# The Vehicle for Hire Problem: A Generalized Kolkata Paise Restaurant Problem

*Complete Research Paper*

**Layla Martin**
Technical University of Munich
layla.martin@tum.de

**Paul Karaenke**
Technical University of Munich
karaenke@in.tum.de

## Introduction

Increasing traffic is a substantial environmental issue in metropolitan areas. New technology has the potential to alleviate this problem. For example, self-driving cars might ultimately lead to a higher utilization of cars and less congestion due to the fact that a central coordinator can position cars optimally based on historical demand patterns. In contrast, nowadays commercial drivers in vehicle for hire markets (e.g., taxis, Uber, and Lyft) self-select their location and this could be the source of inefficiencies; i.e., drivers decide individually where to look for customers. However, the average idle time of taxis is about 25–50% in most cities where data is available (Cramer and Krueger 2016; Linne+Krause Marketing-Forschung 2016). Though excess capacity can partially explain these numbers, utilization could be increased, if drivers were distributed across the city more efficiently. In contrast to underutilization, passengers have to wait for more than 20 minutes in approximately every third case in other cities (Rayle et al. 2014), suggesting that the drivers are not at the locations where needed. To address these inefficiencies in vehicle for hire markets, coordinators could instruct drivers where to wait for customers. In current business models, however, this is not possible, since drivers are not employees of the coordinators. Hence, they try to maximize their individual profits by deciding individually where to look for costumers without considering the social welfare or utilization of the market. In practice, there are approaches like ‘surge pricing’ to respond to expected peaks in demand, though literature on the efficiency of different driver strategies is limited (Chen and Sheldon 2015; Hall et al. 2015; Rogers 2015).

To analyze the fundamental underlying problem, we propose a repeated non-cooperative game model to investigate different strategies in the coordination problem among drivers, which we denote as the vehicle for hire problem (VFHP). We define the VFHP as a generalization of the Kolkata Paise restaurant problem (KPRP) (Chakrabarti et al. 2009) where agents repeatedly compete for a set of resources. We extend the KPRP to better reflect the specifics of vehicle for hire markets; addressing both individual agent preferences and different resource capacities. Besides our model, the contributions of this research are different mixed strategies for this model and an analysis of their impact on car utilization and driver utilities. We assume drivers to share a common ranking of locations by attractiveness. The results of our analysis yield that the more this ranking determines the selection of locations, the lower is the overall utilization. However, if we assume a modest level of randomization in how drivers select their location, we achieve high levels of efficiency of around 80%. These insights constitute building blocks for characterizing favorable agent behavior in order to design mechanisms to distribute drivers efficiently.

The remainder of this paper is organized as follows. We first discuss related work, followed by the presentation of our model. Then we describe the different strategies and their impact on car utilization and driver utilities. We discuss the implications and limitations of our research, before concluding the paper.

## Literature Review

To our knowledge, no other paper builds a game model for vehicle for hire markets. Relevant research is conducted in two fields. First, there is literature in optimization and operations research in the field of vehicle for hire markets. Second, we give an overview of relevant game models in other application areas.

Yet there is only limited research work available on optimal distribution of drivers in vehicle for hire markets. Several studies focus on assigning drivers an optimal district where they await passengers (Lee et al. 2004; Seow et al. 2010); though in most business models, drivers decide independently. Yang et al. (2005) study a model with varying demand and supply. Taxi drivers individually decide when to enter the market and when to leave it, resulting in a market equilibrium. This study does not optimize for utility, but only for utilization. Kim et al. (2011) propose an agent-based model incorporating real-world passenger travel pattern to predict the highest possible utility. Their model also incorporates districts ("areas") and varying utility functions over time, but tests for different criteria: Whilst we analyze utilization and utility for different strategies in a large environment, Kim et al. (2011) study a setting with five nodes and retrieves utilization and passenger wait time for varying fleet sizes. Wong et al. (2015) focus on reduced vacant mileage for taxis as the primary optimization criterion. The authors use a two-step approach in which taxis can only divert to adjacent zones rather than all others. Trigo et al. (2006) apply multiagent markov decision processes to model drivers transporting passengers. This paper's motivation is similar to ours, though instead of stochastic strategies, it applies a two-layered learning process.

Furthermore, there are several papers in the field of operations research which focus on the influence of regulation (taxi medallions, fixed rates) on the market (Cairns and Liston-Heyes 1996; Arnott 1996). In our game model we assume that the market is in a socially efficient state, as there are sufficient agents to carry every customer and sufficient customers such that every agent can carry a customer.

The VFHP is a type of congestion game, a model for games in which agents choose different alternatives to succeed first described by (Rosenthal 1973). Congestion games can be identified by their potential function and thus their pure-strategy Nash equilibria (Monderer and Shapley 1996). Yet, such a Nash equilibrium is usually inefficient, as Correa et al. (2005) prove. Other congestion game models are the KPRP (Chakrabarti et al. 2009), the El Farol bar problem (Arthur 1994), the crowding game (Milchtaich 2000), and the minority game (Challet and Zhang 1998).

## The Model

In our model, we relax two main assumptions of the KPRP (Chakrabarti et al. 2009): (i) Agents do not receive identical utility from a given resource and (ii) customers are clustered in districts. In our model, drivers constitute agents and customers, located in districts, constitute resources. Agents drive to districts to carry customers which await a ride in this district. Agents 'use' (i.e., service) resources up to the resources' capacity limit. We assume agents can divert to other customers inside the district they drove to. Hence, our model consists of three different groups of participants and resources: Agents $I$ ($i \in I = 1 \ldots N$), customers $J$ ($j \in J = 1 \ldots N$), and districts $K$ ($k \in K = 1 \ldots D$). On average $\varphi$ customers are located in one district, thus $N = \varphi D$. We assume that customers randomly 'choose' a district they are located in, thus, the number of customers located in one district is binomially distributed around $\varphi$; other assignments and distributions (such as geometric) make sense as well, but do not affect the fundamental results. $l_{jk}$ denotes if customer $j$ is located in district $k$ ($l_{jk} = 1$) or not ($l_{jk} = 0$). Every customer is located in exactly one district ($\sum_{k \in K} l_{jk} = 1$) and the number of customers that are located in a district $k$ is its capacity $r_k = \sum_{j \in J} l_{jk}$. The utility of a district (and thus the probability to be selected) depends on the customers located in them: If agents decide purely randomly, the expected number of agents driving there equals the number of customers. If (myopic) agents decide strictly depending on the maximum possible utility, they hope for servicing the highest utility customer in a district, therefore, the utility of a district equals the maximum customer utility. Agents thus target customers they drive to according to their strategy (e.g., randomly). If agent $i$ targets customer $j$, we denote this by $d_{ij} = 1$ ($d_{ij} = 0$ otherwise). Every agent targets exactly one customer ($\sum_{j \in J} d_{ij} = 1$); the number of agents driving to customer $j$ is denoted as occupancy $o_j = \sum_{i \in I} d_{ij}$. We assume that agents can divert to other customers that belong to the same district at no cost; $\breve{d}_{ik}$ denotes that agent $i$ drives to district $k$ ($\breve{d}_{ik} = 1$ if $i$ drives to $k$ and $\breve{d}_{ik} = 0$ otherwise). An agent implicitly drives to the district $k$ that the targeted customer $j$ is located in ($\breve{d}_{ik} = \sum_{j \in J} d_{ij} \cdot l_{jk}$). The occupancy $o_k$ of district $k$ is the number of agents $i$ driving to $k$ ($o_k = \sum_{i \in I} \breve{d}_{ik}$). We use $x_{ij} = 1$ to denote that agent $i$ services customer $j$ ($x_{ij} = 0$ otherwise). An agent $i$ can only carry a customer $j$ if $i$ drives to the district $k$ that $j$ is located in ($x_{ij} \leq \sum_{k \in K} \breve{d}_{ik} \cdot l_{jk}$). Agents can carry at most one customer ($\sum_{j \in J} x_{ij} \leq 1$) and a customer can be carried by at most one agent ($\sum_{i \in I} x_{ij} \leq 1$). In every district, agents carry as many customers as possible; i.e., agents do not refuse to carry customers remaining in this district. Thus, the number of customers carried per district is either capacity $r_k$ or occupancy $o_k$ ($\sum_{i \in I, j \in J} x_{ij} \cdot l_{jk} = \min(r_k, o_k)$). The agents' utility from servicing customers is calculated from the valuation (revenue) $v_j$ and transportation cost $c_j$, which are assumed to be equal for all agents but differ for customers $j$ located in the same district, as well as the individual agents' cost dependent on the distances to customers $c_{ij}$. To investigate the effects of both type of costs, we set $c_j, c_{ij} \in \mathscr{U}(0,1)$ and model the total cost as a weighted average of the individual and shared cost components as $\tilde{c}_{ij} = \alpha \cdot c_{ij} + (1-\alpha) \cdot c_j, \alpha \in [0,1]$. The distance (and thus the individual component) is assumed to be equal for all customers located in a district ($c_{ij} = \sum_{k \in K} \check{c}_{ik} \cdot l_{jk}$) and only occurs if the customer is transported due to the agent's myopia. The agent's utility function can then be defined as follows.

$$u_i(x) = \sum_{j \in J} \left( \left( v_j - \tilde{c}_{ij} \right) \cdot x_{ij} \right) \tag{1}$$

For our analysis, we use $v_j = 1$ to allow transportation of all customers by any agent with non-negative utility. If no customer yields a higher utility for agent $i$ than customer $j$, we say that $i$ prefers $j$ ($\forall j' \in J : u(i,j) \geq u(i,j')$). We denote the average agent utility by $\bar{u} = \frac{1}{N} \cdot \sum_{i \in I} u_i(x)$. The agent utilization $f_i$ defines whether an agent $i$ carries any customer ($f_i = \sum_{j \in J} x_{ij}$). The agent utilization fraction, i.e., the average utilization, is denoted by $\bar{f} = \frac{1}{N} \cdot \sum_{i \in I} f_i$. The utilization fraction also yields the fraction of customers carried to another location.

## The Strategies

We consider seven strategies: No Learning (NL), Rank Dependent Choice (RD), Limited Learning (LL), One Period Repetition (OPR), Crowd Avoiding (CA), Stochastic Crowd Avoiding (SCA), and Stochastic Rank Dependent Choice (SRD). RD, LL, OPR, and SRD incorporate the resources' utility in the agents' choices and are therefore utility-based. LL, OPR, CA, and SCA require knowledge about previous iterations and are therefore history-based.

Using the NL strategy, agents randomly choose a resource in every iteration, regardless of history (hence the term "No Learning") or resource utility. This strategy was investigated for the KPRP by (Chakrabarti et al. 2009).

With RD, agents always select one of the resources where they (myopically) expect highest utility. We introduce this strategy, as it mimics simple behavior if limited information is available: If agents do not know about the preferences or behavior of other agents, but assume that only a few agents share the same preference, the most simple approach is to always head for the preferred resource.

Agents using the LL strategy follow a two-step approach: (1) If an agent carried a customer at time $t$, it drives to the highest utility resource at time $t+1$. (2) If an agent did not carry a customer at time $t$, it will randomly choose any other resource at $t+1$. The LL strategy was presented by (Chakrabarti et al. 2009) (named Limited Learning 1). The LL strategy uses information about the preferences of an agent and its own behavior during the previous iteration.

The OPR strategy requires agents to follow a three-step approach: (1) If an agent carried customer $j$ at time $t$ but not at $t-1$, it returns to this resource at $t+1$ (*return*). (2) If an agent served the same resource $j$ at time $t-1$ and $t$, it competes for the highest utility customer at $t+1$ (*improve*). (3) If an agent did not carry any customer at time $t$, it randomly chooses any resource which was vacant at time $t$ in the next iteration (*random*). OPR was also introduced in (Chakrabarti et al. 2009). With this strategy agents require information about their own preferences and their own behavior during the last two iterations.

With the CA strategy agents only drive to resources which were vacant or had remaining capacity at time $t-1$. This strategy originates from a paper by Ghosh et al. (2013). Agents incorporating this strategy need to know which customers $j$ were not carried the previous iteration ($\sum_{i \in I} x_{ij} = 0$). Agents then only select these customers (but can carry any customer in the district $j$ is located in).

Agents using the SCA strategy (Ghosh et al. 2013) stochastically decide whether to return to the same resource or select another one. If a selected resource $j$ does not exceed its capacity at time $t$, agents return there at $t+1$. If the capacity of $j$ is exceeded, agents stochastically either return or drive to any other (randomly chosen) resource at time $t+1$ such that the expected number of agents selecting resource $j$ is equal to its capacity. Agents return with probability $\frac{r_j}{o_j}$ and randomly choose another resource with probability $1-\frac{r_j}{o_j}$.

We present four SRD strategies that combine the RD and SCA strategies: Agents select their preferred resource $k$ if its capacity is not exceeded, that is $r_k \geq \breve{N}_k$. Otherwise, they drive to $k$ with probability $\frac{r_k}{\breve{N}_k}$ and select another resource with probability $1 - \frac{r_k}{\breve{N}_k}$. Thus, the expected number of agents selecting their preferred resource $k$ is $min(\{r_k, \breve{N}_k\})$.

The alternative resource an agent chooses is one of the following: Any customer which is noone's first choice (SRD1); any other customer (SRD2); the agent's second choice (SRD3); or the best customer which is noone's first choice (SRD4). SRD3 and SRD4 are extensions of SRD2 and SRD1 increasing the expected average utility of successful agents, i.e., agents carrying a customer. All SRD strategies require information about the first preferences of all other agents. In addition to the number of agents preferring the same resource, the SRD1 strategy also requires information about the number of agents preferring all other resources. The SRD2 strategy requires less information than the SRD1 strategy; it only requires the number of agents preferring the same resource. It is thus beneficial if the information about other resources cannot be determined easily. The SRD3 strategy requires the same information, but the utility of successful agents is higher, as all successful agents receive a high utility (highest or second highest utility). If the agent utilities are not stochastically independent, there can be a high number of resources that are not selected, neither as first nor as second preference. Depending on $\alpha$, the second preference of an agent can likely be the first preference of another agent. In SRD4, the second preference is only chosen, if no other agent prefers this resource. Thus, the set of first choice resources and the set of alternate choice resources do not intersect, making it impossible that alternate choice agents carry a customer who is preferred by another agent. Yet, SRD4 requires more information about the preferences of other agents than SRD3. Thus, rationale for analyzing the four strategies is their difference in information required.

## Numerical Results

In the following we present our analytical results and our simulations. In (Martin 2017) we calculate the expected average utilization and utility for all given strategies, for this paper we exemplarily present the utilization of the NL and RD strategy.

Using the NL strategy, all agents drive to a randomly selected customer. Thus, individual utility levels are irrelevant. The utilization fraction depends on (1) the capacity $r_k$ of district $k$ (associated with probability $P(r_k)$) and (2) the occupancy $o_k$ of district $k$ (associated with probability $P(o_k|r_k)$). $P(r_k) = \frac{\varphi^{r_k}}{r_k!} \cdot e^{-\varphi}$ is the probability that $r_k$ customers are randomly assigned to the same district given an average of $\varphi$ customers per district. $P(o_k|r_k) = \frac{r_k^{o_k}}{o_k!} \cdot e^{-r_k}$ is the probability of $o_k$ agents randomly driving to district $k$ containing $r_k$ customers (thus, $r_k$ is the expected number of agents driving to $o_k$). If less agents drive to a district $k$ than customers await a ride ($r_k > o_k$), $r_k - o_k$ customers will obviously not be carried. If at least $r_k$ agents drive to district $k$ ($r_k \leq o_k$), all $r_k$ customers will be transported. Thus, the expected remaining capacity is the weighted average of $r_k - o_k$. To calculate the utilization we subtract the expected remaining capacity from the capacity $r_k$ and then calculate the weighted average for all possible capacities. Thus, the expected utilization fraction for the RD strategy $\mathbf{E}\left(\bar{f}_{NL}\right)$ can be calculated as follows.

$$\mathbf{E}\left(\bar{f}_{NL}\right) = \frac{1}{\varphi} \cdot \sum_{r_k=1}^{N} P(r_k) \cdot \left( r_k - \underbrace{\sum_{o_k=0}^{r_k-1} P(o_k|r_k) \cdot (r_k - o_k)}_{\text{expected remaining capacity}} \right) \tag{2}$$

All strategies that include the rank in the decision process depend on the probability $P\left(\check{N}_k\right)$ that a district $k$ is preferred by $\check{N}_k$ agents which depends on the highest utility of a customer $j$ located in $k$. It is calculated as a Poisson distribution around $\mathbf{E}(\check{N}_k)$, the expected number of agents preferring $k$. If all districts have the same $\mathbf{E}(\check{N}_k)$, the utilization of the RD strategy equals the strategy of the NL strategy, as first preferences are randomly drawn. In general we calculate the utilization as follows: For every customer $j$ we calculate the probability $P\left(c_j \leq c_{j'} \forall j' \in J_k | r_k\right)$ that this customer has the lowest shared cost component $c_j$ for the set of all customers $J_k$ located in district $k$ given capacity $r_k$. The utilization yielded in this scenario is then calculated using the average number of customers not carried in this district $k$.

$$\mathbf{E}\left(\bar{f}_{RD}\right) = \frac{1}{N} \cdot \sum_{j=1}^{N} \sum_{r_k=1}^{j} P\left(c_j \leq c_{j'} \forall j' \in J_k | r_k\right) \cdot \left( r_k - \sum_{\check{N}_k=0}^{r_k-1} P\left(\check{N}_k\right) \cdot \left(r_k - \check{N}_k\right) \right) \tag{3}$$

In this paper we restrict ourselves to an analytical evaluation of the two aforementioned strategies. NL and RD are the baseline comparison strategies, their analytical solution suggests the correctness of our simulations. Further, these two strategies mimic basic behavior which can be easily represented in succinct closed form. A more extensive study of the other strategies can be found in (Martin 2017). Table 1 shows our assumptions for numerical experiments and simulations: We assume that there are 1000 agents and customers, an average of 5 customers per district (resulting in 200 districts) and equal weights of 0.5 for individual and shared cost components. We replicate each experimental setup $10^6$ times.

| Parameter | Value |
|---|---|
| $N$ | 1000 |
| $\varphi$ | 5 |
| $\alpha$ | 0.5 |
| replications | $10^6$ |

Table 1: Parameters

In table 2 we present the simulation results for utilization and utility for all strategies we introduced in the previous section. The utilization fraction is the percentage of agents carrying a customer; the utility is the average agent utility of all agents. Obviously, the best strategy with respect to utilization and utility is SCA, meaning that considering only own previous choices is more promising than regarding for utility as well. The second-best strategy with respect to utilization is NL, suggesting that it is rather difficult to reach a higher utilization fraction than with very simple behavior. The second-best strategy regarding utility is OPR, a strategy that combines the waiting of SCA with improving the utility. SRD1 and SRD2 get close to the utilization fraction of NL (and yield a higher utility), as diverting agents behave as if they chose randomly. Also, OPR yields a higher utilization than LL, as agents wait for one iteration in a district where they can be (almost) certain to carry a customer prior to continuing to their preferred resource. Comparing each strategy introduced by (Chakrabarti et al. 2009; Ghosh et al. 2013) for the KPRP (that is, NL, LL, OPR, CA, and SCA) and the VFHP, we observe that these strategies result in a higher utilization and utility in the VFHP, as districts and an individual utility component allow more agents to carry a customer.

We further simulated the effect of varying $\alpha$ and $\varphi$ in the NL (dashed black line) and RD (red line) strategy. The results are presented in figures 1 and 2. We firstly observe that varying $\alpha$ does not have an effect on the NL strategy, as the agent's choice is independent from the utility in this strategy. For RD, the utilization almost reaches the utilization of the NL strategy, whilst it is considerably lower for lower $\alpha$. Obviously, a higher impact of the individual cost component $c_{ij}$ results in a higher utilization, as the preferences of different agents become more independent. A higher $\varphi$ increases the utilization for both NL and RD, as agents can on average divert to more other customers if their preferred customer is carried by another agent. Even for $\varphi = 1$ the utilization is higher than the utilization of the KPRP in which the number of resources also equals the number of agents. In the KPRP there is exactly one customer per district, whilst in a VFHP with $\varphi = 1$ there is on average one customer per district, it is therefore possible that more than one customer is located in a given district.

| Strategy | $\bar{f}$ | $\bar{u}$ |
|---|---|---|
| NL | 83.0% | 0.415 |
| RD | 30.6% | 0.240 |
| LL | 57.0% | 0.357 |
| OPR | 73.9% | 0.457 |
| CA | 49.7% | 0.249 |
| SCA | 93.8% | 0.469 |
| SRD1 | 78.5% | 0.438 |
| SRD2 | 77.3% | 0.432 |
| SRD3 | 36.7% | 0.283 |
| SRD4 | 47.4% | 0.366 |

Table 2: Mean utilization fraction and utility in simulation experiments.

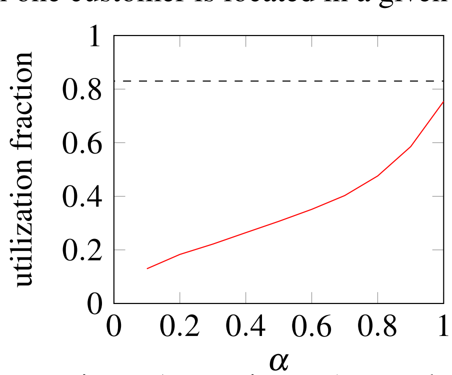

Figure 1: Varying $\alpha$ (NL and RD)

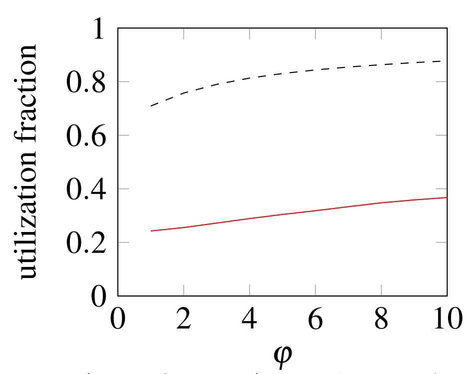

Figure 2: Varying $\varphi$ (NL and RD)

## Discussion

From our simulations, we conclude that it can be useful to introduce randomness in the decision process, as agents then drive to different resources. In the VFHP model, randomness represents different and non-deterministic behavior. A (strict) rank dependent choice as represented by the RD strategy results in lower utilization than a purely random strategy as represented by the NL strategy, as the number of distinct first preferences is reduced. With the added constraints for numerical experiments (in particular shared and individual utilities drawn from a random permutation with step size $\frac{1}{N}$ and influence of the individual utility of $\alpha = 0.5$) we assume that approx. $0.306 \cdot N$ customers are preferred by at least one agent as shown in (Martin 2017). It can also make sense to return to the same resource in order to make the own behavior predictable for other agents which can then randomly drive to other resources. Yet, agents cannot increase their utility by waiting or random behavior but will remain at average utility. The best strategy for our model is SCA, because agents return to the same district (that is: wait) unless the district is overcrowded which reflects the findings of Li et al. (2011) who analyze the benefits of waiting in a given district versus driving to a randomly chosen other location. Further research can include the actual destination of customers, forcing agents to wait in a district unless it is overcrowded.

From comparing our simulation results for the VFHP to the results of the KPRP, we conclude that districts can help increasing the utilization and utility, as agents can divert to other resources if their preferred customer is carried by another agent. For example, the utilization of the NL strategy rises from 63.2% in the KPRP (Chakrabarti et al. 2009) to 83.0% in the VFHP. This increase cannot be caused by different preferences, as agents are ignorant about the utility yielded by a customer in the NL strategy. Also, using districts is reasonable for vehicle for hire markets, as drivers can divert to another customer in close proximity in a short period of time or customers can walk up to central pickup points (i.e. taxi stands).

With this paper we aim at showing that different utilities influence the utilization yielded by implementing given strategies. Yet, the actual utilization of rank dependent strategies such as RD, LL, OPR, and SRD depends on the actual valuation and cost distribution which we currently model by a uniform distribution. The actual utility can depend on the location (for example, the expected valuation might be higher at the airport than downtown) and the distribution is not necessarily uniform. Our simulations can easily be adapted to implement other valuation and cost distributions. Further, the demand is deemed constant and equal to the supply. In the real world, customers who were carried at time $t$ will leave the market, whilst the other customers will either wait for a ride at time $t+1$ or leave the market (i.e. walk to the destination) with some probability. Those customers who wait for a ride increase the demand in the next iteration. Also, the number of customers waiting in some district depends on the time of day and is influenced by randomness.

Changes in the location of agents (due to driving to a district or carrying a customer) influence the individual utility and can worsen the utilization and utility of some of the strategies. In the presented model, returning to the same district yields a high utility (which occurs, i.e., when starting from the garage at the beginning of a shift), but returning to the same district can actually incur higher costs in an intra-day model. We therefore aim at extending the VFHP to include varying demand and agent movement.

## Conclusion

We propose a novel game model addressing issues caused by decentral coordination arising in vehicle for hire markets – the VFHP. We further studied several different strategies that agents, i.e. drivers, can implement. In vehicle for hire markets, all drivers can independently decide where to drive to pick up or await customers. The utilization is, therefore, lower than by a central coordination. Yet, it is impossible to implement a centralized algorithm, as agents are self-employed. We have modeled the VFHP as a generalization of the Kolkata Paise Restaurant Problem (Chakrabarti et al. 2009). Customers, that is resources, are located in districts. Agents choose a customer and then drive to the associated district. In this district they can carry any customer. Our contribution to the game model is this additional layer (districts) and an extended utility function for rank-based strategies which combines a valuation (payoff) with shared and individual costs. For this adapted model we test the utilization for a set of strategies which we derived from literature (Chakrabarti et al. 2009; Ghosh et al. 2013) and strategies we newly introduced. The newly added strategies are strict and stochastic rank dependent and help assessing the influence of the individual utility component.

We found that strict rank dependent behavior as dictated by the RD strategy results is a rather low utilization (in comparison to NL), as the number of distinct top priorities is lower than with randomly distributed priorities. SCA and OPR yield a high utilization. Both strategies are history-based and include a random component in the decision process. OPR also includes a rank-based component. From the strategies ignoring the history, SRD1 and SRD2 reach a utilization close to NL. In highly volatile and dynamic markets it can be beneficial to use history-independent strategies. There, agents as well as customers enter and leave frequently. Further, valuation and costs (and thus, the utility function) are subject to frequent changes. Future work will extend the VFHP to allow for variable utility functions and agents and customers entering and leaving at will.

## References

Arnott, R. 1996. “Taxi travel should be subsidized,” *Journal of Urban Economics* (40:3), pp. 316–333.

Arthur, W. B. 1994. “Inductive reasoning and bounded rationality,” *The American economic review* (84:2), pp. 406–411.

Cairns, R. D., and Liston-Heyes, C. 1996. “Competition and regulation in the taxi industry,” *Journal of Public Economics* (59:1), pp. 1–15.

Chakrabarti, A. S., Chakrabarti, B. K., Chatterjee, A., and Mitra, M. 2009. “The Kolkata Paise Restaurant problem and resource utilization,” *Physica A: Statistical Mechanics and its Applications* (388:12), pp. 2420–2426.

Challet, D., and Zhang, Y.-C. 1998. “On the minority game: Analytical and numerical studies,” *Physica A: Statistical Mechanics and its applications* (256:3), pp. 514–532.

Chen, M. K., and Sheldon, M. 2015. "Dynamic pricing in a labor market: Surge pricing and flexible work on the Uber platform," Tech. rep., Mimeo, UCLA.

Correa, J. R., Schulz, A. S., and Stier-Moses, N. E. 2005. "On the inefficiency of equilibria in congestion games," in *International Conference on Integer Programming and Combinatorial Optimization*, , Springer.

Cramer, J., and Krueger, A. 2016. "Disruptive Change in the Taxi Business: The Case of Uber," *American Economic Review* (106:5), pp. 177–182.

Ghosh, A., Biswas, S., Chatterjee, A., Chakrabarti, A. S., Naskar, T., Mitra, M., and Chakrabarti, B. K. 2013. "Kolkata paise restaurant problem: An introduction," in *Econophysics of Systemic Risk and Network Dynamics*, , Springer, pp. 173–200.

Hall, J., Kendrick, C., and Nosko, C. 2015. "The effects of Uber's surge pricing: A case study," *The University of Chicago Booth School of Business* .

Kim, H., Yang, I., and Choi, K. 2011. "An agent-based simulation model for analyzing the impact of asymmetric passenger demand on taxi service," *KSCE Journal of Civil Engineering* (15:1), pp. 187–195.

Lee, D.-H., Wang, H., Cheu, R., and Teo, S. 2004. "Taxi dispatch system based on current demands and real-time traffic conditions," *Transportation Research Record: Journal of the Transportation Research Board* (:1882), pp. 193–200.

Li, B., Zhang, D., Sun, L., Chen, C., Li, S., Qi, G., and Yang, Q. 2011. "Hunting or waiting? Discovering passenger-finding strategies from a large-scale real-world taxi dataset," in *Pervasive Computing and Communications Workshops (PERCOM Workshops), 2011 IEEE International Conference on*, , IEEE.

Linne+Krause Marketing-Forschung 2016. "Untersuchung zur Wirtschaftlichkeit des Berliner Taxigewerbes," `http://www.stadtentwicklung.berlin.de/verkehr/politik/`

taxi/download/untersuchung_wirtschaftlichkeit_taxi_berlin.pdf, retrieved at 12.02.2017.

Martin, L. 2017. *Extending Kolkata Paise Restaurant Problem to Dynamic Matching in Mobility Markets*, Master's thesis, Technical University of Munich.

Milchtaich, I. 2000. "Generic uniqueness of equilibrium in large crowding games," *Mathematics of Operations Research* (25:3), pp. 349–364.

Monderer, D., and Shapley, L. S. 1996. "Potential games," *Games and economic behavior* (14:1), pp. 124–143.

Rayle, L., Shaheen, S., Chan, N., Dai, D., and Cervero, R. 2014. "App-based, on-demand ride services: Comparing taxi and ridesourcing trips and user characteristics in san francisco," *University of California, Berkeley, United States Rogers, B(2015) The social costs of Uber James E Beasley School of Law, Temple University, Philadelphia, United States* .

Rogers, B. 2015. "The social costs of Uber," *U Chi L Rev Dialogue* (82), p. 85.

Rosenthal, R. W. 1973. "A class of games possessing pure-strategy Nash equilibria," *International Journal of Game Theory* (2:1), pp. 65–67.

Seow, K. T., Dang, N. H., and Lee, D.-H. 2010. "A collaborative multiagent taxi-dispatch system," *IEEE Transactions on Automation Science and Engineering* (7:3), pp. 607–616.

Trigo, P., Jonsson, A., and Coelho, H. 2006. "Coordination with collective and individual decisions," in *Advances in Artificial Intelligence-IBERAMIA-SBIA 2006*, , Springer, pp. 37–47.

Wong, R., Szeto, W., and Wong, S. 2015. "A two-stage approach to modeling vacant taxi movements," *Transportation Research Procedia* (7), pp. 254–275.

Yang, H., Ye, M., Tang, W. H.-C., and Wong, S. C. 2005. "A multiperiod dynamic model of taxi services with endogenous service intensity," *Operations research* (53:3), pp. 501–515.