Capacity, Information and Minority Games in Public Transport

Paul Bouman† Leo Kroon† Gábor Maróti‡ Peter Vervest†
PBouman@rsm.nl LKroon@rsm.nl G.Maroti@vu.nl PVervest@rsm.nl

†Rotterdam School of Management, Erasmus University
Department of Technology & Operations Management
Burgemeester Oudlaan 50, 3062 PA, Rotterdam, The Netherlands

‡Netherlands Railways
Utrecht, The Netherlands
‡VU University, Amsterdam, The Netherlands

ABSTRACT

Many public transport operators are faced with high peak demands. Often this leads to crowded vehicles and discomfort for the passengers. The increasing use of information technologies creates new opportunities for passengers to avoid crowding. However, the role of crowding in the dynamics of a public transport system is not well understood. With the definition and implementation of a model based on the minority games, a class of games that deals with crowding dynamics, we aim to provide public transport operators with insights to deal with crowded situations.

We propose an extension of a minority game where multiple resources and heterogeneous agent preferences are included. We have conducted two simulation studies, aimed at investigating the dynamics of crowding within public transport. In our first experiment we investigate the effect of the availability of information on crowding. In a second experiment we study the dynamic optimization of capacities according to a rolling stock circulation model. We find that both the availability of information disclosed and the chosen capacity optimization mechanism have an impact on the number of agents utilizing resources and their payoffs. As such, these models will allow us to develop new operator policies to deal with crowded situations in the future.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Miscellaneous—*Coherence and coordination*

General Terms

Experimentation, Management, Performance

Keywords

capacity, coordination, information, minority games, public transport, resource allocation

1. INTRODUCTION

Operators in public transport are often faced with peak demands, typically during the morning and afternoon rush hours. As a result, vehicles can become very crowded, greatly reducing the comfort experienced by the passengers. As information technologies enable passengers to have more direct communications with the public transport operators and have more freedom to work at different locations, passengers are gaining more opportunities to avoid crowded situations. However, the impact of crowding on passenger behavior and the interaction between railway operations and passengers is not well understood. In this paper we develop a model, based on the concept of minority games, that allows us to study the dynamics of crowding in public transport through computational experiments and evaluate the impact of operational and behavioral models on a number of performance measures, most importantly the utilization of available capacities.

Since the "El-Farol Bar Game" [1] was first introduced in 1994, the concept of the *minority game* has received a lot of attention from researchers. One of the great strengths of this model lies in the simplicity of its description: a population of agents have to decide every Thursday night whether to go to the bar or not. Once they go the bar, they have a positive payoff if less than 60% of the population goes to the bar, while they have a negative payoff if it is too crowded. As everyone makes this choice every Thursday, the El-Farol Bar Game has an iterative nature. While historic information is provided, the interesting aspect comes from the fact that there is no direct coordination between the agents.

Issues related to limited availability of resources and a lack of explicit coordination occur in many real world systems. The applications of these models include car traffic [2], congestion in computer networks [8] and financial markets [4]. While these types of applications were considered earlier from a game theory perspective, most notably under the name of *congestion games* [11], the novelty from the "El-Farol Bar" study was the application of a complex systems approach enabled by simulation of a repeated game, while game theory is mostly concerned with the properties of equilibria.

In this paper, we focus on minority games where the operator cannot control agent behavior, but has control over the disclosure of information and the system capacities. The main application domain is public transport systems, where passengers share vehicles depending on their chosen route and time of travel. If a connection is operated frequently, passengers with some flexibility in their schedule can try to avoid crowded situations by shifting time of travel. Since it is reasonable that a passenger does not want to travel at any time, we introduce the concept of individual choice sets representing the acceptable choices. To our best knowledge, this type of heterogeneity of the choice sets has not been studied in the context of minority games before.

Within public transport systems, there are many opportunities to provide passengers with additional information: many stations and vehicles have screens with travel information, and many passengers use smart phones to receive information during their journeys. The increasing adoption of smart card ticketing systems allows operators to have accurate data on the utilization of each vehicle. As operators in railway and metro systems can extend or shorten the trains [5] and bus operators can employ different vehicle sizes, adaptive capacity allocation is becoming a possibility.

The main observation in the original "El-Farol Bar Game" simulations [1] is that even though individual agents keep switching their preferred predictive model, the aggregate utilization of the bar converges to the efficient level. In order to explain this phenomenon the minority game was introduced, where the utilization history was replaced with a history of binary values indicating whether the bar was overcrowded or not. The main idea of this approach is that the set of all possible deterministic strategies can be characterized so that methods from statistical mechanics can be applied [3].

The remainder of this paper is organized as follows: in Section 2 we introduce our class of minority games. In Section 3 we discuss the architecture of our simulation and agents. This simulation framework is then applied in order to investigate the effect of different information policies in Section 4. In a second simulation study we evaluated the effect of rolling stock optimization in the context of public transport (Section 5). In Section 6 we show that the inclusion of individual choice sets and scoring functions leads to NP-hardness of maximizing the efficiency of a given system. We discuss our findings and plans for future research in Section 7.

Related Work

A variation of the minority games are the *resource allocation games*, introduced by [7]. This extension of minority games introduces multiple resources and capacities that vary over time. Conditions are given under which the agents can use a social network structure in order to adapt efficiently to variations of the capacities. The fluctuations of the capacities considered in the studies associated with the *resource allocation games* only depend on time and do not depend on the distribution of agents over the resources during the game.

While the body of knowledge on learning techniques for agents in minority games [10] is very useful for the engineering and design of artificial agents, it is a question whether it is applicable within systems where real humans are involved. Selten et al. [13] conducted a laboratory experiment involving route-choice. The participants could be divided into three groups: participants who had the tendency to switch away from a road if it was congested during the previous round, participants who had the tendency to stay on their current road regardless of it being congested during the previous round, and participants who were harder to classify. Although the participants showed different types of behavior, the distribution of the participants over the roads approached the equilibrium very closely.

2. A MODEL FOR CROWDING DYNAMICS

The general scheme of model is that in each round every agent decides whether he will use one or more resources or refrains from doing so. Using a resource gives the possibility to gain a positive payoff or a negative payoff depending on the utilizations encoutered. If he does not use any resource, the payoff will be neutral, i.e. zero.

We define symbols for the resources, the agents and payoffs. The resources will be defined in the following way:

- A set $\mathcal{R} := \{1, 2, \ldots, m\}$ of m resources.
- A *soft* capacity function $\mathrm{cap} : \mathcal{R} \to \mathbb{Z}^+$.

Thus there are m resources, each of which having an associated capacity. Note that we define *soft* capacities: they can be violated, but everyone in such a situation should have a negative payoff. Based on the capacity we define the utilization of a resource as the fraction of its capacity that is occupied. The typical example in public transport is the number of passengers divided by the number of seats. As the game is played iteratively, the transport operator can adapt the capacities based on observations recorded during earlier rounds of the game. We also define the preferences and payoffs of the agents that play the game:

- A set $N = \{1, 2, \ldots, n\}$ of n agents.
- A non-empty collection C_i of subsets of $\mathcal{R}$.
- A scoring function $s_i : \mathbb{Q}^m \times C_i \to \mathbb{R}$ for each agent $i \in N$

During each round, every agent should choose one of the options in its choice set. We assume that every choice set contains the empty set as a neutral option, but this is not strictly necessary. We can describe the outcome of a round based on the choices made by all agents. If an agent i chooses to use a set of resources $c \in C_i$, we set the indicator variable x_{ic} to 1. The set of all vectors of x_{ic}'s describing a valid outcome is thus defined by

$$\mathcal{O} = \{x \mid \forall i \in N : \sum_{c \in C_i} x_{ic} = 1, x_{ic} \in \{0, 1\}\}. \tag{1}$$

Given the outcome vector x for a round, we can calculate the utilization of the resources. We define a vector $u(x) \in \mathbb{Q}^m$ that contains an entry for each resource. The entry $u_r(x)$ for resource $r \in \mathcal{R}$ is calculated as follows:

$$u_r(x) = \frac{\sum_{i \in N} \sum_{c \in C_i : r \in c} x_{ic}}{\mathrm{cap}(r)}. \tag{2}$$

While in principle s_i can be a general scoring function, for ease of analysis we will use the restricted class of *threshold based scoring functions*. These scoring functions have a payoff of -1, 1 or 0 depending on a individual threshold θ_i and the maximum encountered utilization. The scoring function itself is then defined as follows:

$$s_i(u, c) = \begin{cases} 0 & \text{if } c = \emptyset, \\ 1 & \text{if } \max_{r \in c} u_r \leq \theta_i, \\ -1 & \text{otherwise.} \end{cases} \tag{3}$$

Performance Measures

Since we want to analyze the behavior of an agent population, we will introduce some measures that are of analytic interest and can be recorded during a simulation. We define the # symbol to denote the cardinality of a set (e.g. $\#\{6, 9\} = 2$). Given an outcome $x \in \mathcal{O}$ during any of the rounds of the game, we can calculate the following observations:

- $\mathrm{utl}(x) = \frac{1}{n}\#\{i \in N : x_{ic} = 1, c \neq \emptyset\}$, i.e. the fraction of agents utilizing a resource.
- $\mathrm{pos}(x) = \frac{1}{n}\#\{i \in N : x_{ic} = 1, s_i(y, c) > 0, y = u(x)\}$, i.e. the fraction of agents with a positive payoff.
- $\mathrm{posc}(x) = \frac{\mathrm{pos}(x)}{\mathrm{utl}(x)}$ is the fraction of agents with a positive payoff among the agents who utilize a resource.
- $\mathrm{avg}(x) = \frac{1}{n}\sum_{i\in N}\sum_{c\in C_i} s_i(u(x), c)x_{ic}$ is the average payoff of the agents.

3. ARCHITECTURE OF THE AGENTS AND SIMULATION

Given an instance of the game, a simulation still depends on two more aspects: the way the agents make their decisions and to which extent the agents can observe the outcome of the previous rounds. As we want to be able to evaluate the effect of different types of agent behavior, we will allow different types of agents in the population. We will introduce a number of types in Section 3.1. We first define the main steps that will be executed in each round of the simulation:

1. Let every agent $i \in N$ choose one option $c \in C_i$ from its choice set according to its agent type.
2. Calculate the outcome vector x and corresponding utilization vector $u(x)$ accordingly.
3. Let every agent $i \in N$ observe, learn and process its score $s_i(u(x), c)$ based on its agent type.
4. The operators lets every agent $i \in N$ observe, learn and process information based on the active information policy and the utilization vector u.

From these steps we can see the necessary ingredients for an agent implementation within this simulation scheme: an agent needs a choice function and can optionally implement a method to process incoming scores and information.

3.1 Agent Types

The most simple agent type is the *random agent*, who selects a choice from its choice set uniformly at random in each round. This agent type is useful for both benchmarking purposes, validating the simulation architecture analytically, and to model noisy behavior within the population.

The more complicated agent types will make decision based on observations during earlier rounds of the game. For these agent types, step 4 of the simulation process in a round can have an effect on step 1 in the next round. The number of rounds the agents look back is referred to as the *memory length*. An important finding in the minority game model is that the most efficient utilization is reached when the memory length of the agents is proportional to the logarithm of the total number of agents [12].

The second type of agent, the *average payoff agent*, applies a simple reinforcement learning heuristic. Reinforcement learning strategies have received notable attention in the literature, and we take one of the most simple ones as an example. As such, our *average payoff* agents perform exploration during 10% of the rounds by making a random choice, while they exploit the observed average payoff values during 90% of the rounds. In case multiple choices have the best average payoff, the tie is broken by picking one option uniformly at random.

A variation of average payoff agent is the *average utilization agent*, who uses the same reinforcement learning heuristic to learn the average utilization of the resources. The main difference is that this agent uses the information received to learn the average utilization and pick the choice with the lowest average utilization, or the neutral option if this choice has still higher average utilization than its threshold.

The last type of agent, the *predictive agent*, aims to predict future utilizations in order to find the best choice. If the agent can predict future utilizations, the agent can generate a fictitious utilization vector and evaluate the expected score of each choice. This agent type is similar to the one studied in original El Farol Bar paper [1]. In a round with index t, the agent checks which of its personal heuristics was most accurate in round $t-1$ and uses this one to predict utilizations in round t. As our model introduces the concept of multiple resources, there can be situations where an agent does not know all historic utilizations of each resource. We calculate the accuracy of each heuristic based only on the information that is available. As availability of utilization information is defined on the agent level, an agent can compare the heuristics using the same data set.

We implemented the following predictive heuristics: replicate the oldest utilization in memory, take the average of the utilizations in memory or fit a linear regressive model on the utilizations in memory.

3.2 Information Policies

At the end of each simulation round, we let each agent process information and observations on the utilization of the resources. We define a unit of information as a 3-tuple (t, r, u) consisting of the round of the game t, a resource r and a utilization vector u. As the agent can have multiple resources in its choice set, it should be able to receive and process multiple pieces of information each round. In general, an *information policy* is a set of rules that determine the information offered to each agent in each round. While there are very many information policies possible, we propose four basic ones.

In public transport, the fact that an agent is using a resource allows it to observe the utilization. Thus in our most basic information policy, *private information*, an agent receives exact information for the resources in its choice.

On top of private information, the entity or agent controlling the resources could monitor the utilizations and try to attract more agents in case a resource r has a low utilization, say less than 40% of its capacity. In such a situation the information policy can state that additional information regarding resource r should be provided to all agents. We will refer to this type of policy as *adaptive information*.

In some situations there are information systems that provide information on the crowding of a resource. A real life example one can think of is a smart phone application of a

Table 1: Results of the simulation study where different information policies are evaluated. The minimum, average and maximum utl (fraction of agents utilizing a resource) and posc (fraction of agents who have a positive payoff among those that utilize a resource) values measured for each of the 66 population mixtures are reported.

utl ($n = 50$)	**private**	**adaptive**	**estimate**	**full**
Minimum	0.63	0.75	0.75	0.75
Average	0.77	0.90	0.91	0.91
Maximum	0.91	0.97	0.99	0.99

posc($n = 50$)	**private**	**adaptive**	**estimate**	**full**
Minimum	0.90	0.90	0.89	0.87
Average	0.95	0.93	0.91	0.90
Maximum	0.99	0.97	0.93	0.92

utl ($n = 100$)	**private**	**adaptive**	**estimate**	**full**
Minimum	0.45	0.66	0.70	0.73
Average	0.60	0.70	0.73	0.77
Maximum	0.75	0.75	0.82	0.85

posc ($n = 100$)	**private**	**adaptive**	**estimate**	**full**
Minimum	0.46	0.46	0.41	0.30
Average	0.77	0.65	0.60	0.48
Maximum	0.98	0.84	0.72	0.54

utl ($n = 200$)	**private**	**adaptive**	**estimate**	**full**
Minimum	0.20	0.40	0.46	0.51
Average	0.42	0.48	0.52	0.56
Maximum	0.75	0.75	0.75	0.78

posc ($n = 200$)	**private**	**adaptive**	**estimate**	**full**
Minimum	0.02	0.02	0.02	0.02
Average	0.49	0.35	0.27	0.19
Maximum	0.94	0.68	0.42	0.26

public transport operator, that shows one, two or three icons based on the forecasted crowdedness of a vehicle, reducing the utilization level provided to the agent to a few discrete values. This idea is capture by the *estimate information policy*, where the utilization of each resource is rounded up to either 0, $\frac{1}{4}$, $\frac{2}{4}$, and so on, similar to the 3 symbol crowding indicators provided by some operators. This rounded utilization is then provided to all agents. We should take care that we send the rounded utilizations in case the agent did not observe the utilization by itself, and use exact utilization otherwise.

In the final template, we send out exact information on every resource to every agent in each round – thus in this situation the agents have *full information*.

4. EVALUATING INFORMATION POLICIES

In our first experiment, we evaluate the four information policies in a population of agents that use public transport to travel from a single origin to a single destination, but can choose for different times of travel. As such their choice sets contain only singleton resources, reflecting the departure times a public transport service is scheduled and the empty set as a neutral option, reflecting a journey by car or staying at home. We find that increasing the available information leads to a greater number of agents utilizing the public transport system, but at the cost of the average payoff. However, the magnitude of this effect is influenced by the ratio of population size and available capacity.

4.1 Experimental Setup

In our experiments, we work with $m = 10$ resources representing the departure times. Every choice set C_i contains $\emptyset$ and 3 different singleton sets picked uniformly at random from $\mathcal{R}$ without replacement. For each agent i we use a threshold based scoring function with $\theta_i \in \{\frac{5}{10}, \frac{6}{10}, \ldots, 1\}$ picked uniformly. The capacity of each resource is fixed to 10, i.e. $\text{cap}(r) = 10$.

As we have 100 units of capacity available each round, we consider a high capacity scenario with $n = 50$ agents, a regular scenario with $n = 100$ agents and a low capacity scenario with $n = 200$ agents. For each of these scenarios, we vary the population by picking all pairs $p, q \in 0, 1, \ldots, 10$ such that $p + q \leq 10$. Our population then consists of $10p$ *random* agents, $10q$ *average utilization* agents and $10(10 - p - q)$ *predictive* agents. In total 66 population mixtures are evaluated. For each mixture of agent implementations we regenerate the choice sets and thresholds 100 times. For a given instance of the choice sets we run the experiment 25 times, regenerating the predictive agents 5 times if they are part of the population. Thus, in total we run 2500 simulations per combination of population mixture and population size. As we want to ignore the warm-up period of the simulation and like to interpret the rounds as days, the measures are recorded from round 10 to 40 during each simulation run.

The predictive agents each have an individual randomly selected set of 3 random predictive heuristics from the following list: average heuristic with memory lengths of either 4, 5, 6 or unlimited, linear regression with memory lengths of 4, 5, 6 or unlimited, replicate the oldest observation with memory length either 1, 2 or 3.

4.2 Results and Discussion

The results of our simulation experiments are presented in Table 1. If we look at the left column of Table 1, we can verify that when we increase the level of information provided to the agents, the number of agents utilizing a resource increases. If we look at the average values from private to adaptive, we can see that the 0.13 increase for $n = 50$ scenario is greater than the 0.06 increase for $n = 200$. These numbers suggest that the effect of information depends on the units of capacity available per agent in the population.

If we look at the fraction of agents utilizing a resource with a positive payoff (this can be interpreted as customer satisfaction) in the right column of Table 1, we can see that increasing the level of information decreases the posc value. This seems intuitive, since adding information attracts more agents, and having more agents increases the likelyhood of crowding. Again, the amount to which the posc value decreases when we move from the *private* to the *adaptive* case is impacted by the amount of capacity available per agent: for $n = 50$ the decrease of 0.02 is less dramatic than the 0.14 decrease of in the $n = 200$ case.

For future work we will investigate whether better information policies can be designed. The are also questions regarding the effect of noise in communications, such as technical problems at the side of the operator or agents ignoring information sometimes. We are curious to learn whether such noise could lead to less correlated agent behavior and whether this can lead to better system efficiency.

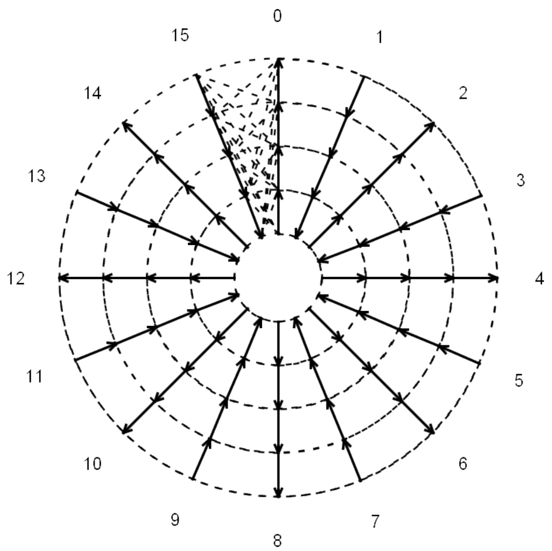

Figure 1: The network of trips along a line during the 16 time slots with the possible movements of rolling stock units outside the regular trips represented by dashed lines.

5. CAPACITY OPTIMIZATION IN PUBLIC TRANSPORT

In our second experiment, we want to evaluate the effect of rescheduling capacities on the crowding dynamics in the system. Consider a public transport scenario where a train moves back and forth a line of 5 stops. The train drives 8 full cycles per day and as moving along the line in one direction gives us 4 trips between the stops, the timetable consists of $4 \cdot 2 \cdot 8 = 64$ trips are offered each day. As individual travelers want to travel between two stops that are not necessarily connected by a single trip, a journey can consist of one or multiple trips. We will assume passengers always want to travel in the direction of their destination and as such for each origin-destination (OD) pair there are 8 different time slots at which passengers can make their journeys.

In order to facilitate the flow of passengers, the trains need to be long enough in order to allow comfortable transportation. To achieve this, the operator monitors utilization of vehicles and adapts the assigned number of rolling stock units to each train accordingly. The operator can decide how often the observed utilizations are evaluated to build a new rolling stock model. In this experiment, we will assume that this will happen periodically. The number of rounds after which the operator produces a new rolling stock schedule will be referred to as the *reschedule period*, denoted by an integer k.

5.1 Capacity Allocation

As the use of rolling stock units determines a significant amount of the operational costs of a public transport operator, they try to monitor the utilization of the train vehicles and adapt the capacities if necessary. The typical model used to determine the rolling stock allocation in these situations is by constructing the network of possible train movements, specifying a minimum demand on the arcs that correspond to passenger trips and look for a minimum cost circulation [6] based on operational costs.

We implemented a module in the simulation that represents an operator which dynamically optimizes demand. During each round of the simulation, train utilization for each trip is recorded. After k rounds, the demand of a trip is set to $\mu + 2\sigma$, where μ is the mean utilization and σ its standard deviation during those k rounds. We chose this rule because similar rules are employed by real operators. The capacities of each trip are then calculated according to a rolling stock circulation, where we define a cost of 1000 per unit used and a costs of 1 for moving a unit between consecutive stops on the line. These numbers represent that buying and maintaining rolling stock units is a lot more costly than moving them around. We also assume that storing a unit at a station does not impose any costs. As a result, the minimum cost rolling stock circulation will minimize the number of units required before minimizing the movement costs, given that the defined demand must be met.

We use a minimum cost circulation algorithm [6] (which shares quite a lot of similarity with the well known augmenting path methods for max flow) to obtain the capacities. Although more efficient algorithms exist for this problem, the augmenting path method is straightforward to implement and fast enough for our simulations. The input network is visualized in Figure 1. The straight arcs represent movements between the stops and must carry the determined demands. The circular arcs represent storing a vehicle at a stop. The overnight arcs represent the purchase costs of the vehicles and the overnight balancing movements.

While the algorithms employed by operators need to take many different types of rolling stock and regulations into account [5], for reasons of simplicity and interpretability we assume that we have only one type of rolling stock with a nominal capacity of 10 seats.

5.2 Experimental Setup

In order to set up the simulation, we define a resource set that consists of the trips, so based on the 5 stops and 16 timeslots, we get $m = 64$ resources. The choice set of an individual agent is generated as follows: we pick two stops $o \neq d$ from among the five stops. By choosing o as the origin and d as the destination, the direction along the line is defined. We then pick 3 from the 8 available time slots corresponding to this direction in order to define the acceptable journeys. The choice set then consists of the empty set and the sets of trips corresponding to the journeys drawn randomly. Again we work with threshold based scoring functions where the threshold is picked uniformly from $\{\frac{5}{10}, \frac{6}{10}, \dots, 1\}$.

For the purpose of simplicity, we use only one type of agent during this experiment: the *average payoff agent*. One of the reasons to choose this agent implementation is that software packages for dynamic traffic equilibrium computations with feedback use this approach. We pick the number of agents simulated as $n = 1000$. The reason to take a relatively large agent population is because we have 64 trips and as each trip should have at least one unit of rolling stock available, the available capacity is at least 640. In order to have a high probability to facilitate all the demand during the first rounds of the simulation, we set the initial demand of rolling stock units for each trip to 5. We also checked initial rolling stock counts of 1 and 10 units and our findings were robust under these variations.

Our goal is to evaluate the effect of different rescheduling periods. As the demand observed depends on the length of

Table 2: Results of the simulation study where the effect of rolling stock optimization on average payoff, operator costs and rolling stock units required is evaluated. The measures at round 100 of each simulation are reported, for different rescheduling periods (k) of 1, 5 and 10.

$k = 1$	min	mean	(±std.)	max
utl	0.38	0.42	(±0.02)	0.46
posc	0.60	0.73	(±0.05)	0.84
avg	0.08	0.19	(±0.04)	0.28
cost	2092	2532	(±498.5)	3112
units	2	2.43	(±0.497)	3

$k = 5$	min	mean	(±std.)	max
utl	0.58	0.60	(±0.015)	0.66
posc	0.74	0.85	(±0.03)	0.92
avg	0.28	0.43	(±0.04)	0.50
cost	3166	4120	(±314)	4120
units	3	3.9	(±0.31)	5

$k = 10$	min	mean	(±std.)	max
utl	0.62	0.66	(±0.01)	0.70
posc	0.76	0.86	(±0.03)	0.92
avg	0.36	0.48	(±0.04)	0.56
cost	4188	4219	(±142)	5208
units	4	4.02	(±0.14)	5

the rescheduling period, we evaluated periods of 1 round, 5 rounds and 10 rounds. For each of these rescheduling policies, we generate 50 agent populations of choice sets with corresponding thresholds. For each population we then run 2 simulations of 50 rounds.

5.3 Results and Discussion

The time series distributions of the observed values of utl and posc during the 100 simulation runs for each of the three policies are presented in Figure 2. The distributions of average payoffs, the operational costs (according to the minimum cost rolling stock circulation) and the number of rolling stock units utilized during the last round of the simulation are presented in Table 2.

As we increase the length of the period, we can observe that utl increases. If the reschedule period is 1 round, we can observe from Figure 2a that it converges to a mean of 0.42. For a reschedule period of 5 rounds it converges to a mean of 0.6 (Figure 2c) and for a reschedule period of 10 rounds it converges to a mean of 0.66 (Figure 2e). One possible explanation for the fact that longer periods give a higher value for utl is that a longer reschedule period has the potential to yield a more stable mean and possibly more accurate standard deviation (except for the case where the period is 1; then the standard deviation is always 0). As the mean and standard deviation have direct effects on the demand and thus the capacities that are calculated, they seem likely causes for the observed behavior.

For the fraction of agents that utilize a resource and have a positive payoff posc, we can observe that it converges to a value of 0.73 for a reschedule period of 1 round (Figure 2b), to a value of 0.85 for a reschedule period of 5 rounds (Figure 2d) and to a value of 0.86 for a reschedule period of 10 rounds (Figure 2f). The 1 round scenario has a higher standard deviation than the other scenarios .

While both the utl and posc have higher averages for longer reschedule periods, Figure 2 also shows slower convergence for longer reschedule periods. Table 2 also suggests that longer reschedule periods lead to higher costs and a higher number of rolling stock units required. However, this can be explained by the increase of the utl value. A final interesting observation in Table 2 is that the average payoff for the agents also increases if we use longer periods.

Our results suggest that there are many disadvantages for the single round reschedule period. Increasing the period may lead to higher costs, but the number of passengers using one of the trains increases as well, which can lead to extra revenue. For future research we aim to search for different approaches to determine the demand for the rolling stock circulation based on the utilizations observed in the simulation. A different approach to the $\mu + 2\sigma$ rule would be to adapt demand based on observed scores.

6. COMBINATORIAL ASPECTS

In the original "El-Farol Bar" model, it is not difficult to see that the ideal utilization of the bar lies at 60%, because all agents have the same payoff. In our extension it is not easy to determine the ideal utilization, as we are allowed to have agents with different scoring functions assigned to the same resource. As a result, it can be the case that for a single resource, some agents have a positive payoff and others have a negative one. The individual choice sets complicate matters even further. As a result, it is a combinatorial problem to maximize pos(x). We will show this by proving the NP-completeness of the related decision problem.

THEOREM 1. *For a given instance of the game, deciding whether there exists a valid outcome $x \in \mathcal{O}$ such that all agents have a positive payoff (i.e. whether* pos(x) *is equal to* 1*) is NP-complete, even if we have threshold scoring functions with* 2 *different thresholds and we allow only singleton resources in the choice sets.*

PROOF. We will show NP-hardness by reduction from the k–SET COVER problem [9]. In the k–SET COVER problem we are given a collection $\mathcal{A} = \{A_1, \ldots, A_n\}$ of n sets, a set of all elements $U = \bigcup_{i \in N} A_i$ and a positive integer k. We have to decide whether there exists a subset $\mathcal{A}' \subseteq \mathcal{A}$ such that $|\mathcal{A}'| \leq k$ and $\bigcup_{A \in \mathcal{A}'} A = U$.

We now introduce $|U|$ regular agents and $|\mathcal{A}| - k$ grumpy agents. We introduce a mapping between the sets in $\mathcal{A}$ and the resources. Each element in $e \in U$ is represented by a regular agent which has a choice set that consists of singleton resources corresponding to the sets in $\mathcal{A}$ containing e. The grumpy agents have a choice set with a singleton for every resource. We define the payoff functions such that the regular agents have a positive payoff as long as they have chosen a resource, and the grumpy agents have a positive payoff if they are exclusively assigned to a resource (if we fix all cap$(r) = 1$, then $\theta_i = n$ if i is a regular agent and $\theta_i = 1$ if i is a grumpy agent). As a result the grumpy agents can only have a positive payoff if they are assigned to resources in such a way that all the other agents can be assigned to the remaining resources. By construction of the choice sets, this is only possible if the remaining k resources that are not utilized by the grumpy agents correspond to sets that are able to cover all elements. Thus, we have reduced the k-SET COVER problem into our decision problem with 2 threshold scoring functions and singleton choice sets.

NP-completeness then follows from the fact that given a vector x, we can easily check whether it is feasible and whether indeed pos$(x) = 1$. □

In order to understand how the reduction works, we provide an example in Table 3. Here the A's and e's represent the sets and elements of the k-SET COVER instance. The

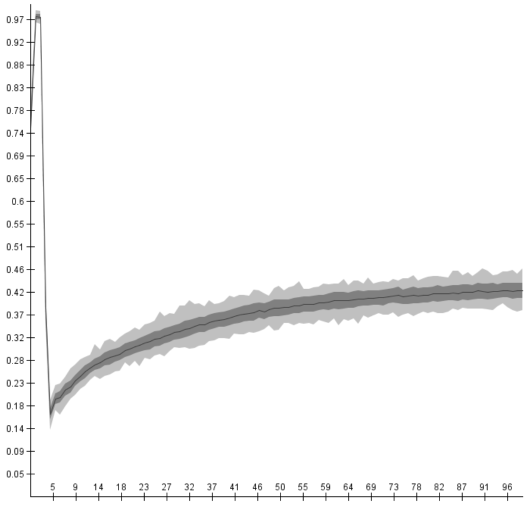

(a) utl when rescheduling period $k = 1$

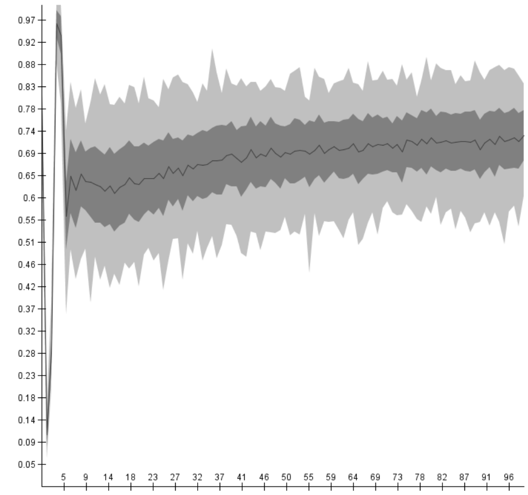

(b) posc when rescheduling period $k = 1$

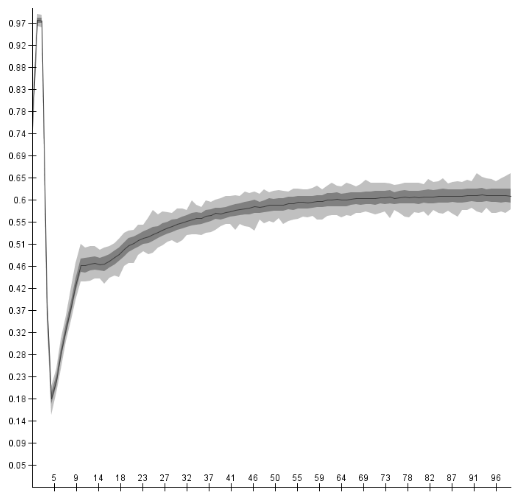

(c) utl when rescheduling period $k = 5$

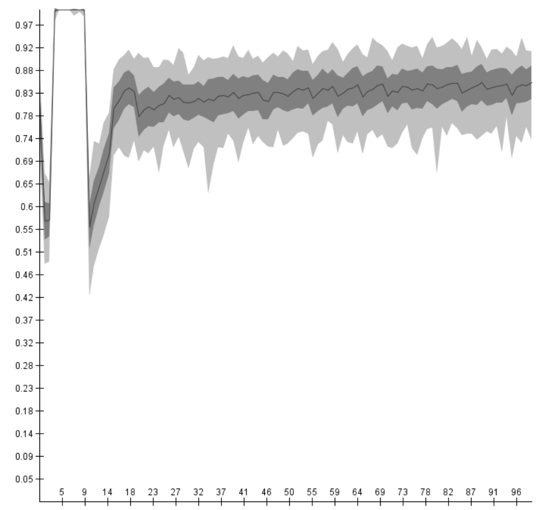

(d) posc when rescheduling period $k = 5$

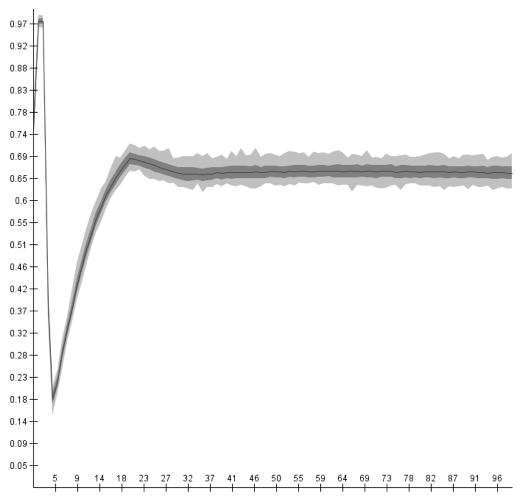

(e) utl when rescheduling period $k = 10$

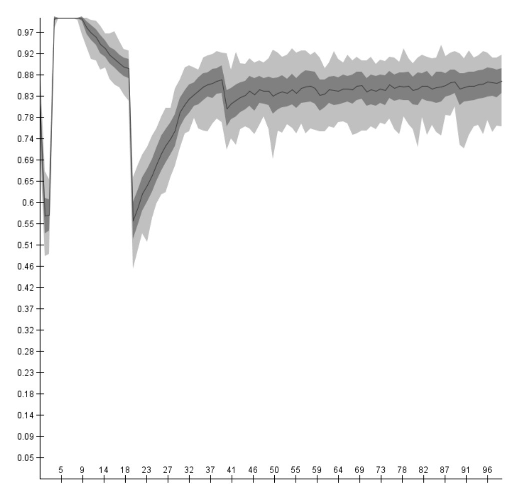

(f) posc when rescheduling period $k = 10$

Figure 2: Results of the capacity rescheduling experiments. The dark line shows the mean over all 100 experiments, the dark gray area is one standard deviation away from the mean and the light gray area shows the minimum and maximum values observed.

Table 3: An example reduction from k-Set Cover to a game instance.

	e_1	e_2	e_3	e_4	e_5
A_1	×	×			
A_2				×	×
A_3		×	×		
A_4			×	×	
A_5	×				×

(a) An example instance of a k-Set Cover problem.

Agent	C_i
a_1	$\{\emptyset, \{1\}, \{5\}\}$
a_2	$\{\emptyset, \{1\}, \{3\}\}$
a_3	$\{\emptyset, \{3\}, \{4\}\}$
a_4	$\{\emptyset, \{2\}, \{4\}\}$
a_5	$\{\emptyset, \{2\}, \{5\}\}$
g_1	$\{\emptyset, \{1\}, \{2\}, \{3\}, \{4\}, \{5\}\}$
g_2	$\{\emptyset, \{1\}, \{2\}, \{3\}, \{4\}, \{5\}\}$

(b) The corresponding choice sets for $k = 3$

corresponding game instance contains the a agents for the elements and $5 - k = 2$ grumpy agents denoted by the g agents. For $k = 3$, we can assign the two grumpy agents to resource 4 and resource 5, as the other agents are covered by the remaining resources. If we would now change k to 2, we would need to add an additional grumpy agent. However, we cannot give a positive payoff to both this additional grumpy agent and all the regular agents at the same time. This is consistent with the fact that there is no solution for the k-Set Cover instance if $k = 2$.

7. CONCLUSION AND FUTURE WORK

We have evaluated the effect of information disclosure and capacity optimization in a minority game designed to study crowding effects in public transport. The inclusion of heterogeneous agents poses many new challenges. From the theoretical perspective there are questions to what extend observations for the original minority game, such as the relation between memory length and efficiency, still apply. From the practical perspective, the question is whether an operator can influence and manage the cooperation of the agents in order to stimulate the efficient utilization of the vehicles. We have conducted two simulation studies where we focused on the practical challenges. In the first study we evaluate the effect of different information policies in a scenario where every agent uses at most a single trip every round. We find that disclosing more information attracts more agents, but that this comes at the cost of lower payoffs. This trade-off is influenced by the number of agents and the available capacity in the system.

In the second simulation study, we evaluate the effect of adaptive capacity management in the context of railway transportation. Here the agents make a journey along a line. They have to choose a time to travel between an individually assigned origin and destination stop every round. As such journeys can cross multiple stops and thus overlap on the line, more complex patterns of agent interaction can emerge. We find that the number of rounds utilizations are recorded before capacities are re-optimized has an impact on the number of agents utilizing the system and their payoffs. Rescheduling every round seems to lead to worse system performance than rescheduling every 5 or 10 rounds.

Our studies show that we are able to evaluate and compare the effects of different policies for information and capacity. The question remains whether we can improve on the policies we evaluated. We think that policies that act on agents that repeatedly have a low payoff are an interesting area for further research.

8. REFERENCES

[1] W. Arthur. Inductive reasoning and bounded rationality. *The American economic review*, 84(2):406–411, 1994.

[2] A. Bazzan, R. Bordini, G. Andrioti, R. Vicari, and J. Wahle. Wayward agents in a commuting scenario (personalities in the minority game). In *Proceedings of the Fourth International Conference on MultiAgent Systems*, pages 55–62. IEEE, 2000.

[3] D. Challet, M. Marsili, and R. Zecchina. Statistical mechanics of systems with heterogeneous agents: Minority games. *Physical Review Letters*, 84(8):1824, 2000.

[4] D. Challet, M. Marsili, and Y.-C. Zhang. Minority games: interacting agents in financial markets. *OUP Catalogue*, 2004.

[5] P. Fioole, L. Kroon, G. Maróti, and A. Schrijver. A rolling stock circulation model for combining and splitting of passenger trains. *European Journal of Operational Research*, 174(2):1281–1297, 2006.

[6] L. R. Ford and D. R. Fulkerson. *Flows in networks.* Princeton university press, 1962.

[7] A. Galstyan, S. Kolar, and K. Lerman. Resource allocation games with changing resource capacities. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pages 145–152. ACM, 2003.

[8] B. A. Huberman and R. M. Lukose. Social dilemmas and internet congestion. *Science*, 277(5325):535–537, 1997.

[9] R. M. Karp. Reducibility Among Combinatorial Problems. In R. E. Miller and J. W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, 1972.

[10] W. Kets. Learning with fixed rules: The minority game. *Journal of Economic Surveys*, 26(5):865–878, 2012.

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

[12] R. Savit, R. Manuca, and R. Riolo. Adaptive competition, market efficiency, and phase transitions. *Physical Review Letters*, 82(10):2203, 1999.

[13] R. Selten, T. Chmura, T. Pitz, S. Kube, and M. Schreckenberg. Commuters route choice behaviour. *Games and Economic Behavior*, 58(2):394–406, 2007.