

Emergent effects in multi-agent simulations of road pricing

Dominik Grether, Yu Chen, Marcel Rieser, Ulrike Beuck
and Kai Nagel

Transport Systems Planning and Transport Telematics
Technical University Berlin
D-10587 Berlin
+49-30-314 23308
+49-30-314 26269

30.06.2008

Abstract

Road pricing is debated as an option of transportation policy. Especially in metropolitan areas congestion pricing is promising to reduce congestion and to protect the environment. In order to reach the promised results the choice and design of a policy is very important, especially in a "second-best" context. Therefore it is worth to attempt detailed predictions of the effects and implications of the planned pricing scheme. Most if not all state-of-the-practice methodologies forecasting those effects are

- aggregate and in consequence do not consider social and economic characteristics of individual travelers.
- static in time and in consequence do not consider temporal effects such as toll avoidance

In order to bridge this gap, multi-agent microsimulations can be used. Our large-scale multi-agent traffic simulation is capable to simulate a complete day-plan of up to seven million individuals (agents). In contrast to other approaches, our simulation truly traces the synthetic travelers through their day, thus enabling us (at least in principle) to model emergent effects such as complex re-scheduling across the whole day.

This paper describes the implementation of a toll-scheme for the bigger Zurich area and presents the results of the simulation. We point out how agents (population) react to changed prices of transportation by modifying their consumption patterns. The analysis of the policy is based on the performance of simulated day-plans of the agents. This performance is directly given by a utility function, which is used to measure gains and losses of different groups of inhabitants in the research area. Based on these measurements we provide an economic interpretation of the policy and highlight emergent phenomena like changes in route choice and time reactions.

1 Introduction

Economic textbook theory states that optimal tolls should be based on marginal social cost. Yet, as Verhoef (2001) states, in the real world “all kinds of constraints may exist that prevent a regulator from charging the prices that she ideally would like to charge. Under such conditions, the regulator has to resort to ‘second-best’ pricing: setting the prices that are available optimally, under the constraints applying.” Examples for such constraints are (Verhoef, 2001):

- Optimal tolls may be highly variable both spatially and temporally; yet, this may be difficult to implement technologically, and it may be too complex to be useful as an incentive.
- To increase acceptability, a toll may (initially) only be introduced in a small area, and with a cap on the maximum amount.
- A neighboring jurisdiction may not do optimal pricing.
- There may be legal constraints (such as on the real time variability of the toll).
- There may be financial constraints (such as cost recovery for the tolling infrastructure).
- Marginal social cost pricing is theoretically only optimal if all other market, including, say, the labor market, are optimally priced. This is almost certainly not the case.
- Political bargaining may impose constraints.

In all these cases, the optimal solution under the constraints will be different from marginal social cost based pricing.

For those reasons, it is plausible to assume that second-best pricing will be the rule rather than the exception. Unfortunately, optimal second-best pricing is a lot more complicated than optimal first-best pricing. In some cases, it may be possible to make progress using analytical models (Verhoef, 2001). In many other, more complicated, situations, simulation may be the only applicable method. This paper will present a simulation approach where synthetic travelers are represented as “agents” that individually attempt to optimize their daily lives. The approach allows, in principle, arbitrary choice dimensions, such as choice of workplace or choice of residency. This paper will concentrate on route choice, time choice, and mode choice as the only choice dimensions. In contrast to most other work (e.g. de Palma and Marchal, 2002; Ettema et al., 2003), this paper will consider full daily plans. Consequences of this are, for example, that travelers that chose a non-car mode in the morning cannot chose a car mode to travel back, or travelers that start work late in the morning probably will also return later in the evening. In the same vein, the chaining of activities (such as home–work–shop–leisure–home) may put the whole activity chain under pressure when opening times are reduced.

This paper will look at the effect of an *afternoon* toll. This particular toll is not selected for optimality, but in order to demonstrate the emergent capabilities of the simulation: Not only will the toll in the afternoon change the afternoon travel patterns, but it will also change the morning travel patterns. This demonstrates not only that the choices of the synthetic travelers are connected along the time axis, but it also shows that the synthetic travelers are able to *look ahead*: Some of them switch to public transit in the *morning*, knowing that they would pay a car toll in the *afternoon*. This is achieved by learning iterations, i.e. agents repeatedly go through the same day over and over again, incrementally optimizing their daily plan under the constraints that they find, including emergent congestion effects.

The approach described in this paper is meant to evaluate/appraise approaches that are designed by other means. There are other approaches, such as by Markose et al. (2007) or Zhang and Levinson (2007), which concentrate on finding optimal designs, using agent-based approaches on the supply side. The emphasis of the present paper, in contrast, is on the travelers' behavioural realism, which we intend to improve much further in the near future. Ultimately, however, it is probable that the two approaches, agents on the supply-side, and improved behavioural realism on the demand side, will merge.

The paper first describes the simulation structure (Sec. 2). This is followed by a section on the scenario, which is a realistic simulation of regular daily workday traffic in the Zurich metropolitan area (Sec. 3). Sec. 4 introduces the toll scheme and its results, including an interpretation in monetary terms. The paper ends with a conclusion.

2 Simulation Structure

The following describes the structure of the simulation that is used. It is essentially the standard structure of MATSim, as described at many places (e.g. Balmer et al., 2005; Rieser et al., 2007). It differs, however, in some new elements concerning mode choice and minimum activity durations.

2.1 Overview

Our simulation is constructed around the notion of agents that make independent decisions about their actions. Each traveler of the real system is modeled as an individual agent in our simulation. The overall approach consists of three important pieces:

- Each agent independently generates a so-called *plan*, which encodes its intentions during a certain time period, typically a day.
- All agents' plans are simultaneously executed in the simulation of the physical system. This is also called the *traffic flow simulation* or *mobility simulation*.

- There is a mechanism that allows agents to *learn*. In our implementation, the system iterates between plans generation and traffic flow simulation. The system remembers several plans per agent, and scores the performance of each plan. Agents normally chose the plan with the highest score, sometimes re-evaluate plans with bad scores, and sometimes obtain new plans by modifying copies of existing plans.

A **plan** contains the itinerary of activities the agent wants to perform during the day, plus the intervening trip legs the agent must take to travel between activities. An agent’s plan details the order, type, location, duration and other time constraints of each activity, and the mode, route and expected departure and travel times of each leg. Two modes are considered car and non-car. The former is used for all plans using car while the latter covers all plans where travelers go by public transit, bike, foot, etc.

The task of generating a plan is divided into sets of decisions, and each set is assigned to a separate **module**. An agent strings together calls to various modules in order to build up a complete plan. To support this “stringing”, the input to a given module is a (possibly incomplete) plan, and the output is a plan with some of the decisions updated. This paper will make use of two constructive modules only: “activity times generator” and “router”. Mode choice will be simulated by giving each agent *two* initial plans, one using the car mode, and the other one using the non-car mode. During the iterations, the simulation makes sure that at least one car plan and one non-car plan are retained for every agent.

Once the agent’s plan has been constructed, it can be fed into the **traffic flow simulation**. This module executes all agents’ plans simultaneously on the network, allowing agents to interact with one another, and provides output describing what happened to the agents during the execution of their plans.

The outcome of the traffic flow simulation (e.g. congestion) depends on the planning decisions made by the decision-making modules. However, those modules can base their decisions on the output of the traffic flow simulation (e.g. knowledge of congestion). This creates an interdependency (“chicken and egg”) problem between the decision-making modules and the traffic flow simulation. To solve this, **feedback** is introduced into the multi-agent simulation structure (Kaufman et al., 1991; Bottom, 2000). This sets up an iteration cycle which runs the traffic flow simulation with specific plans for the agents, then uses the planning modules to update the plans; these changed plans are again fed into the traffic flow simulation, etc, until consistency between modules is reached.

The feedback cycle is controlled by the **agent database**, which also keeps track of multiple plans generated by each agent, allowing agents to reuse those plans at will. The repetition of the iteration cycle coupled with the agent database enables the agents to learn how to improve their plans over many iterations.

In the following sections we describe the used modules in more detail.

2.2 Activity Time Allocation Module

This module is called to change the timing of an agent’s plan. At this point, a very simple approach is used which just applies a random “mutation” to the duration and end time attributes of the agent’s activities. For each such attribute of each activity in an agent’s plan, this module picks a random time from the uniform distribution $[-30 \text{ min}, +30 \text{ min}]$ and adds it to the attribute. Any negative duration is reset to zero; any activity end time after midnight is reset to midnight.

Although this approach is not very sophisticated, it is sufficient in order to obtain useful results. This is consistent with our overall assumption that, to a certain extent, simple modules can be used in conjunction with a large number of learning iterations (e.g. Nagel et al., 2004). Since each module is implemented as a “plugin”, this module can be replaced by a more enhanced implementation if desired.

MATSim contains already a more sophisticated activity scheduling module (Meister et al., 2006). After appropriate testing, this will be used in future studies.

2.3 Router

The router is implemented as a *time dependent Dijkstra algorithm*. It calculates link travel times from the events output of the previous traffic flow simulation (see next section). The link travel times are encoded in 15 minute time bins, so they can be used as the weights of the links in the network graph. Apart from relatively small and essential technical details, the implementation of such an algorithm is straightforward (Jacob et al., 1999; Lefebvre and Balmer, 2007). With this and the knowledge about activity chains, it computes the fastest path from each activity to the next one in the sequence as a function of departure time.

2.4 Traffic Flow Simulation

The traffic flow simulation simulates the physical world. It differentiates between non-car plans and car-plans. The duration of every non-car trip is assumed to take twice as long as the car mode at free speed. This is based on the (informally stated) goal of the Berlin public transit company to generally achieve door-to-door travel times that are no longer than twice as long as car travel times. This, in turn, is based on the observation that non-captive travelers can be recruited into public transit when it is faster than this benchmark (Reinhold, 2006). For the purposes of the present paper, it is assumed that all non-car modes very roughly have the shared characteristics that they are slower than the (non-congested) car mode—this will be further disaggregated in future work. In the same vein, both for car and for non-car trips there are no separate considerations of access and egress.

Car plans are simulated by a queue simulation, which means that each street (link) is represented as a FIFO (first-in first-out) queue with two restrictions (Gawron, 1998; Cetin et al., 2003). First, each agent has to remain for a certain time on the link, corresponding to the free speed travel time. Second, a link storage capacity is defined which limits the number of agents on the link. If it is filled up, no more agents can enter this link.

Even though this structure is indeed very simple, it produces traffic as expected and it can run directly off the data typically available for transportation planning purposes. On the other hand, there are some limitations compared to reality, e.g. the number of lanes, weaving lanes, turn connectivities across intersections or signal schedules cannot be included into this model.

The output that the traffic flow simulation produces is a list of events for each agent, such as entering/leaving link, left/arrived at activity, and so on. Data for an event includes which agent experienced it, what happened, what time it happened, and where (link/node) the event occurred. With this data it is easy to produce different kinds of information and indicators like link travel time (which will be used by the router), trip travel time, trip length, percentage of congestion, and so on.

2.5 Agent Database – Feedback

As mentioned above, the feedback mechanism is important for making the modules consistent with one another, and for enabling agents to learn how to improve their plans. In order to achieve this improvement, agents need to be able to try out different plans and to tell when one plan is “better” than another. The iteration cycle of the feedback mechanism allows agents to try out multiple plans. To compare plans, the agents assigns each plan a “utility” based on how it performed in the traffic flow simulation.

Our framework always uses *actual plans performance* for the utility. This is in contrast to all other similar approaches that we are aware of. These other approaches feed back some aggregated quantity such as link travel times and reconstruct performance based on those (e.g. URBANSIM [www page](#), accessed 2007; Ettema et al., 2003). Because of unavoidable aggregation errors, such an approach can fail rather badly, in the sense that the performance information derived from the aggregated information may be rather different from the performance that the agent in fact experienced (Raney and Nagel, 2004).

The procedure of the feedback and learning mechanism is described in detail by Balmer et al. (2005). For better understanding, the key points are restated here.

1. The agent database starts with one complete plan per agent, which is marked as “selected”.
2. The simulation executes these marked plans simultaneously and outputs events.

3. Each agent uses the events to calculate the score of its “selected” plan, and decides which plan to select for execution by the next traffic flow simulation. When choosing a plan, the agent database can either:
 - create a new plan by sending an existing plan to the router, adding the modified plan as a new plan and selecting it,
 - create a new plan by sending an existing plan to the time allocation module, adding the modified plan and selecting it,
 - pick an existing plan from memory, choosing according to probabilities based on the scores of the plans. The probabilities are of the form $p = e^{\beta U_j} / \sum_i e^{\beta U_i}$, where U_j is the score (utility) of plan j , and β is an empirical constant. This is the familiar logit model (e.g. Ben-Akiva and Lerman, 1985).
4. Next, the simulation executes the newly selected plans, that is, it goes back to 2.

This circle continues until the system has reached a relaxed state. At this point, there is no quantitative measure of when the system is “relaxed”; we just allow the cycle to continue until the outcome seems stable. In the long run, measures to stop the iterations automatically and consistently should be found.

2.6 Scores (= utilities) for plans

In order for adaptation to work in a meaningful way, it is necessary to be able to compare the performance of different plans. This is easiest achieved by assigning scores to plans. This is the same as the fitness function in genetic algorithms, or the objective function in optimization problems. Note once more that every agent has its own scoring function, and attempts to optimize for her-/himself.

In principle, arbitrary scoring schemes can be used (e.g. prospect theory Avineri and Prashker, 2003). In this work, a utility-based approach is used, and therefore the word “score” is replaced by “utility” in the following. The approach is related to the Vickrey bottleneck model (Arnott et al., 1990), but is modified in order to be consistent with our approach based on complete daily plans (Charypar and Nagel, 2005; Raney and Nagel, 2006). The elements of our approach are as follows:

- The total utility of a plan is computed as the sum of individual contributions:

$$U_{total} = \sum_{i=1}^n U_{perf,i} + \sum_{i=1}^n U_{early.dp,i} + \sum_{i=1}^n U_{travel,i} ,$$

where U_{total} is the total utility for a given plan; n is the number of activities, which equals the number of trips; $U_{perf,i}$ is the (positive) utility earned for performing activity i ; $U_{early.dp,i}$ is the (negative) utility earned for leaving activity i too early; and $U_{travel,i}$ is the (negative) utility earned for traveling during trip i . In order to work in plausible real-world units, utilities are measured in Euro.

- A logarithmic form is used for the positive utility earned by performing an activity:

$$U_{perf,i}(t_{perf,i}) = \beta_{perf} \cdot t_{*,i} \cdot \ln \left(\frac{t_{perf,i}}{t_{0,i}} \right)$$

where t_{perf} is the actual performed duration of the activity, t_* is the “typical” duration of an activity, and β_{perf} is the marginal utility of an activity at its typical duration. β_{perf} is the same for all activities, since in equilibrium all activities at their typical duration need to have the same marginal utility.

$t_{0,i}$ is a scaling parameter that is related both to the minimum duration and to the importance of an activity. If the actual duration falls below $t_{0,i}$, then the utility contribution of the activity becomes negative, implying that the agent should rather completely drop that activity. A $t_{0,i}$ only slightly less than $t_{*,i}$ means that the marginal utility of activity i rapidly increases with decreasing $t_{perf,i}$, implying that the agent should rather cut short other activities. This paper uses

$$t_{0,i} = t_{*,i} \cdot \exp(-\zeta/t_{*,i}) .$$

where ζ is a scaling constant set to 10 hours. With this specific form, $U_{perf,i}(t_{*,i}) = \beta_{perf} \cdot \zeta$ is independent of the activity type.¹

Performing an activity only accumulates utility during the opening hours of the facility where the activity is performed. For example, if an agent’s work facility opens at 08:00 am and the agent arrives already at 07:45 pm, no utility is generated for the first 15 minutes of “being at the workplace”. Thus the agent will try to arrive more punctually to avoid the opportunity costs of waiting.

- The (dis)utility of leaving early is uniformly assumed as:

$$U_{early.dp,i} = \beta_{early.dp} \cdot t_{early.dp,i} ,$$

where $\beta_{early.dp}$ is the marginal utility (in Euro/h) for leaving an activity too early, and $t_{early.dp,i}$ is the number of hours an activity is performed shorter than its “typical” duration $t_{*,i}$

- To include alternative modes the (dis)utility of traveling is dependent on the mode. Our simple approach to do this is to use different valuations of the time for the two modes:

$$U_{travel,mode,i} = \begin{cases} \beta_{car} \cdot t_{travel,i} & \text{if trip } i \text{ is by car} \\ \beta_{non-car} \cdot t_{travel,i} & \text{if trip } i \text{ is not by car} \end{cases}$$

where β_{car} and $\beta_{non-car}$ are the marginal utilities of traveling by car or not by car (in Euro/h), respectively and $t_{travel,i}$ is the number of hours spent traveling during trip i . Clearly, for the time being this leaves out all more complicated aspects of non-car travel valuations, such as changing, schedule restrictions, waiting times, etc.

¹This “consequence” is actually the motivation for the specific mathematical form of the activity performance utility contribution. The reason for this motivation is not relevant to this paper, but is described in Charypar and Nagel (2005).

In principle, arriving early could also be punished. There is, however, no immediate need to punish early arrival, since waiting times are already indirectly punished by foregoing the reward that could be accumulated by doing an activity instead (opportunity cost). In consequence, the effective (dis)utility of waiting is already $-\beta_{perf}$.

Similarly, that opportunity cost has to be added to the time spent traveling, arriving at an effective (dis)utility of traveling of $-|\beta_{travel}| - \beta_{perf}$ where $\beta_{travel} \in \{\beta_{car}, \beta_{non-car}\}$.

Late arrivals are not punished by the above utility function. There is, however, a penalty for staying shorter than the typical duration, $t_{*,i}$. Together with land use constraints (workplaces and shopping facilities only open until 8pm, leisure facilities only open until midnight, see Tab.1), this can still lead to a pressure to arrive early. Assume, for example, a plan of “home–work–shop–shop–home”, and typical durations of 8h for work and 2h for shopping. Even without travel, any work starting time after 8am would lead to incurring the penalty for leaving early at some point. Therefore, the “penalty for leaving early” is a replacement for the more typical “penalty for arriving late”. This was done since our initial demands contains rather complex activity chains (such as, say, “home–work–work–leisure (=lunch)–work–work–shop–leisure–home”), and it is rather difficult to see which of these activities should incur a late penalty, and which not. – No opportunity cost needs to be added to the penalty for leaving early, since the time that is left early is spent at some other activity.

The values (β_{perf} , $\beta_{perf} + |\beta_{travel}|$, and $|\beta_{early.dp}|$) are the values that would correspond to the consensus values of the parameters of the Vickrey model (Arnott et al., 1990) if MATSim would just look for late arrival.

3 Scenario

3.1 Network

The scenario covers the metropolitan area of Zurich, Switzerland, which has about 1m inhabitants. It is shown in Fig. 1a. The network is a Swiss regional planning network, which includes the major European transit corridors. It consists of 24 180 nodes and 60 492 links. The original links have attributes (flow capacity, free speed, number of lanes, ...) suitable for static traffic assignment, but not for our dynamic agent-based simulation. This led to bad results in the simulation especially within the city of Zurich with its dense road network. Thus, all links within a circle with radius 4 kilometers around the center of Zurich had their attributes modified as follows:

- links corresponding to primary roads in OpenStreetMap² get a capacity of at least 2000 vehicles per hour. If the original capacity is higher than that, the capacity is not changed.

²see <http://www.openstreetmap.org>

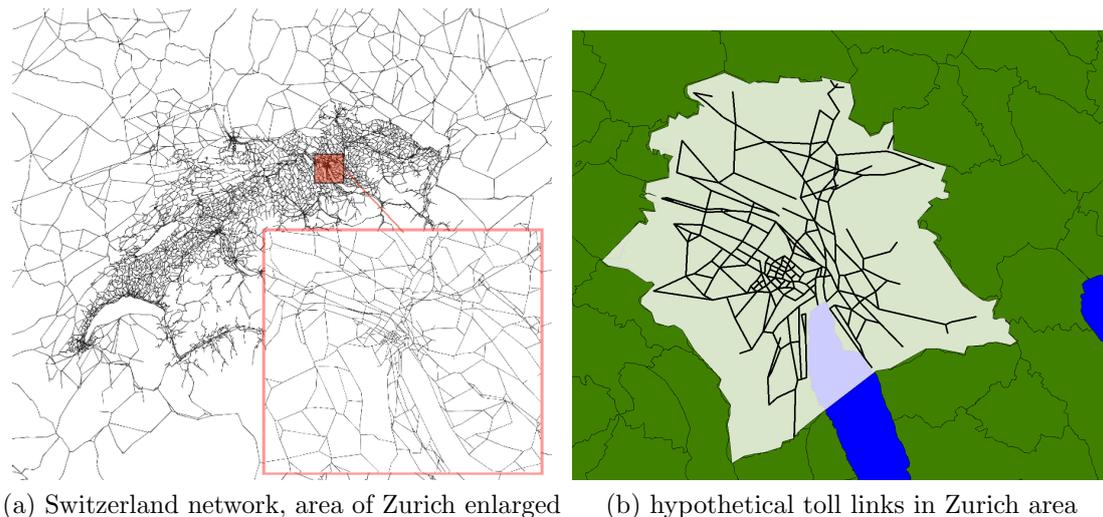


Figure 1: Scenario: Switzerland network with toll links for Zurich.

- links corresponding to secondary roads in OpenStreetMap keep their original capacity (usually between 1000 and 2000 veh/h).
- all other links get a capacity of at most 600 veh/h. If the original capacity is lower, it is not changed.
- a few single links are manually adjusted based on local knowledge.

3.2 Initial demand

The simulated demand consists of all travelers within Switzerland that cross at least once during their day an imaginary boundary around Zurich. This boundary is defined as a circle with a radius of 30 kilometers and with its center at “Bellevue”, a central place in the city of Zurich. To speed up computations, a random 10% sample was chosen for simulation, consisting of 181 725 agents.

The travelers have complete daily activity patterns based on microcensus information. Such activity patterns can include activities of type *home*, *work*, *education*, *shopping*, and *leisure*. The typical durations for those activities are derived from the microcensus data and are specified individually for each member of the synthetic population. Each agent gets two plans based on the same activity pattern. The first plan uses only “car” as transportation mode, while the second plan uses only “non-car”.

This demand was then extended with people crossing the borders of Switzerland and travelling within the region of Zurich, either because they live in neighboring countries but work in Switzerland, because they live in Switzerland but work outside, or because

they travel through Switzerland on transit. Again, a 10% sample was taken, adding 5759 agents to the demand. This part of the demand is important to get more realistic traffic volumes especially on highways. These agents of our population have no option to switch from mode car to non-car. We will refer to them as “transit” traffic in the following paragraphs.

To specify opening and closing times for the facilities where activities are performed, activities are classified by type, i.e. it is distinguished between home, work, education, shop and leisure activity types. Opening and closing times for the facilities where those types are performed are shown in Table 1.

Activity type	Opening time	Closing time
Home	00:00	24:00
Work	06:00	20:00
Education	06:00	20:00
Shop	08:00	20:00
Leisure	00:00	24:00

Table 1: Activity opening and closing times used in the scenario.

3.3 Simulation runs

The simulation is run for 500 iterations to retrieve a relaxed state in which the initial plans are adapted to the traffic conditions. In each iteration 10% of the agents adapt routes and 10% adapt activity times. This is done until iteration 500 is reached. Based on the 500th iteration we run two simulations with further 300 iterations. The first simulation is run with the same settings and will be referred as base case, i.e. the situation in that no policy is implemented. The second run implements the policy, i.e. each iteration 10% of the agents adapt routes and 10% times. With a probability of 80% agents select one of the existing plans as described in Sec. 2.5. In doing so they can choose between modes car or non-car. This run will be referred as city toll. The following values for the parameters were used:

Parameter	Value
β_{perf}	6 Euro/h
$\beta_{early.dp}$	-18 Euro/h
β_{car}	-6 Euro/h
$\beta_{non-car}$	-3 Euro/h
β (existing plans)	4

Table 2: Behavioral parameters used in the scenario.

Although it is not obvious at a first glance, these values mirror the standard values of the Vickrey scenario (e.g. Arnott et al. (1993)): An agent that arrives early to an activity must wait for the activity to start, therefore forgoing the $\beta_{perf} = +6$ Euro/hour that it could accumulate instead. An agent that travels by car forgoes the same amount, plus a loss of 6 Euro/hour for traveling. Finally, an agent that leaves an activity before he has performed this activity at least for its typical duration receives a penalty of 18 Euro/hour.

With this setup the simulation is run with $\beta_{non-car} = -3$ Euro/hour that results in a car usage rate of $\approx 42\%$ that is around the real mode distribution of the Zurich area.

For the Zurich region data from 159 traffic counting stations is available. The hourly measured traffic volumes can be compared with the amount of traffic of the base case scenario simulation runs. This comparison is shown in Fig. 2. Most important is the red curve which is calculated for each hour by following formula:

$$\frac{\text{Simulated traffic volume} - \text{Real traffic volume}}{\text{Real traffic volume}} * 100$$

During the night, i.e. from 00:00 am till 06:00 am, the simulation deviates from reality with 70% to 130%. However the simulation results for the daytime, i.e. from 06:00 am till midnight, have a relative deviation of about 30%.

4 Results

4.1 Toll

For this scenario a “city toll” is implemented which is a distance-based charge for a certain area of Zurich. The toll area covers the Zurich city area of administration, but not the motorways that lead into and partially around the city. The exclusion of the motorways is plausible because they are owned by the Swiss Confederation and not by the city of Zurich. This is a plausible scenario based on the status of the political discussion in Switzerland (Bundesrat (Government) of Switzerland, 2007). We are not aware of more specific politically discussed scenarios for Zurich; otherwise we would have used those. Fig. 1b shows the area and the tolled links. The diameter of the toll area is about 11 km. The toll is levied during evening rush hour and is set to 2 Euro/km. The toll starts exactly at 03:00 pm in the afternoon and lasts till 07:00 pm in the evening. The tolled area has a high density of offices and other work places, so the in-bound traffic is larger in the morning than the out-bound traffic, and vice versa in the evening.

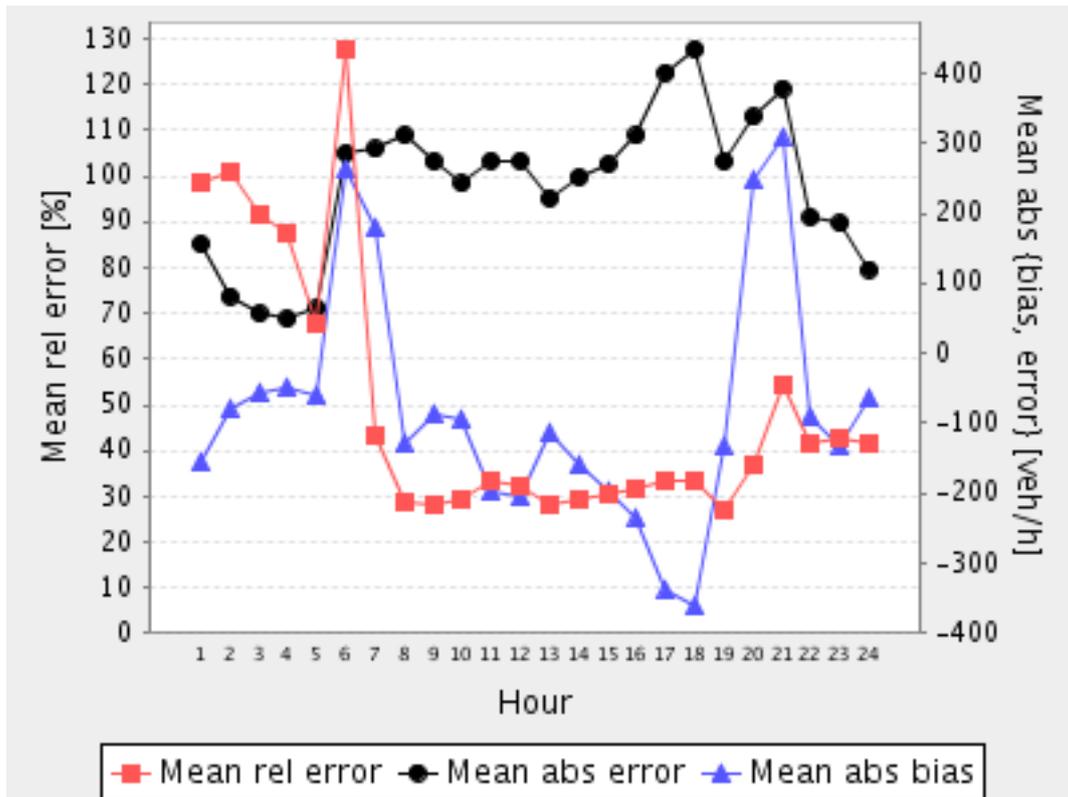


Figure 2: Realism of the base case. 159 traffic counting stations provide real traffic counts for the Zurich area. The three curves show mean relative error (red), mean absolute error (black) and the mean absolute bias (blue) when comparing the traffic volumes of the base case with the real values.

4.2 Traffic results

This paragraph will analyze effects of the toll on the traffic in the network. Thus, it will be shown which modifications in travel behavior are evoked when the policy is implemented in the Zurich scenario.

The results of the simulation runs of the Zurich scenario are summarized in Fig. 3. The first and second column contain the data for the base case and the city toll scenario, respectively. The third column contains the difference between the city toll run and the base case values.

The first line contains the number of agents used for simulation. The next line contains an indicator for the performance of the system as a whole. The average duration of trips per agent is 737 seconds in the base case scenario. Due to the toll the average time spent for each trip increments by 11 seconds. This suggests that the toll has a negative effect for travelers. However if we look at the next three lines, that differ between car

	Base case	City toll	Difference
Size of population	187,484	187,484	0
Trip duration [avg. per trip, s]	737	748	11
Car trip duration [avg. per trip, s]	708	699	-8
Non-car trip duration [avg. per trip, s]	540	572	33
Transit trip duration [avg. per trip, s]	7266	7285	19
Car rate [%]	43.49	40.74	-2.75
non-car rate [%]	53.44	56.19	2.75
Transit rate [%]	3.07	3.07	0

Figure 3: Results for the Zurich scenarios. The third column displays the difference, i.e. values of the toll scenario are subtracted from the base case values.

trips, non-car trips and transit traffic, the effects of the toll can be analyzed in a more differentiated way. Due to the toll the average trip duration for cars is reduced by 8 seconds from 708 seconds to 699 seconds. Therefore the average duration of non-car trips gets longer by 33 seconds, while the trip durations for the transit also increases by 19 seconds.

To explain those effects it is worth taking a look at the last three lines of Fig. 3. Due to the toll the percentage of car travelers is reduced from 43.49% to 40.74%. Accordingly the number of non-car travelers increases by 2.75%. The percentage of transit traffic is not varied because the mode-choice option is switched off for those travelers.

The toll increases costs of car travel, so the stimulus to switch to non-car modes is stronger in the policy scenario. Following the stimulus, agents switch to non-car modes thus the agents that still travel by car suffer less congestion. However by excluding the highways from the tolled area car traffic is more likely to prefer the highways. As a lot of the transit traffic uses the highways travel times for transit get worse.

Looking at the effects of the toll one could ask who are the agents that switch their mode. Fig. 4. depicts the home locations of the agents affected by the toll, i.e. those that have at least one activity in the toll area. Dots are colorized by the mode change due to the toll. Agents living at red or green dots stay at modes car or non-car, respectively. Yellow dots symbolize home locations of agents switching from car to non-car modes, while blue dots depict homes of the agents changing from non-car to car mode.

Looking at Fig. 4 one suggests that agents living closer to the metropolitan area are more stimulated by the toll to change mode from car to non-car. Residents living quite at the borders of the metropolitan tend to choose mode car and to keep this mode even if the toll is implemented. Their travel times to the city center are long thus they have a

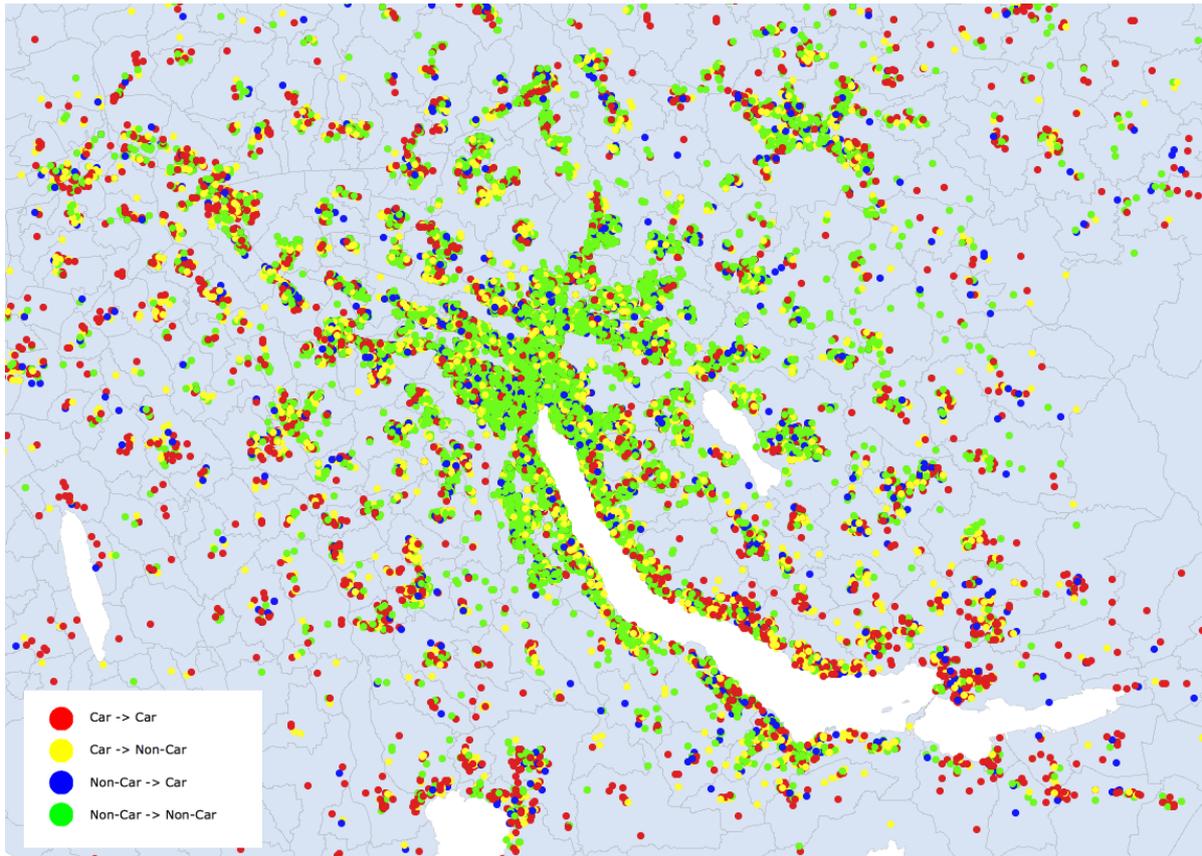


Figure 4: Map of the greater Zurich area. Dots symbolize home locations of agents that are affected by the toll, i.e. those that have at least one activity in the toll area. Dots are colorized by the mode change due to the toll. Note that agents living at the border of the metropolitan area are most likely forced to switch from non-car to car (blue dots).

stronger incentive to travel car mode than non-car mode even when prices for mode car are rised due to the policy.

Mode changes however are not the only reaction to the toll. Fig. 5 shows the average speed in the toll area for the base case and the city toll scenario. Average speeds are calculated per hour. Despite the toll is only implemented between 03:00 pm and 07:00 pm average speeds increase in the tolled area during the whole day. The best increase of speed is when the toll is implemented. However between 07:00 pm and 08:00 pm average speed falls below the value of the average speed in the base case. The reason is a time reaction of agents to the toll.

Fig. 6 shows the number of travelers simultaneously on the road for the two scenarios. The figure subdivides the travelers by the mode, transit traffic is not included in the diagram. The area below the curves can be interpreted as the total time agents spend on the road. One can clearly see the reduction of car traffic and the increase of non-car

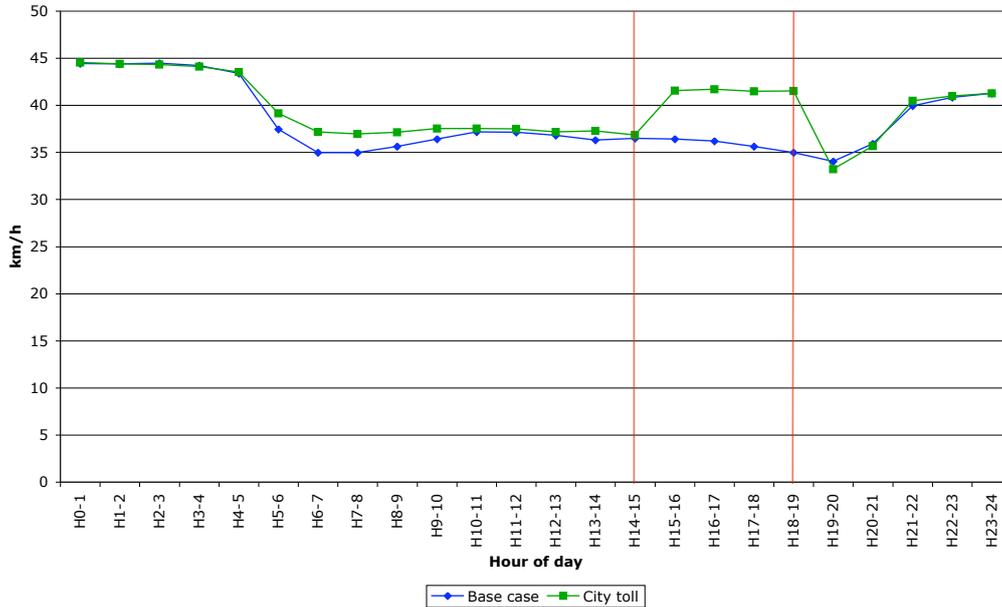


Figure 5: Average speed in the city area of Zurich for the base and policy case scenario.

traffic during the tolled time-frame. At 03:00 pm, when the toll starts, car traffic goes back quickly. At 07:00 pm, when the toll period ends, one can see the strong rise of car traffic. This explains the lower average speed in the tolled area that can be seen in Fig. 5. People try to avoid the toll and arrange their activity schedule in a way that they are able to depart out of the tolled area directly after the end of the policy. If they are not able to avoid the toll they tend to switch to mode non-car. Note that the mode changes are not only present during the tolled time frame. Also during the morning peak we can observe a rise of non-car travelers and a lower number of car users in the city toll scenario.

4.3 Winners and losers

Standard economic appraisal, as is for example used for cost-benefit-analysis (e.g. Pearce and Nash, 1981), would now take the above traffic patterns as input, and attach economic valuations to them. It should be clear, however, that this is too simplistic for the above scenario. In particular, the standard approach only counts travel time gains, but not schedule delay effects (caused by people choosing a non-preferred time of travel to avoid

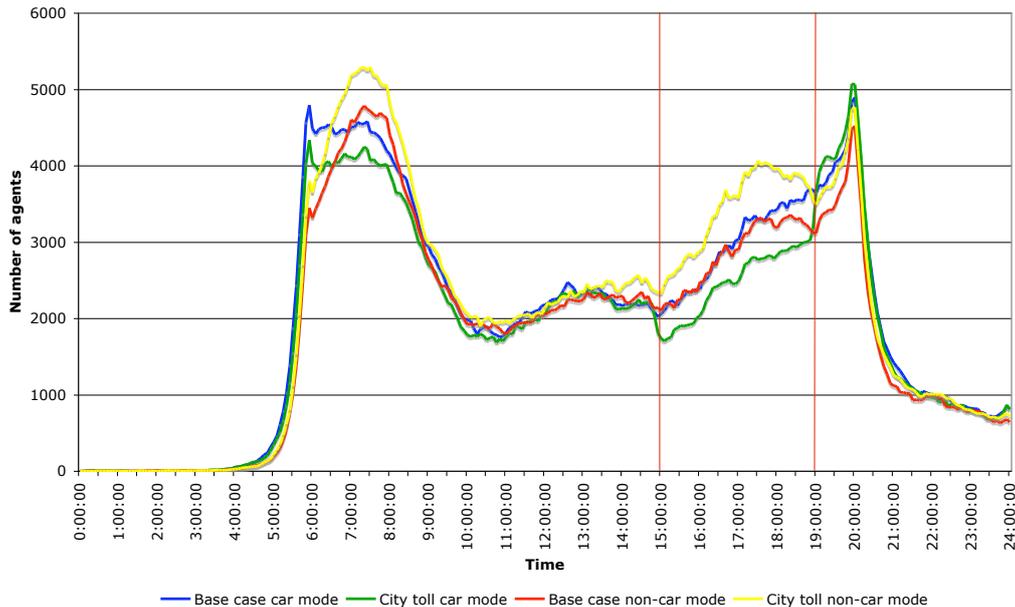


Figure 6: Number of travelers on the road over the time of day. The red bars mark the toll period.

the toll). However, research indicates that schedule delay effects might contribute more than half of the economic effects of well-chosen time-dependent tolls (Arnott et al., 1990).

Fortunately, with our multi-agent simulation, it is, in fact, not necessary to add the economic appraisal as it is conventionally done, since the agent utilities already are the economic performance indicators of the system. This is because the utility is the measure that every agent attempts to improve, and a higher utility directly measures the amount of improvement that an agent was able to reach. This automatically includes the schedule delay effects, since every agent will have optimally adjusted to any trade-off between time-dependent congestion, time-dependent toll, and schedule delay, including any personal restrictions that an agent may have, such as specific opening times. The multi-agent approach could also include different values of time, since they would be included as person-specific values of β_{perf} , β_{travel} , and β_{late} .

It is, thus, immediately possible to identify winners and losers of a policy. An example of such an analysis is Fig. 7, which allocates the gains and losses after the introduction

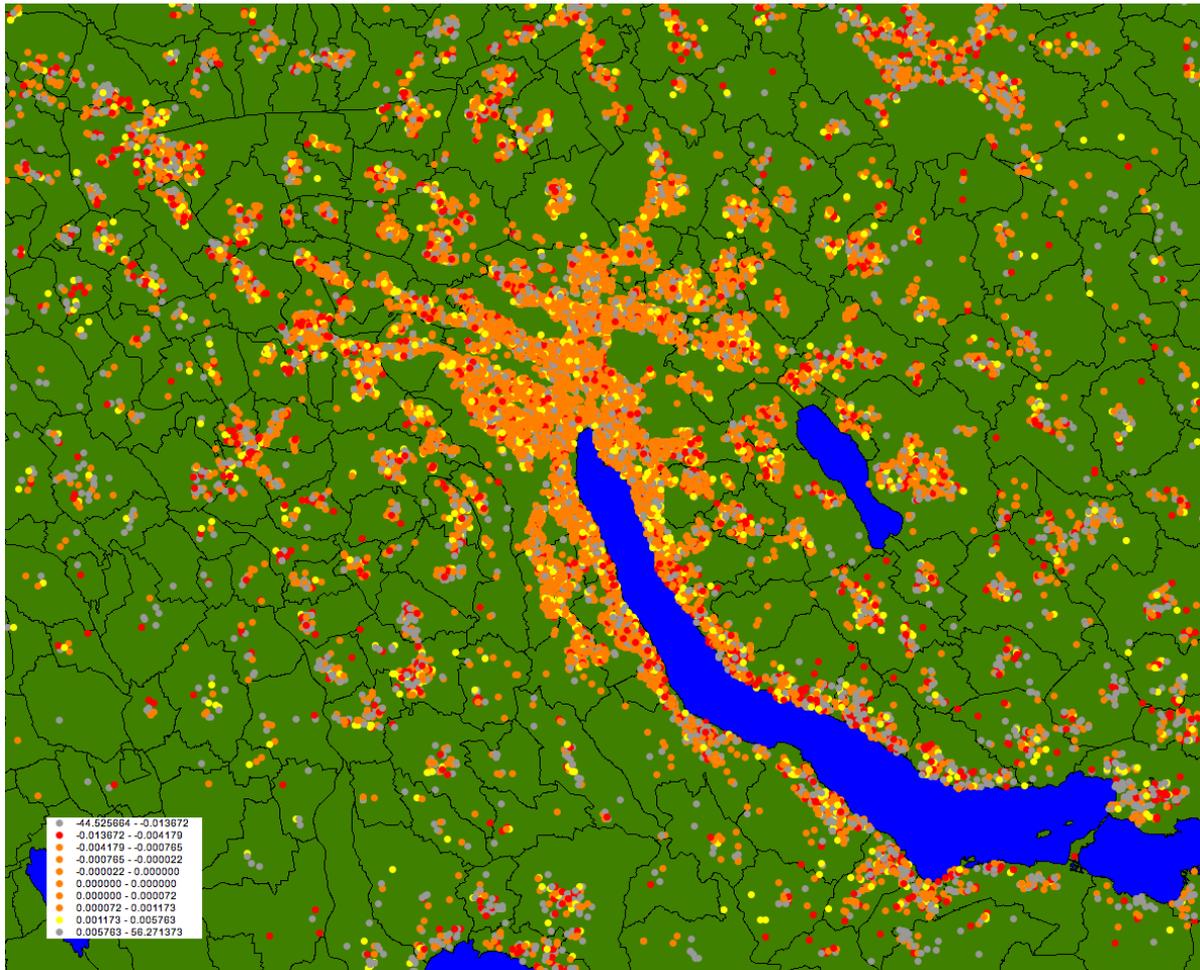


Figure 7: Map of the greater Zurich area. Each dot locates a home location of an agent. The relative deviation of the base and policy case scores is partitioned in decils.

of the city toll to the residential locations of the agents. The analysis is done without redistribution of the toll revenues. Each dot symbolizes the home location of an agent that is affected by the toll, i.e. that has at least one activity in the toll area. The colors reflect a classification based on the relative deviation score comparing the city toll scenario with the base case. The relative deviation is partitioned in decils. The grey dots depict the borders of the distribution where the score deviation is quite high. Recall that in each iteration 10% of the agents try to find better routes and another 10% tries to improve their schedule. If the performance of the new plans turns out to be bad the deviation gets high. One could consider this cases as outliers, but for completeness of data they are included in the study. Red dots depict the agents in the second decil, i.e. that lose 1.36 to 0.41%. Yellow dots symbolize agents in the ninth decil which gain 0.11 to 0.57%. All other decils are colored orange, those dots stand for the agents that are more or less indifferent to the policy.

In this case, there are both winners and losers close to Zurich, and mostly losers further away (red dots). The reason is that the non-car mode is attractive for short trips and becomes less attractive the longer the trips are. A price rise for car trips thus leads to a mode reaction that stimulates the short trips more to use the non-car mode than the long trips. As all trips are charged equally the agents with long trips and no possibility to avoid the toll by using non-car mode will have the highest loss. Apart from this, there does not seem to be any structure, implying is that any structure that there may be is not geographically oriented.

4.4 Aggregated benefits

For economic analysis, the gains and losses need to be aggregated. The simplest — utilitarian — way to do this is to sum up the utilities. The reasoning behind this would be that the monetarized utilities reflect willingness to pay of each individual for the change of the system by the policy (including the losers, which would need to receive money to accept), and thus the sum of these utilities reflects the aggregated willingness to pay. Table 4.4 shows one such analysis, for the scenario described above. On the left, one finds different entries for the base case. This is followed by two columns for the city toll, one depicting the new numbers after the policy introduction, the other the differences to the base case. For example, when switching on the city toll the average utility increases by 0.35 Euro, including the utility loss from the toll payment. Since, however, 4359 agents pay toll and average toll payments are 4.40 Euro, after redistribution of the toll revenues there is an average utility gain of 0.45 Euro. Multiplied by the number of agents this results in 84,009 Euro of utility gains per day.

An advantage of the agent-based approach is that the simulation and the appraisal are automatically consistent. A similar consistency argument is made by de Jong et al. (2005), where it is suggested to use the utility functions estimated for discrete choice models (or more precisely the expectation value, commonly called the “logsum term”) directly for appraisal. This would have the same effect as our approach, i.e. that the model that is used for predicting the behavioral response is the same as the one that is used for evaluation/appraisal. The difference is that in our approach, the reactions of the travelers are directly and microscopically computed. This becomes clear when one, say, attempts to predict the effects of an area toll. In that situation, for the logsum approach it will be quite difficult to differentiate who will be affected by that toll and to what extent. One option would be to look at every person separately. Then, however, one ends up very close to a multi-agent simulation.

Finally, any aggregation by sub-groups is possible since the individual utilities are attached to every individual of the synthetic population, thus allowing filtering and aggregation by arbitrary criteria as was shown in Figures 4 and 7.

	Base case	City toll	Difference
Size of population	187,484	187,484	0
Utility [avg. per agent, Euro]	174.72	175.07	0.35
Utilities [sum/Euro]	32,757,715.99	32,822,561.74	64,845.75
Number of paying agents	0	4359	4359
Toll paid [avg. per paying agent, Euro]	0.00	4.40	4.40
Toll paid [sum, Euro]	0.00	19,164.07	19,164.07
Utility, after redistribution of toll [avg. per agent, Euro]	174.72	175.17	0.45
Utility, after redistribution of toll [sum/Euro]	32,757,715.99	32,841,725.81	84,009.82

Figure 8: Results for the Zurich scenarios. The third column displays the difference, i.e. values of the toll scenario are subtracted from the base case values.

5 Conclusion

It was argued in the introduction that it will not be possible to base real-world toll schemes on theoretically optimal marginal social cost pricing, since in reality many constraints exist that will render such an approach infeasible. The paper has demonstrated that multi-agent simulation is a viable option in such a situation. Starting from a realistic base case for the Zurich metropolitan area, route choice, mode choice, and time choice were included as choice dimensions in reaction to the toll, and these choices were not trip-based but based on complete daily activity plans.

This was demonstrated by investigating an *afternoon/evening* toll. The results demonstrate that the approach is capable of adjusting the full daily plans, most importantly to switch to public transit in the *morning* in order to avoid the car toll in the *afternoon/evening*. It was also demonstrated that, because every synthetic traveller optimizes according to an individual utility function, it is straightforward to identify winners and losers of a policy measure. These groups can then be analyzed with arbitrary methods, including subgroup aggregation, or spatial analysis.

Finally, unweighted addition of the monetarized individual utilities leads to the relevant utility term that is used in conventional cost-benefit-analysis. The advantage of the approach is, however, that the utilities that drive the decision-making in the model are directly available for the cost-benefit-analysis, i.e. no post-processing is necessary with respect to the internal part of the benefits.

6 Acknowledgments

This work was funded in part by the Volvo Research and Educational Foundations within the research project "Environmentally-oriented Road Pricing for Livable Cities" and by the "Bundesministerium für Bildung und Forschung" (BMBF) within the research project "Adaptive Verkehrssteuerung" (Advest).

References

- R. Arnott, A. de Palma, and R. Lindsey. Economics of a bottleneck. *Journal of Urban Economics*, 27(1):111–130, 1990.
- R. Arnott, A. De Palma, and R. Lindsey. A structural model of peak-period congestion: A traffic bottleneck with elastic demand. *The American Economic Review*, 83(1): 161–179, 1993.
- E. Avineri and J.N. Prashker. Sensitivity to uncertainty: Need for paradigm shift. *Transportation Research Record*, 1854:90–98, 2003.
- M. Balmer, B. Raney, and K. Nagel. Adjustment of activity timing and duration in an agent-based traffic flow simulation. In H.J.P. Timmermans, editor, *Progress in activity-based analysis*, pages 91–114. Elsevier, Oxford, UK, 2005.
- M. Ben-Akiva and S. R. Lerman. *Discrete choice analysis*. The MIT Press, Cambridge, MA, 1985.
- J.A. Bottom. *Consistent anticipatory route guidance*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, 2000.
- Bundesrat (Government) of Switzerland. Grünes Licht für Versuche mit Road Pricing. press conference, Dec/7 2007.
- N. Cetin, A. Burri, and K. Nagel. A large-scale agent-based traffic microsimulation based on queue model. In *Proceedings of Swiss Transport Research Conference (STRC)*, Monte Verita, CH, 2003. See www.strc.ch. Earlier version, with inferior performance values: Transportation Research Board Annual Meeting 2003 paper number 03-4272.
- D. Charypar and K. Nagel. Generating complete all-day activity plans with genetic algorithms. *Transportation*, 32(4):369–397, 2005.
- G. de Jong, A.J. Daly, M. Pieters, and A. van der Hoorn. The logsum as an evaluation measure: review of the literature and new results. In *presented at the 45th Congress of the European Regional Science Association*, Amsterdam, August 2005.
- A. de Palma and F. Marchal. Real case applications of the fully dynamic METROPOLIS tool-box: An advocacy for large-scale mesoscopic transportation systems. *Networks and Spatial Economics*, 2(4):347–369, 2002.
- D. Ettema, G. Tamminga, H. Timmermans, and T. Arentze. A micro-simulation model system of departure time and route choice under travel time uncertainty. In *Proceedings of the meeting of the International Association for Travel Behavior Research (IATBR)*, Lucerne, Switzerland, 2003. See www.ivt.baug.ethz.ch.
- C. Gawron. *Simulation-based traffic assignment*. PhD thesis, University of Cologne, Cologne, Germany, 1998. URL www.zaik.uni-koeln.de/AFS/publications/theses.html.
- R. R. Jacob, M. V. Marathe, and K. Nagel. A computational study of routing algorithms for realistic transportation networks. *ACM Journal of Experimental Algorithms*, 4 (1999es, Article No. 6), 1999.

- D.E. Kaufman, K.E. Wunderlich, and R.L. Smith. An iterative routing/assignment method for anticipatory real-time route guidance. Technical Report IVHS Technical Report 91-02, University of Michigan Department of Industrial and Operations Engineering, Ann Arbor MI 48109, May 1991.
- N. Lefebvre and M. Balmer. Fast shortest path computation in time-dependent traffic networks. Arbeitsbericht Verkehrs- und Raumplanung 439, IVT, ETH Zürich, Zürich, 2007.
- S. Markose, A. Alentorn, D. Koesrindartoto, P. Allen, P. Blythe, and S. Grosso. A smart market for passenger road transport (SMPRT) congestion: An application of computational mechanism design. *Journal of Economic Dynamics and Control*, 31(6): 2001–2032, 2007.
- K. Meister, M. Balmer, K.W. Axhausen, and K. Nagel. planomat: A comprehensive scheduler for a large-scale multi-agent transportation simulation. In *Proceedings of the meeting of the International Association for Travel Behavior Research (IATBR)*, Kyoto, Japan, 2006. See www.iatbr.org.
- K. Nagel, M. Strauss, and M. Shubik. The importance of timescales: Simple models for economic markets. *Physica A*, 340(4):668–677, 2004.
- D. Pearce and C.A. Nash. *Social appraisal of projects: A text in cost-benefit analysis*. Wiley & Sons, London, 1981.
- B. Raney and K. Nagel. Iterative route planning for large-scale modular transportation simulations. *Future Generation Computer Systems*, 20(7):1101–1118, 2004.
- B. Raney and K. Nagel. An improved framework for large-scale multi-agent simulations of travel behavior. In P. Rietveld, B. Jourquin, and K. Westin, editors, *Towards better performing European Transportation Systems*. Routledge, London, 2006. Similar version TRB preprint number 05-1846.
- T. Reinhold. Konzept zur integrierten Optimierung des Berliner Nahverkehrs. In *Öffentlicher Personennahverkehr*. Springer Berlin Heidelberg, 2006. doi: 10.1007/3-540-34209-5_8.
- M. Rieser, K. Nagel, U. Beuck, M. Balmer, and J. Rügenapp. Truly agent-oriented coupling of an activity-based demand generation with a multi-agent traffic simulation. *Transportation Research Record*, 2021:10–17, 2007. doi: 10.3141/2021-02.
- URBANSIM www page, accessed 2007. URL www.urbansim.org.
- E.T. Verhoef. Marginal cost based pricing in transport. key implementation issues from the economic perspective. See http://www.imprint-eu.org/public/EV_V2.PDF, 2001. Paper presented at Seminar One of IMPRINT-EUROPE.
- L. Zhang and D. Levinson. Agent-based model of price competition, capacity choice, and product differentiation on congested networks. In *Proceedings of the World Conference on Transport Research*, Berkeley, CA, 2007.