).

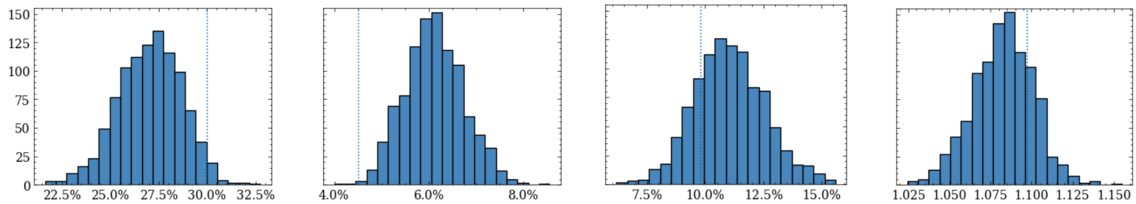

(a) Mileage reduct. $\mathcal{D}$ (b) Utility gain $\mathcal{U}$ (c) Trip detour $\mathcal{T}$ (d) Profitability $\mathcal{P}$

Figure 2: Distribution of ride-pooling performance resulting from 1000 replications. Dotted lines show the deterministic benchmark.

Regardless of behavioural variability, pooling always reduces the travel distance ($\mathcal{D}$). 90% of the observations range between 24.3% and 29.7% with the average of 27.1%, which is significantly worse than 30% reductions obtained in the deterministic setting (Figure 2a). While this is a significant reduction, which should be hugely appreciated from the city's and environmental's perspective, it also varies significantly, with observations ranging from 20% to 32.5%. The right tail seems to be fatter, with occasional outliers greater than 30% vehicle kilometres reductions.
The mean of average utility gain (which is user-subjective and depends on behavioural profile) is approximately 6.1%, which, in turn, is now significantly greater than in the deterministic benchmark. The 90% two-sided confidence interval spans between 5.1% and 7.2% utility gains.
A more physical measure of pooling performance, from the user's perspective, is the trip detour, which does not depend on traveller behaviour. Notably, the acceptable detour (and delay), unlike the fixed time windows used in other studies, depends on VoT and PfS and can vary significantly with the behaviour. The mean travel time of pooled ride is approximately 11% longer than for the solo-rides only system (Figure 2c) and 90% of the observations fit between 8.6% and 13.6%, which, again, is worse than in the deterministic benchmark (9.8%).

Class dependent ride-pooling performance

Considering the importance of the behavioural preferences of individuals on pooling performance, we analyse further the data with respect to individuals and their classes. First, we show the cumulative distributions of individual indicators observed in 1000 replications. Along with overall distributions, we present the profiles associated with four distinct classes. While the CDF profiles serve for illustrative purposes mainly, we enhance them with tables showing differences between classes in the four consecutive indicators as they reveal intriguing patterns among the introduced behavioural classes.
Here, we aggregate all the rides, irrespective of in which replication they appear. In the Fig. 3, we present the cumulative distribution functions of the utility gains and detours. For classes C1 and C3, the utility gains are below average and detours are above average. For the class C4 the opposite is true, while class C2 in on average with the mean performance.
In the Table 2 we report the resulting trip detour for the passengers in the four respective behavioural classes. Apart from mean and standard deviation, we report values of 75th, 90th percentile and 95th percentile. First, we present the values obtained for all the rides and then, for illustrative purposes, only for shared rides (excluding the travellers who were not matched with anyone and travelled alone). While on average, pooled rides are 8% longer than private rides, it varies from 5% for class C1 to 15% for class C4. Notably, 5% of travellers in C4 decided to pool while having at least 61% longer travel times. Those trends are further pronounced when we restrict the analysis to the rides which were successfully pooled (right columns).
Similarly to travel time analysis in Table 2, we present data of relative utility gain in ride-hailing system with ride-pooling (relative to system without pooling). Surprisingly, the class which has the longest detours (C4) has also the greatest utility gains (11% on average). For the class unwilling to pool the benefits are only 4% (C1). For 10% of all the travellers ride-pooling led to at least 25% increase in perceived utility.

Ride-pooling performance vs ride-pooling behaviour

In this section, we provide an insight into the relation between sampled behavioural parameters (VoT, PfS) and performance of the individual travellers ($\mathcal{U}_r$, $\mathcal{T}_r$). In Figure 4, we scatter the values of VoT and PfS (first and second row respectively) on x-axis against the performance indicators,

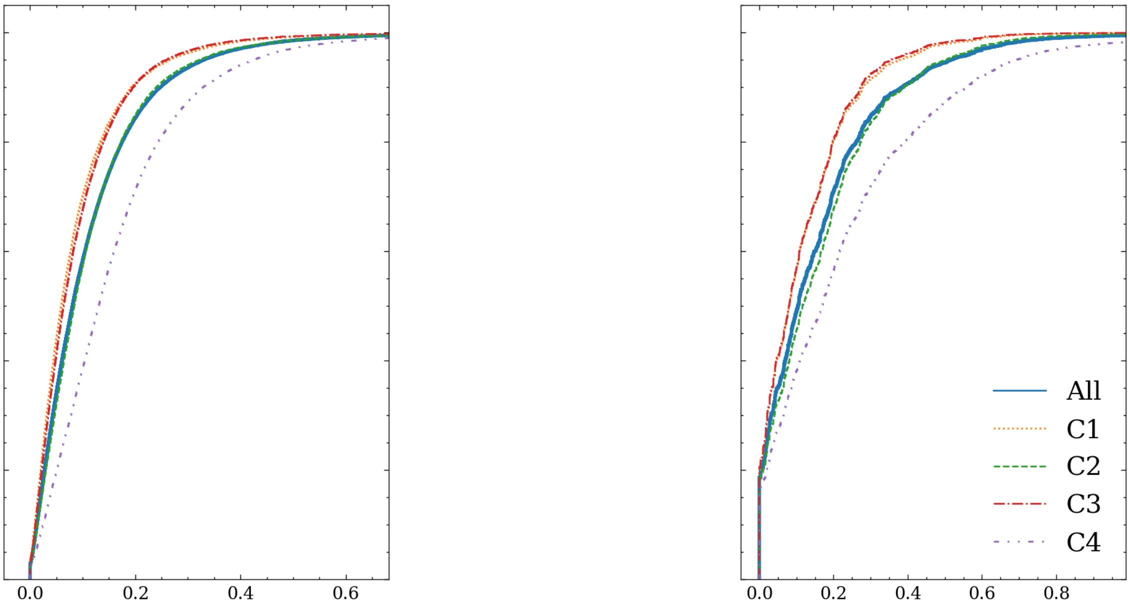

(a) Utility gain ($\mathcal{U}_r$)

(b) Detour ($\mathcal{T}_r$)

Figure 3: Cumulative distribution functions of utility gains and detours for the trips observed in 1000 replications, stratified into four classes of behaviour.

Table 2: Statistical properties of trip detour (relative) $\mathcal{T}_r$ for travellers of four respective behavioural classes. Mean, variance and significant percentiles are shown first for all the travellers (left) and then (right) only for those who were pooled.

	All rides					Shared rides				
	Means	St.dev.	75	90	95	Means	St.dev.	75	90	95
All	0.08	0.16	0.11	0.27	0.39	0.16	0.19	0.22	0.39	0.52
C1 “It’s my ride”	0.05	0.11	0.07	0.19	0.27	0.12	0.13	0.18	0.28	0.38
C2 “Sharing is saving”	0.09	0.15	0.14	0.28	0.41	0.16	0.17	0.23	0.39	0.50
C3 “Time is gold”	0.05	0.11	0.07	0.19	0.27	0.11	0.13	0.18	0.27	0.35
C4 “Cheap and half empty"	0.15	0.23	0.23	0.45	0.61	0.23	0.25	0.33	0.56	0.69

Table 3: Statistical properties of travellers gains (in relative increase of utility) $\mathcal{U}$ for travellers of four respective behavioural classes. Mean, variance and significant percentiles are shown first for all the travellers (left) and then (right) only for those who were pooling.

	All rides					Shared rides				
	Means	St.dev.	75	90	95	Means	St.dev.	75	90	95
All	0.06	0.10	0.09	0.18	0.25	0.11	0.12	0.15	0.25	0.33
C1 “It’s my ride”	0.04	0.08	0.05	0.12	0.18	0.09	0.10	0.12	0.19	0.26
C2 “Sharing is saving”	0.06	0.10	0.09	0.18	0.25	0.11	0.12	0.15	0.24	0.32
C3 “Time is gold”	0.04	0.08	0.06	0.14	0.19	0.09	0.10	0.12	0.20	0.26
C4 “Cheap and half empty"	0.11	0.14	0.16	0.28	0.37	0.16	0.15	0.22	0.33	0.42

i.e. $\mathcal{T}_r$ and $\mathcal{U}_r$ (first and second column) on y-axis. Each dot represents an individual traveller, coloured accordingly to her behavioural class.

We can observe a clear trend of ride-pooling detours decreasing with an increasing value of time (Fig. 4a), despite the longer detours for travellers with low value of time, their benefits of pooling (relative increase in utility) remains high (Fig. 4b), whereas for travellers with high value of time utility gains are significantly lower. The three classes (C1, C2 and C4) have relatively low variance and do not overlap, while for class C3 the variance of VoT is very big. Mind that in our experiment each traveller was assigned to different classes across replications, nonetheless the class membership strongly correlate with the resulting ride-pooling performance, despite the fixed spatiotemporal trip characteristics.

Those clear trends become blurred when we plot against penalty-for-sharing (PfS in Fig. 4c and 4d) where the Gaussian shape is observed and class memberships are indistinguishable. Both highest

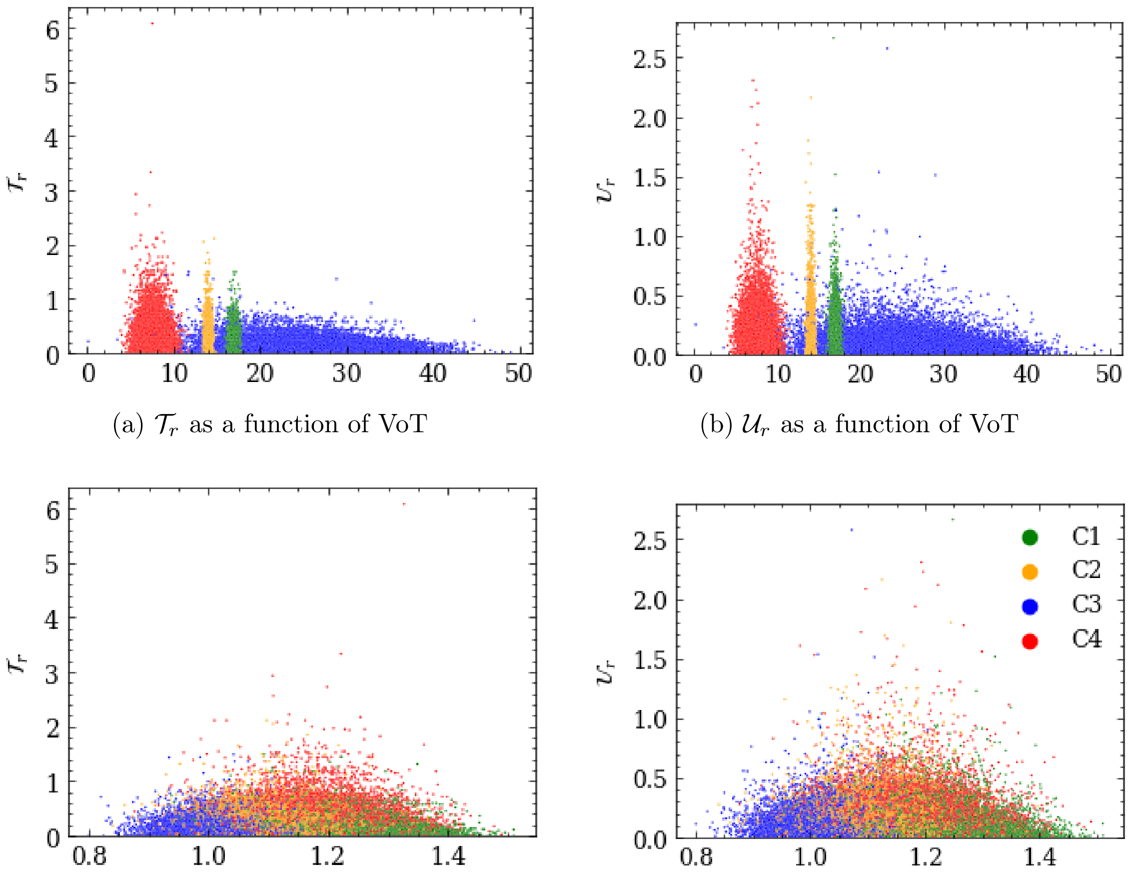

(a) $\mathcal{T}_r$ as a function of VoT

(b) $\mathcal{U}_r$ as a function of VoT

(c) $\mathcal{T}_r$ as a function of PfS

(d) $\mathcal{U}_r$ as a function of PfS

Figure 4: Scatter plot of travellers' behavioural parameters (x-axis) against the resulting ride-pooling performance. Each dot denotes a single traveller, coloured accordingly to the class assigned to her in the respective replication. y-axis are scaled to represent relative change, e.g. 1 corresponds to increase of 100%.

detours and utility gains are obtained for intermediate values of PfS, while extreme cases have the lowest benefits. The travellers with a low penalty for sharing (even below 1 for a some of class C3 travellers) seem to be often exploited by the system and their utilities are sacrificed to match preferences for travellers with lower preferences for sharing. Yet for those with high penalty the benefits are also limited, since they can hardly be pooled with others. The pooling seems to be best performing in the penalties ranging from 1.1 to 1.2 where it can reach the greatest expected benefits (Fig. 4d).

Scaling with the demand levels

The notion of critical mass is crucial to the ride-pooling and number of studies reported how the ride-pooling grows non-linearly with the demand size. Yet, up to now, those findings were reported in the deterministic setting only. Here, apart from reporting how system performance improves with growing demand, we provide insights on how the performance variability changes.
Demand set on which we run the experimental was selected to be slightly beyond the critical mass needed to induce the effective pooling, we extend it here from two sides: sub-critical (99 requests) and super-critical (198 requests). We report the previously introduced performance measures for three levels of demand in Figure 5. Unsurprisingly, we observe the critical mass effect as the performance significantly increases when the demand grows from 100 to 150 requests and then somehow stabilises when it reaches 200 requests (Fig. 3). Yet more importantly, we observe the trends with variabilities, which are alternative: either narrowing with growing demand (utility in Fig. 5b), widening (profitability in Fig. 5b), or remaining roughly constant (travel times and mileage savings in Fig. 5a and Fig. 5c).

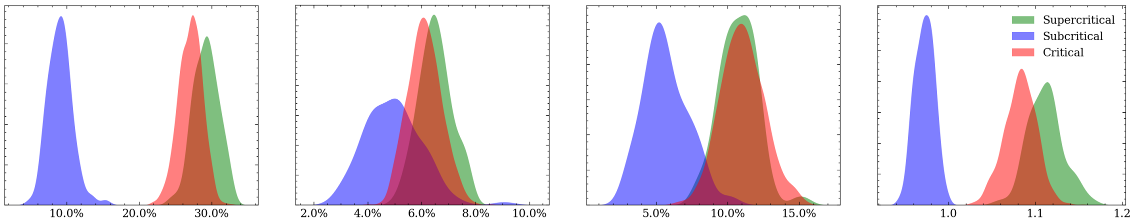

(a) Mileage save ($\mathcal{D}$) (b) Utility gain ($\mathcal{U}$) (c) Trip detour ($\mathcal{T}$) (d) Profitability ($\mathcal{P}$)

Figure 5: Distributions of four performance indicators with the three levels of demand: subcritical, critical and supercritical.

4 Conclusions

Most ride-pooling algorithms rely on the fixed constraints (usually time windows). To more objectively assess the attractiveness of a shared ride from the traveller's perspective, we previously proposed the utility-based ExMAS algorithm. In this study, we take the matter further. We no longer consider travellers homogeneous and deterministic. Using recent empirical findings, we assume that while we still do now know preferences of the individuals, we know the general distributions.
To properly include this in our calculations, we introduced random variables representing the value of time and perceived sharing discomfort. Such reformulation allowed to reproduce the four distinct classes of travellers' behaviour towards ride-pooling, as revealed in Alonso-González et al. (2020). Conducting sufficient number of replications, we obtain reliable estimations of the ride-pooling problem solution, which is no longer deterministic. Apart from mean values, we can now estimate the lower and upper bounds of the performance indicators, which substantially differ from the deterministic benchmark.
The additional probabilistic layer introduces a variability at the level of individuals and at the system performance. We find that despite the high variability of individuals, pooling performance remains stable and within reasonably narrow confidence intervals. Notably, probabilistic results are shifted compared to the deterministic benchmark. While the primary objective of minimising mileage is better met in the deterministic scenario, we observe a much higher satisfaction with the service in the heterogeneous setting.
Our study also provides an insight into how ride-pooling performs for travellers of certain preferences. We analyse the impact of the value of time and perceived sharing discomfort on trip detour and satisfaction with the service. We find that people with low value of time can be considered both the most flexible and the most beneficial travellers in the pooling system. However, those with intermediate penalty for sharing not only benefit more than those with high penalties (who does not want to share in general), yet also more than those with low penalties (willing to share with anyone and often exploited by the system).
Similarly to the deterministic case, we observe the critical mass effect and pooling becomes effective only when the demand levels reach the so-called critical mass. Now we enrich this notion with findings on variability, which may either decrease (in terms of utility gains) or increase (in terms of profitability for the provider) with growing demand levels.
The proposed method is general and can be easily applied to new cases, both for general demand patterns and different behavioural models. Also, the specific experimental setting used in this study may be reformulated, e.g. when we know individuals' class membership, the demand is not predicted properly or when the behaviour is assumed fixed, but demand is varying (like in Kucharski et al. (2021)). In the future studies, those additional dimensions of variability may be included for even richer assessments. Finally, the proposed method may be valuable in the pandemic-analyses, when virus-averse behaviour drives the pooling behaviour of individuals.

Acknowledgements

This research was funded by National Science Centre in Poland program OPUS 19 (Grant Number 2020/37/B/HS4/01847)

References

Alonso-González, M., Cats, O., & van Oort, N. e. a. (2020). What are the determinants of the willingness to share rides in pooled on-demand services? *Transportation*, 1733–1765. doi: https://doi.org/10.1007/s11116-020-10110-2

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

Bilali, A., Engelhardt, R., Dandl, F., Fastenrath, U., & Bogenberger, K. (2020). Analytical and agent-based model to evaluate ride-pooling impact factors. *Transportation Research Record*, *2674*(6), 1–12.

Chavis, C., & Gayah, V. V. (2017). Development of a mode choice model for general purpose flexible-route transit systems. *Transportation Research Record*, *2650*(1), 133–141.

Ke, J., Zheng, Z., Yang, H., & Ye, J. (2021). Data-driven analysis on matching probability, routing distance and detour distance in ride-pooling services. *Transportation Research Part C: Emerging Technologies*, *124*, 102922.

Kucharski, R., & Cats, O. (2020). Exact matching of attractive shared rides (exmas) for system-wide strategic evaluations. *Transportation Research Part B: Methodological*, *139*, 285–310.

Kucharski, R., Cats, O., & Sienkiewicz, J. (2021). Modelling virus spreading in ride-pooling networks. *Scientific Reports*, *11*, 1–11.

Lavieri, P. S., & Bhat, C. R. (2019). Modeling individuals' willingness to share trips with strangers in an autonomous vehicle future. *Transportation research part A: policy and practice*, *124*, 242–261.

NYC Taxi & Limousine Commission, u. (2022). *Taxi fare.* Retrieved from `https://www1.nyc.gov/site/tlc/passengers/taxi-fare.page`

Seltman, H. (2012). *Approximations for mean and variance of a ratio. unpublished note.*

Shah, S., Lowalekar, M., & Varakantham, P. (2020). Neural approximate dynamic programming for on-demand ride-pooling. In *Proceedings of the aaai conference on artificial intelligence* (Vol. 34, pp. 507–515).

Shaheen, S., & Cohen, A. (2019). Shared ride services in north america: definitions, impacts, and the future of pooling. *Transport reviews*, *39*(4), 427–442.

Wang, J., Wang, X., Yang, S., Yang, H., Zhang, X., & Gao, Z. (2021). Predicting the matching probability and the expected ride/shared distance for each dynamic ridepooling order: A mathematical modeling approach. *Transportation Research Part B: Methodological*, *154*, 125–146.