Income-Contingent User Preferences in Policy Evaluation – Application and Discussion based on Multi-Agent Transport Simulations

Benjamin Kickhöfer, Dominik Grether,
and Kai Nagel

Transport Systems Planning and Transport Telematics
Berlin Institute of Technology
D-10587 Berlin
+49-30-314 23308
+49-30-314 26269

March 25, 2011

Abstract

Standard economic policy evaluation allows the realisation of projects if the aggregated economic benefit overweights their costs. In democratically organized societies, the implementation of measures that have negative impacts on some part of the population tends to be complicated due to low public acceptance, even when only a minority is worse off. The microscopic multi-agent simulation approach presented in this paper is capable to help designing better solutions in such situations. In particular, it is shown that income can and needs to be included in utility calculations for a better understanding of problems linked to acceptability. This paper shows how multi-agent approaches can be used in policy evaluation when including income in the user preferences. Therefore, an income-dependent utility function is estimated based on survey data. Subsequently, using the MATSim framework, the implementation is tested in a test scenario. Furthermore, and going beyond Franklin (2006), it is shown that the approach works in a large-scale real world example. Based on a hypothetical speed increase of public transit, effects on the welfare distribution of the population are discussed. It is shown that the identification of winners and loosers seems to be quite robust when using individual utility functions. However, results indicate that a conversion or aggregation of individual utility changes for welfare analysis is highly dependent on the functional form of the utility functions as well as on the choice of the aggregation process.

1 Introduction

Policy measures in transportation planning aim at improving the system as a whole. In democratically organized societies, however, it is quite difficult to realize projects when they have a negative impact on some part of the population even if this is a minority – presumably, this has something to do with the fact that losses are weighted more than gains (Kahneman and Tversky, 1979). In addition, changes to the system that result in an unequal distribution of the overall welfare gain are, however, hard to implement in democratically organized societies. Studies indicate that, e.g., tolls tend to be regressive if no redistribution scheme is considered at the same time, and may so increase the inequality in welfare distribution (e.g. Franklin, 2006). An option to reach broader public acceptance for such policies may be to include the redistribution of total gains into the scheme. Hence, methods and tools are needed that simulate welfare changes due to policies on a highly granulated level, e.g. considering each individual of the society. With such tools, policy makers are able to consider impacts of different proposed measures on the welfare distribution. In addition, it is possible to estimate the support level within the society and, if necessary, to evaluate alternatives for further discussion.

Traditional transport planning tools using the four-step process combined with standard economic appraisal methods (e.g. Pearce and Nash, 1981) are not able to provide such analysis. In order to bridge this gap, multi-agent microsimulations can be used. Large-scale multi-agent traffic simulations are capable of simulating the complete daily plans of several millions of individuals (agents) (Meister et al., 2008). In contrast to traditional models, all attributes that are attached to the synthetic travelers are kept during the simulation process, thus enabling highly granulated analysis (Nagel et al., 2008). Being aware of all attributes enables the possibility to attach to every traveler an individual utility function that is used to maximize the individual return of travel choices during the simulation process. Another advantage of the multi-agent simulation technique is the connection of travelers' choices along the time axis when simulating time dependent policies (Grether et al., 2008).

In the context of policy evaluation, simulation results can immediately be used to identify winners and losers, since the utility of the individual agents are kept and can be compared between scenarios agent-by-agent. They can also be aggregated in arbitrary ways, based on available demographic attributes including spatial information of high resolution. Welfare computations, if desired, can be done on top of that, without having to resort to indirect measures such as link travel times or inter-zonal impedances. The usual problems when aggregating or monetizing the individual utility still apply (Bates, 2006).

This paper shows how multi-agent approaches can be used in policy evaluation. It studies why income should be included in utility calculations when considering issues linked with public acceptance. Then, it describes implications on the simulation model and focuses on the measurement of welfare effects resulting from a policy measure. Note that this paper is an extension of Grether et al. (2009b), who considered three policy measures: a public transit (pt) price increase, a pt speed increase, and a combination of the two. The results of the combined measure are also reported in Grether et al. (2010). In contrast to the latter, the present paper concentrates on the pt speed increase only. In contrast to both papers, the present paper provides more profound insights, particularily in the discussion.

2 Simulation Structure

The following describes the structure of the simulation that is used. It is the standard structure of MATSim[1], as described at many places (Raney and Nagel, 2006; Balmer et al., 2005). Readers familiar with the MATSim approach can skip this section.

2.1 Overview

In MATSim, each traveler of the real system is modeled as an individual agent. The overall approach consists of three important parts:

- Each agent independently generates a so-called *plan*, which encodes its preferences 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 the 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 choose 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 that 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.

A plan can be modified by various **modules**. In the test scenario, the Time Adaptation module is used, while the large-scale application additionally uses a Router module. The *Time Adaptation* module changes the timing of an agent's plan. A very simple approach is used which just applies a random "mutation" to the duration attributes of the agent's activities (Balmer et al., 2005). The *Router* is a time-dependent best path algorithm (Lefebvre and Balmer, 2007), normally using the link travel times of the previous iteration as the link's generalised costs. *Mode choice* will not be simulated by a module per se, but instead by making sure that every agent has at least one "car" and at least one "public transit" plan (Grether et al., 2009a; Rieser et al., 2009).

One plan of every agent is marked as "selected". The **traffic flow simulation** executes all agents' selected plans simultaneously on the network and provides output describing what happened to each individual agent during the execution of its plan. The *car traffic flow* simulation is implemented as a queue simulation, where each street (link) is represented as a 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. The *public transit simulation* simply assumes that travel by public transit takes twice as long as travel by car on the fastest route in an empty network (Grether et al., 2009a; Rieser et al., 2009), and that the travel distance is 1.5 times the beeline distance. Public transit is assumed to run continuously and without capacity restrictions.

The modules base their decisions on the output of the traffic flow simulation (e.g. knowledge of congestion) using **feedback** from 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

[1] Multi-Agent Transport Simulation, see www.matsim.org

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.

In every iteration, 10 % of the agents generate new plans by taking an existing plan, making a copy of it, and then modifying the copy with the Time Adaptation or the Router module. The other agents reuse one of their existing plans. The probability to change the selected plan is calculated according to

$$p_{change} = min(1, \alpha \cdot e^{\beta \cdot (V_{random} - V_{current})/2}) \ , \tag{1}$$

where α is the probability to change if both plans have the same score, set to 1%; β is a sensitivity parameter, set to 20 for the tests and to 2 for the large-scale Zurich simulations; and $V_{\{random,current\}}$ is the score[2] of the current/random plan (see later). In the steady state, this model is equivalent to the standard multinomial logit model

$$p_j = \frac{e^{\beta \cdot V_j}}{\sum_i e^{\beta \cdot V_i}} \ ,$$

where p_j is the probability for plan j to be selected. In consequence, V corresponds to the systematic component of utility in Random Utility Models (RUM) (e.g. Ben-Akiva and Lerman, 1985; Train, 2003), where utility is defined as $U = V + \epsilon$. In RUM, the ϵ is called random component of utility. In the steady state and assuming a Gumbel distribution for ϵ, the choice model used in this paper is thus equivalent to the standard multinomial logit model.

The repetition of the iteration cycle coupled with the agent database enables the agents to improve their plans over many iterations. As the number of plans is limited for every agent by memory constraints, the plan with the worst performance is deleted when a new plan is added to a person which already has the maximum number of plans permitted. If agents have several plan types in their memory, e.g. one plan using car and another using public transit mode only, at least one plan of each type is kept. The iteration cycle 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 is stable.

2.2 Scoring Plans

In order to compare plans, it is necessary to assign a quantitative score to the performance of each plan. In this work, in order to be consistent with economic appraisal, a simple utility-based approach is used. The elements of our approach are as follows:

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

$$V_{total} = \sum_{i=1}^{n} \left(V_{perf,i} + V_{late,i} + V_{tr,i} \right) \ , \tag{2}$$

  where V_{total} is the total utility for a given plan; n is the number of activities; $V_{perf,i}$ is the (positive) utility earned for performing activity i; $V_{late,i}$ is the (negative) utility earned for arriving late to activity i; and $V_{tr,i}$ is the (usually negative) utility earned for traveling during trip i. Activities are assumed to wrap around the 24-hours-period, that is, the first and the last activity are stitched together. In consequence, there are as many trips between activities as there are activities.

[2] The terms "score" and "utility" refer to the same absolute value. "Utility" is the common expression in economic evaluation and is therefore used in this paper.

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

$$V_{perf,i}(t_{perf,i}) = \beta_{perf} \cdot t_{*,i} \cdot \ln\left(\frac{t_{perf,i}}{t_{0,i}}\right) \tag{3}$$

  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. As long as dropping activities from the plan is not allowed, $t_{0,i}$ has essentially no effect.

- The (dis)utility of being late is uniformly assumed as:

$$V_{late,i}(t_{late,i}) = \beta_{late} \cdot t_{late,i} \ , \tag{4}$$

  where β_{late} is the marginal utility (in 1/h) for being late, and $t_{late,i}$ is the number of hours late to activity i. β_{late} is usually negative.

- The (dis)utility of traveling used in this paper is estimated from survey data. It will be explained in an extra section.

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 $-\beta_{perf}\, t_{*,i}/t_{perf,i} \approx -\beta_{perf}$. Similarly, that opportunity cost has to be added to the time spent traveling.

No opportunity cost needs to be added to late arrivals, because the late arrival time is spent somewhere else. In consequence, the effective (dis)utility of arriving late remains at β_{late}.

3 Estimation of the Income-Contingent Utility Function

3.1 Estimation Data

Data for estimation of the travel related part of the utility function presented in this paper is taken from stated preference surveys run by the Institute for Transport Planning and System at ETH Zurich (Vrtic et al., 2008). Estimation of the late arrival penalty is retrieved from a time and route choice survey. All other estimations use data from a mode and route choice survey.

3.2 Functional Form

There is some agreement in literature that income should be considered in transport policy analysis, see, e.g., Small (1983); Herriges and Kling (1999); Kockelman (2001); Mackie et al. (2001); Bates (2006, 1987); Franklin (2006). The argument essentially is that monetary price changes affect people with different income differently. Conversely, if different income groups need to be compensated for losses or should be taxed for gains from non-monetary policy measures, the valuation of these offsetting payments is income-contingent. This paper demonstrates how these insights can be used constructively in an agent-based approach.

The starting point for the travel related part of the (dis)utility functions used in this paper is loosely based on Franklin (2006) and Franklin (2007) and is similar to Kickhöfer (2009).[3] In that paper, two transport modes are available: car and public transit , resulting in the following utility functions:

$$\begin{aligned} V_{car,i,j} &= \beta_{cost} \cdot \ln(y_j - c_{i,car}) + \beta_{tt_{car}} \cdot t_{i,car} \\ V_{pt,i,j} &= \beta_{cost} \cdot \ln(y_j - c_{i,pt}) + \beta_{tt_{pt}} \cdot t_{i,pt} \ , \end{aligned} \tag{5}$$

where y_j is the daily income of person j, c_i is monetary cost for the trip to activity i, and t_i the corresponding travel time. Monetary cost and travel time are mode dependent, indicated by the indices. Utilities are computed in "utils"; a possible conversion into units of money or "hours of leisure time" (Jara-Díaz et al., 2008) needs to be done separately. Daily income y_j is obtained by the following calculation:

$$y_j = \frac{y_{year,HH}}{n_{HH} \cdot 240} \ ,$$

where $y_{year,HH}$ depicts the income of the household per year, n_{HH} the number of persons in the household and 240 the number of working days per year.

It was, however, not possible to use this form directly, since the survey data contains relatively long trips, meaning that $y_j - c_i$ can become negative, in which case the logarithm does not work.[4] To circumvent this problem, Taylor's theorem is used to approximate the logarithm,

$$\ln(y_j - c_i) \approx \ln(y_j) - c_i \cdot [ln(y_j)]' = ln(y_j) - \frac{c_i}{y_j} \ , \tag{6}$$

which results into the quite normal $1/y$ dependency of the cost term and thus seems quite plausible. Applying (6) to (5) and setting the estimated parameters

$$\beta_{cost} = 4.58 \ , \qquad \beta_{tt_{car}} = -2.83/h \ , \qquad \text{and} \qquad \beta_{tt_{pt}} = -1.86/h \ ,$$

leads to the functional form:

$$\begin{aligned} V_{car,i,j} &= + \ 4.58 \ln(y_j/CHF) - 4.58 \frac{c_{i,car}}{y_j} - \frac{2.83}{h} t_{i,car} \\ V_{pt,i,j} &= + \ 4.58 \ln(y_j/CHF) - 4.58 \frac{c_{i,pt}}{y_j} - \frac{1.86}{h} t_{i,pt} \end{aligned} \tag{7}$$

It might be a bit surprising that the disutility of travel time comes out higher for car than for public transit. It is, however, consistent with the higher costs of $c_{pt} = 0.28$ CHF/km[5] assumed for public transit than for car ($c_{car} = 0.12$ CHF/km), which were used in the survey (Vrtic et al., 2008) and will be used in the simulations. Clearly and somewhat unusual, for Switzerland, public transit is the higher value mode compared to car.

Due to this specification, Values of Time (VoT) are obviously income *and* mode dependent. The VoT for the median income of the sample ($y_{median} = 155$ CHF per person and day) turn out to be 96 CHF/h for car and 63 CHF/h for public transit , respectively. These values are two to three times higher as those in Vrtic et al. (2008). Thus, the inclusion of income in the utility function seems to have unintended impacts on the VoT. This effect should be addressed in future research. However, note that e.g. the VoT for public transit varies from 7 to 330 CHF/h

[3] Two different forms for the alternative specific constants were also estimated. Both, the income-contingent bias term (Franklin, 2006) and the general alternative specific constant (Train, 2003), were estimated not significantly different from zero and are therefore not considered in the functional form of the utility functions. This essentially means that neither an income-contingent nor a general a-priori preference is found for one of the transport modes.

[4] One may argue that in such cases the model should reject the journey completely, at least if it is a regular journey (M. Wegener, personal communication).

[5] 1 CHF = 0.92453057 US$, exchange rate at 28.09.2008.

along the income range, which naturally includes the values from the linear model in Vrtic et al. (2008).

Another open question at this point is how much of the estimated disutility from traveling are opportunity costs of time, and how much is an additional disutility caused by traveling in the corresponding mode. This approach is consistent with economic approaches where there is an inherent opportunity cost of time and additional utilities or disutilities depending on how the time is spent (e.g. Jara-Díaz et al., 2008). Unfortunately, these values cannot be obtained from the survey as it was taken. Because of this, it was assumed that traveling in public transit neither adds nor subtracts from the opportunity cost of time. This implies $\beta_{perf} = 1.86/h$ in (3), and modifies the travel related part of the utility functions to:

$$\begin{aligned} V_{car,i,j} &= + \; 4.58 \ln(y_j/CHF) \; - \; 4.58 \frac{c_{i,car}}{y_j} \; - \; \frac{0.97}{h} t_{i,car} \\ V_{pt,i,j} &= + \; 4.58 \ln(y_j/CHF) \; - \; 4.58 \frac{c_{i,pt}}{y_j} \end{aligned} \tag{8}$$

Applying (8), (4) and (3) to (2) results in the two final utility functions used in this paper, which are selected depending on the mode of the ith trip:

$$\begin{aligned} V_{car,i,j} &= +\frac{1.86}{h} t_{*,i} \cdot \ln(\frac{t_{perf,i}}{t_{0,i}}) \; -\frac{1.52}{h} t_{late,i} \; -4.58 \frac{c_{i,car}}{y_j} \; -\frac{0.97}{h} t_{i,car} \\ V_{pt,i,j} &= +\frac{1.86}{h} t_{*,i} \cdot \ln(\frac{t_{perf,i}}{t_{0,i}}) \; -\frac{1.52}{h} t_{late,i} \; -4.58 \frac{c_{i,pt}}{y_j} \end{aligned} \tag{9}$$

The income-related offset $+4.58 \ln(y_j/CHF)$ in (8) can be interpreted as the utility earned from daily income. It is therefore calculated once for each individual and added to the overall utility score of daily plans and was removed from the activity related functions in (9). The marginal utility for being late β_{late} was computed similar to Kickhöfer (2009): in a time and route choice survey from ETH Zurich (Vrtic et al., 2008) people stated their willingness to pay in order to reduce the probability of being late. Based on this data, β_{late} was estimated and re-scaled with respect to the cost related behavioral parameter β_{cost} in (7), resulting in $\beta_{late} = 1.52/h$. The parameter for late arrival will only be used in the following test scenario (Sec. 4), but not for the real-world scenario of Zurich metropolitan area (Sec. 5).

3.3 Income Generation

Income is generated based on a Lorenz curve. Due to the lack of exact data the functional form of the Lorenz curve is approximated. Then the income curve, the first derivative of the Lorenz curve, is calculated (Kämpke, 2008).[6] To generate personal incomes for the agents, a random number between 0 and 1 is drawn from a uniform distribution. For this number, the corresponding value on the income curve is calculated and multiplied by the median income. Doing this for all members of the synthetic population, an income distribution is derived, similar to the distribution in reality.

Adding income at an individual level results in a personalized utility function for each agent. In the test scenario described in the following section, income is the only varying attribute between the agents. The real-world scenario in the subsequent section, however, includes varying trip distances and daily plans so that demographic attributes of each agent are strongly personalized.

[6] The Lorenz curve is $L(x) \propto \int_0^x y(\xi)\, d\xi$. Therefore, $L'(x) \propto y(x)$. The correct scaling is given by the fact that $y(0.5)$ is the median income.

4 Test Scenario

The goal of this section is to verify the correctness and plausibility of the estimated choice model and the underlying implementation. A simple setup is used in order to test the plausibility of traveler choice reactions resulting from a policy change.

4.1 Network

The test network (see Fig. 1) consists of a cycle of one-way links with (unrealistically) high capacities so as to minimize their influence on traffic patterns, essentially making it possible for most agents to drive with free speed. One link between home and work location has reduced capacity of 1000 veh/h, building a bottleneck.

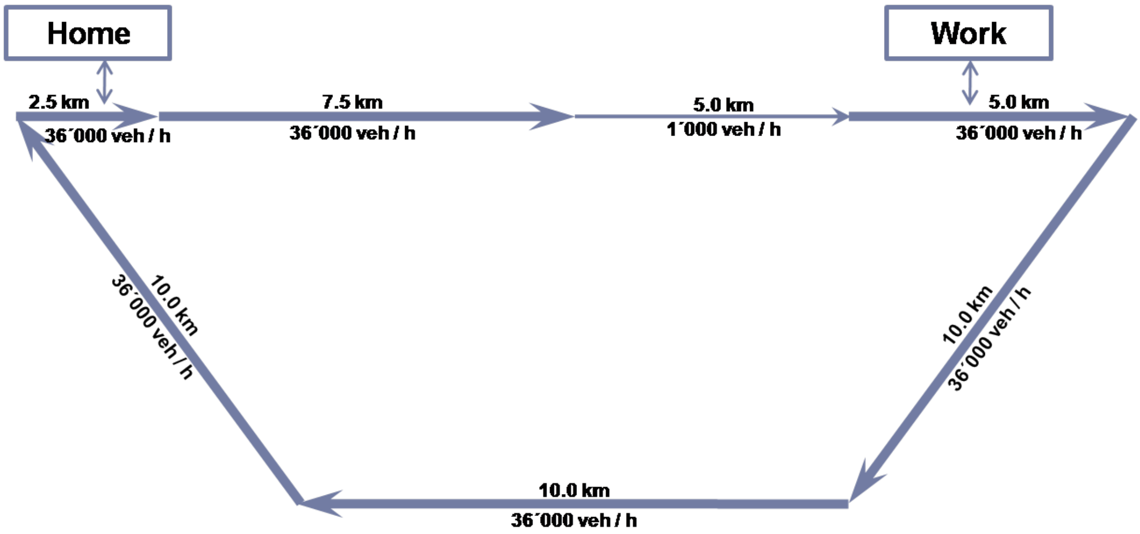

Figure 1: The layout of the testnetwork with link attributes. Traffic runs clockwise starting at the home location. Between home and work location lies a bottleneck link with a capacity limited to 1000 veh/h.

4.2 Initial Plans

The synthetic population consists of 2000 agents. All agents start at their home activity, which they initially leave at 6:00 a.m. They initially drive to work with a car, where they initially stay for 8 hours, after which they drive home again. The home to work trip has a length of 17.5 km while the way back is 32.5 km long. Speed limit is at 50 km/h so the free speed travel time from home to work by car is 21 minutes while 39 minutes are needed for the way back home. Thus, the total free speed travel time driving by car is 60 minutes. As the agents are forced to remain on that route, the scenario is similar to the well-known Vickrey bottleneck scenario (Arnott et al., 1990; Vickrey, 1969); also see below for more details.

In addition, each agent possesses an initially non-active plan that uses the public transit mode for both trips. These trips take twice as long as by car at freespeed, i.e. 42 minutes from home to work, and 78 minutes for the way back. The total public transit travel time is 120 minutes. In contrast to the car travel times, these transit travel times are not affected by congestion. Since public transit is assumed to run continuously and without capacity restrictions, a home departure at time t will always result in a work arrival at $t + 42min$. Work opens at 7:00 a.m. and closes at 6:00 p.m. In order to obtain the similarity to the Vickrey bottleneck scenario, an additional behavioral parameter of $\beta_{late} = -1.52/h$ is used, i.e. deducting $-1.52/h \cdot t_{late}$ for arriving late. Any arrival time after 7:00 a.m. is directly considered as late.

Estimation of income for the synthetic population, as described in Sec. 3.3, is based on values for the Kanton Zurich in 2006.[7] The median income for that year is 46 300 CHF.

4.3 Behavioral Parameters

The behavioral parameters are set and can be interpreted as follows:

- marginal utility of performing an activity at its typical duration: $\beta_{perf} = 1.86/h$
- marginal disutility of arriving late: $\beta_{late} = -1.52/h$
- marginal utility for traveling by car: $\beta_{tr,car} = -0.97/h$
- marginal utility for traveling by public transit : $\beta_{pt} = 0$
- factor for the logit choice (eq. 1): $\beta = 20$
- "typical" durations of $t_{*,w} = 8$ and $t_{*,h} = 12$ hours for work and home mean that work and home times have a tendency to arrange themselves with a ratio of 8:12 (i.e. 2:3).

  The activity of the home activity is "wrapped around", i.e. a departure at 6:00 a.m. and a return at 5:00 p.m. results in a home activity duration of 13 hours.

A work start time at exactly 7:00 a.m. means that (a) no utility can be accumulated from an arrival earlier than 7:00 a.m., and (b) any late arrival is immediately punished with $\beta_{late} = -1.52/h$.

Because of the argument made earlier regarding the opportunity cost of foregone activity time when arriving early, the *effective* marginal disutility of early arrival is $\beta_{early_{eff}} = -\beta_{perf}\, t_{*,i}/t_{perf,i} \approx -\beta_{perf} = -1.86/h$ which is equal to the effective marginal disutility of traveling with public transit $\beta_{tr,pt_{eff}}$. The effective marginal disutility of traveling with car is, by the same argument, $\beta_{tr,car_{eff}} = -\beta_{perf}\, t_{*,i}/t_{perf,i} - |\beta_{tr,car}| \approx -\beta_{perf} - |\beta_{tr,car}| \approx -2.83/h$. The return trip has no influence since there is no congestion. Overall, the *effective* values of car travel time of our study would correspond to the values $(\beta_{early_{eff}}, \beta_{tr,car_{eff}}, \beta_{late_{eff}}) = (-1.86, -2.83, -1.52)$ of the Vickrey bottleneck scenario (Vickrey, 1969; Arnott et al., 1990).

4.4 Simulation Runs

First, a "preparatory run" is performed by running the base case for 4000 iterations. During the first 2000 iterations, 10% of the agents perform "time adaptation", i.e. they make a copy of an existing plan and shift each element of their time structure by a random amount between zero and 7.5 minutes. The other 90% of the agents switch between their existing plans according to (1), which means that they potentially also switch the mode. During the second 2000 iterations, time adaptation is switched off; in consequence, agents only switch between existing options according to (1). That is, their choice set now remains fixed to what they have found in the first 2000 iterations, and they choose within this set according to a logit model.

After this, the speed increase for public transit is introduced. It now takes only 1.8 (instead of 2.0) times as long as traveling by car on an empty network. This corresponds to a pt speed increase of 10%. The policy case is run for another 2000 iterations, starting from the final iteration of

[7] `http://www.statistik.zh.ch/themenportal/themen/index.php`, last access 01.08.2009 `http://www.statistik.zh.ch/themenportal/themen/aktuell_detail.php?id=2752&tb=4&mt=0`, last access 01.08.2009

the preparatory run. For the first 1000 iterations, the time adaption module is again switched on, with the same 10% replanning fraction. The final 1000 iterations are once more with a fixed choice set. The following policy measure is investigated: For further analysis, iteration 4000 of the base case is then compared to the final iteration of the policy case.

4.5 Results

Since car is the low value and public transit the high value mode, low income people predominantly use the car while high income people predominantly use public transit (Fig. 2a). When the pt speed is increased, the modal split predictably shifts from car to pt, from 54% : 46% to 42% : 58% (car:pt). Also predictably but importantly, this happens through a shift of the income level that divides the two regimes – this level, naturally, moves to lower incomes (Fig. 2b).

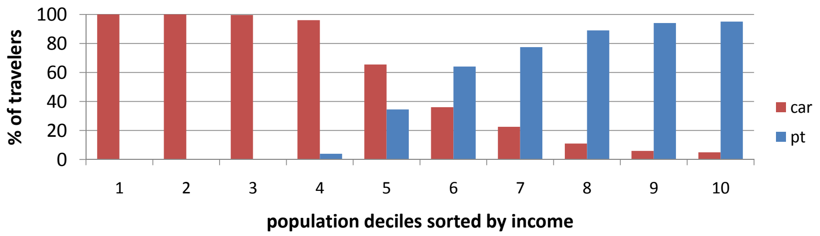

(a) Base case

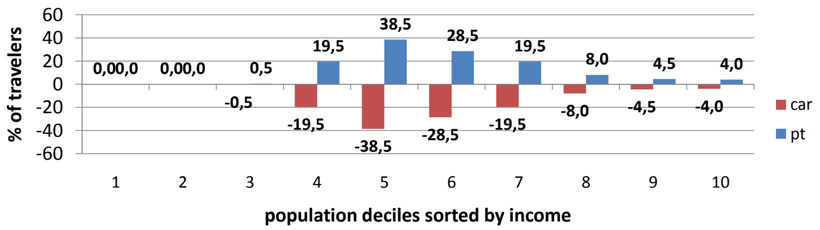

(b) Changes due to the public transit speed increase

Figure 2: Modal Split over deciles of the population sorted by income. Absolute values for the base case, changes in percentage points for the pt speed increase. Red bars depict car drivers, blue bars public transit users.

Fig. 3 shows, agent-by-agent, the utility differences between the base case and the policy case as a scatter plot over deciles of the population. Every decile, summed up over all three plots, contains the same number of agents, sorted by their income. For example, the first decile from 0% to 10% includes the 10% of agents with the lowest incomes. Fig. 3a shows synthetic travelers that choose the same transport mode before and after the policy change, red for the car mode, green for the pt mode. One notices a homogeneous utility increase of about 0.35 for all pt users, and a smaller increase of, in the average, about 0.25 for the car users. This is a plausible consequence of the fact that the policy measure benefits the pt users directly, while the car users benefit from congestion relief because of reduced car mode share.

The car users, who, because of stochastic congestion effects, face rather strong fluctuations of their utilities from iteration to iteration, and therefore also from base case to policy case. These fluctuations are much less pronounced for pt users, where the transport mode is assumed to be completely reliable. Figs. 3b and 3c show synthetic travelers that change their transport mode, Fig. 3b from pt to car, and Fig. 3c from car to pt. With a pt improvement measure, a switch

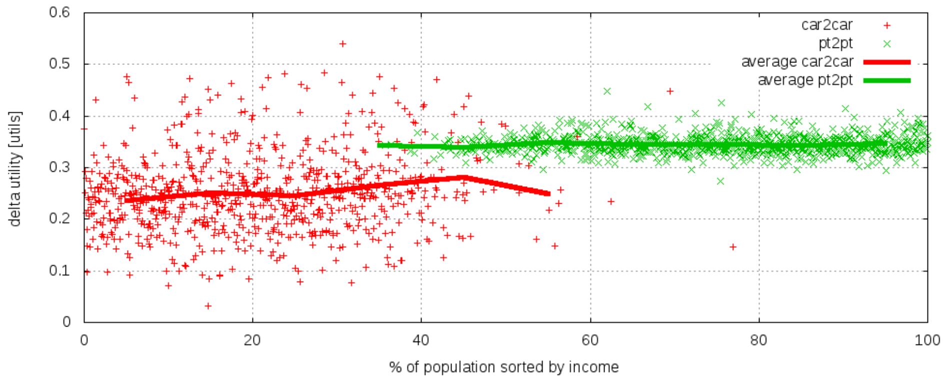

(a) Utility changes for synthetic travelers who keep their mode

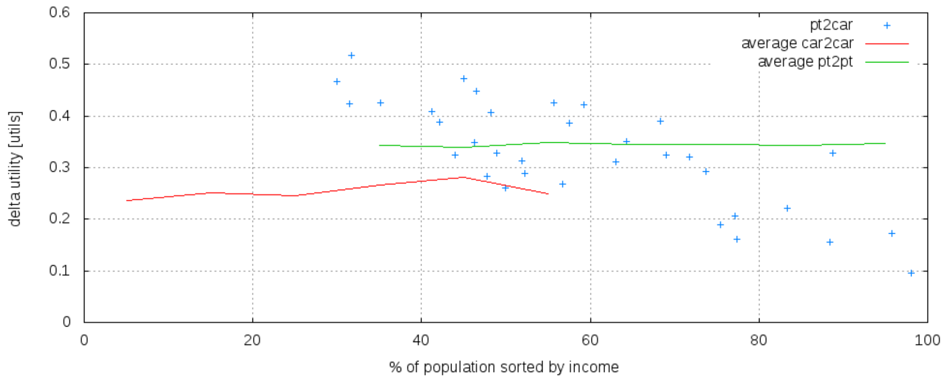

(b) Utility changes for synthetic travelers who switch from pt to car

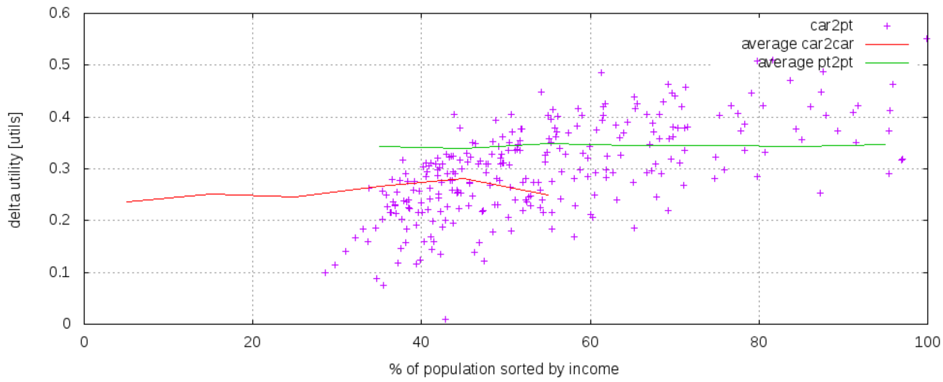

(c) Utility changes for synthetic travelers who switch from car to pt

Figure 3: Utility changes, sorted by income. Every synthetic traveler is plotted according to her relative position on the income scale. The green and red lines in Figs. 3b and 3c denote the average utility changes from Fig. 3a.

from pt to car (Fig. 3b) is not what should be expected. These synthetic travelers are "logit switchers", in the following sense:

- Those that gain more than the average by the switch, towards the lower income scales, are travelers that would have had a higher score already in the base case when using car.
- Those that gain less than the average by the switch, towards the higher income scales, are travelers that would have gained even more by staying with pt.

However, in both cases the score computation is done according to the systematic component of utility as described in Sec. 2. The switches are therefore caused by changes in the random components (the "epsilons").

Fig. 3c also contains those "logit switchers", this time from car to pt. But in addition, it also contains "systematic" switchers who change the mode due to a utility gain in the systematic component. The density of switchers is largest towards middle income groups, since there, the systematic switchers are located. One could, in principle, also attach the random elements as random but fixed to every alternative of every traveler; see, e.g., Horni, A. et al for an example. Karlstrom and Morey (2004) and therefore Franklin (2006) assume such an interpretation in their welfare computations.

An interpretation of the results according to the well-known "rule-of-a-half" (Jara-Díaz, 2007), where the pre-existing users of an improved facility obtain the full utility gains, whereas the users switching towards the improved facility in average obtain half of those gains. In the present case, the situation gets more complicated because of substitution effects: Also users not using the improved facility gain because of congestion relief. Differentiated by user groups, we obtain:

Average utility gain pt2pt:	0.345
Average utility gain car2pt:	0.297
Average utility gain car2car:	0.250

That is, the switchers from car to pt indeed gain, in the population average, the mean value between the (direct) pt gains and the (indirect) car gains. Nevertheless, the situation is confounded by the fact that there are also significant gains for the car users, which could not be computed by considering the pt facility alone.

Overall, the results demonstrate that the approach picks up the distributional effects of transport policy measures. As is plausible, a quality-of-service change affects the higher income groups more when assuming higher VoT with increasing income. Thus, the plausibility test can be regarded as successful. The approach is therefore applied to a real-world scenario of Zurich metropolitan area in the next section.

5 Scenario – Zurich Switzerland

The income-contingent utility function is now applied to a large-scale, real-world scenario. The area of Zurich, Switzerland, is used which counts about 1 million inhabitants. The following paragraphs give a simplified description of the scenario and focus on differences to similar simulations done by Chen et al. (2008).

5.1 Network and Population

The network is a Swiss regional planning network that includes the major European transit corridors. It consists of 24 180 nodes and 60 492 links.

The simulated demand consists of all travelers within Switzerland that are inside an imaginary 30 km boundary around Zurich at least once during their day (Chen et al., 2008; Vrtic et al., 2007). All agents have complete daily plans with activities like *home*, *work*, *education*, *shopping*, *leisure*, based on microcensus information (SFSO, 2000, 2006). The time window during which activities can be performed is limited to certain hours of the day: *work* and *education* can be performed from 7:00 a.m. to 6:00 p.m., *shopping* from 8:00 a.m. to 8:00 p.m., while *home* and *leisure* have no restrictions. Each agent initially gets two plans based on the same activity pattern. The first plan uses only car as transportation mode, while the second plan uses only public transit.

Unlike the test scenario described above, there is no punishment for being late. This is not possible because agents can have more than one work activity, e.g. one in the morning and one in the afternoon. In such a case it is complicated to specify when an agent starts an activity late or not.

To speed up computations, a random 10% sample is taken from the synthetic population for simulation, consisting of 181 725 agents. In this large-scale application, the agents can, in addition to the previously described time adaptation, also perform route adaptation, which is essential for the car mode. Mode adaptation is implicitly included as described in Sec. 2.1.

5.2 Income Generation

Income for the large-scale scenario is generated as described in Sec. 3.3. Region specific data is used for the Kanton[8] Zurich since here, income medians are available for each municipality. [9]

For every person living in the Kanton Zurich, the municipality of the person's home location is determined. Then the median income specific for this municipality is used for income calculation in conjunction with a Lorenz curve for the Kanton Zurich.[10] Incomes for persons living outside the borders of Kanton Zurich are computed with the help of the median income and the Lorenz curve of the Swiss Confederation.[11] The median income used for the Swiss Confederation is 43 665 CHF per household and year.

The resulting distribution with focus on the Kanton Zurich is shown in Fig. 4. While outside the Kanton's borders, income is equally distributed, one can see some geospatial differences in the area where detailed data is available. The structural pattern has similarities to official data of Zurich.[12]

5.3 Simulation Runs

In order to maintain consistency with the test scenarios, the total amount of iterations is reduced but the proportion of the different simulation steps is held constant. This means for the base case:

- For 1000 iterations, 10% of the agents perform "time adaptation" and 10% adapt routes. The other 80% of the agents switch between their existing plans, which implicitly includes mode choice as explained in Sec. 2.1.
- During the second 1000 iterations, time and route adaption are switched off; in consequence, agents only switch between existing options.

After this, the pt speed improvement is introduced. The policy case is run for another 1000 iterations, starting from the final iteration of the base case. Again, during the first 500 iterations 10% of the agents perform "time adaptation" while another 10% of agents adapt routes. Agents neither adapting time nor route switch between existing plans and such eventually switch between

[8] A Swiss "Kanton" is similar to a federal state.
[9] `http://www.statistik.zh.ch/themenportal/themen/daten_detail.php?id=759`, last access 30.10.2009
[10] `http://www.statistik.zh.ch/themenportal/themen/aktuell_detail.php?id=2752&tb=4&mt=0`, last access 30.10.2009
[11] `http://www.bfs.admin.ch/bfs/portal/de/index/themen/20/02/blank/dos/01/02.html`, last access 30.10.2009
[12] `http://www.stadt-zuerich.ch/content/dam/stzh/prd/Deutsch/Statistik/Publikationsdatenbank/Steuerstudie1.pdf`, last access 30.10.2009

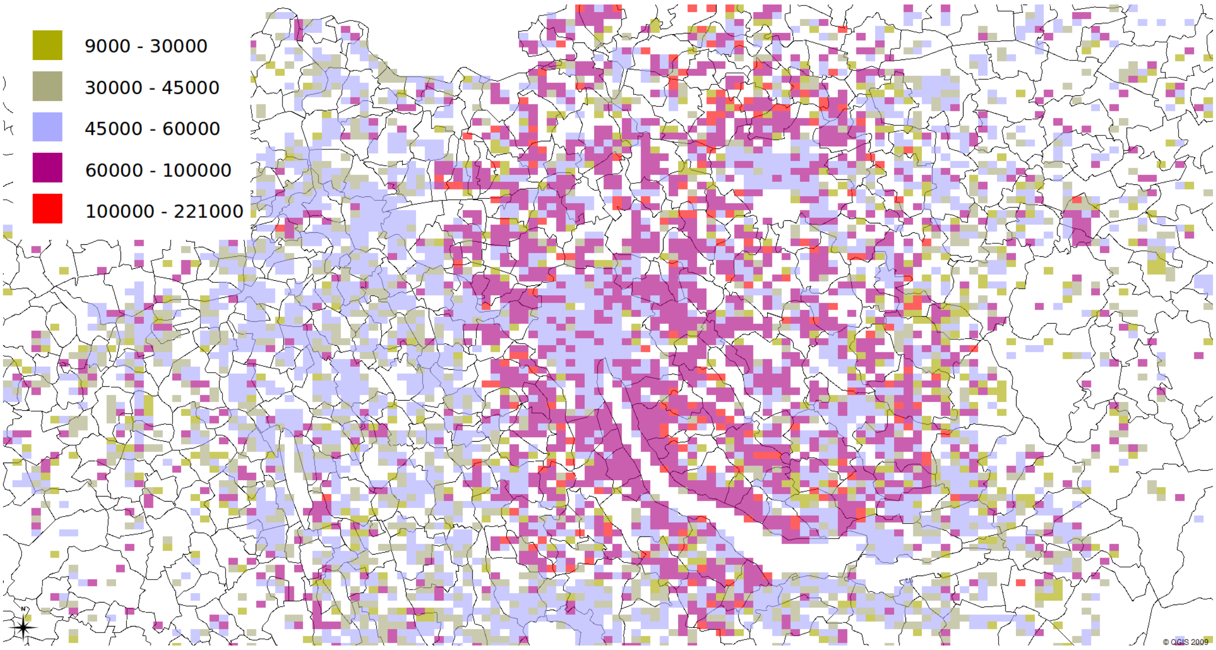

Figure 4: Computed income distribution for the Zurich scenario. Each cell is colored by the mean income of inhabitants [CHF/Year & Person]

transport modes. For the final 500 iterations only a fixed choice set is available. Different paramter combinations were tested, up to an overall 30% public transit speed increase.

For evaluating the impact of the pt speed increase, iteration 2000 of the base case is compared to the final iteration of the policy case.

5.4 Validation

Simulated traffic volumes are compared with the hourly traffic volumes from 159 real-world counting stations. Fig. 5 shows, in blue, the mean relative error of the reference Zurich scenario (Chen et al., 2008) between hourly flows in reality and hourly flows from the simulation. That run was based on $\beta_{perf} = 6/h$, $\beta_{tr,car} = -6/h$, $\beta_{pt} = -3/h$, and no dependence on travel distance or income was assumed.

The red curve depicts the same for the estimated income-contingent utility function. One notices a slight improvement of the mean relative error, especicially during day time in comparison with Chen et al. (2008). This underlines the advantage of using estimated behavioral parameters for more realistic results. Nevertheless, it is a bit surprising that, at least at the aggregated level, there is so little difference between the simulations. Presumably, this is due to the fact that the activity patterns, preferred activity durations, opening times, and transportation network structure are dominating the results. In particular, given the fact that the traffic counts are reproduced much better between 8:00 a.m. and 7:00 p.m. than the remainder of the day, one may speculate that the need to squeeze all activities into the available opening times is, in fact, the dominating force.

Together with the analysis of other traffic condition indicators, such as peak hours, modal split or the average trip duration or length, it can be stated that the base case seems to be a good starting point for investigating different transport policy measures.

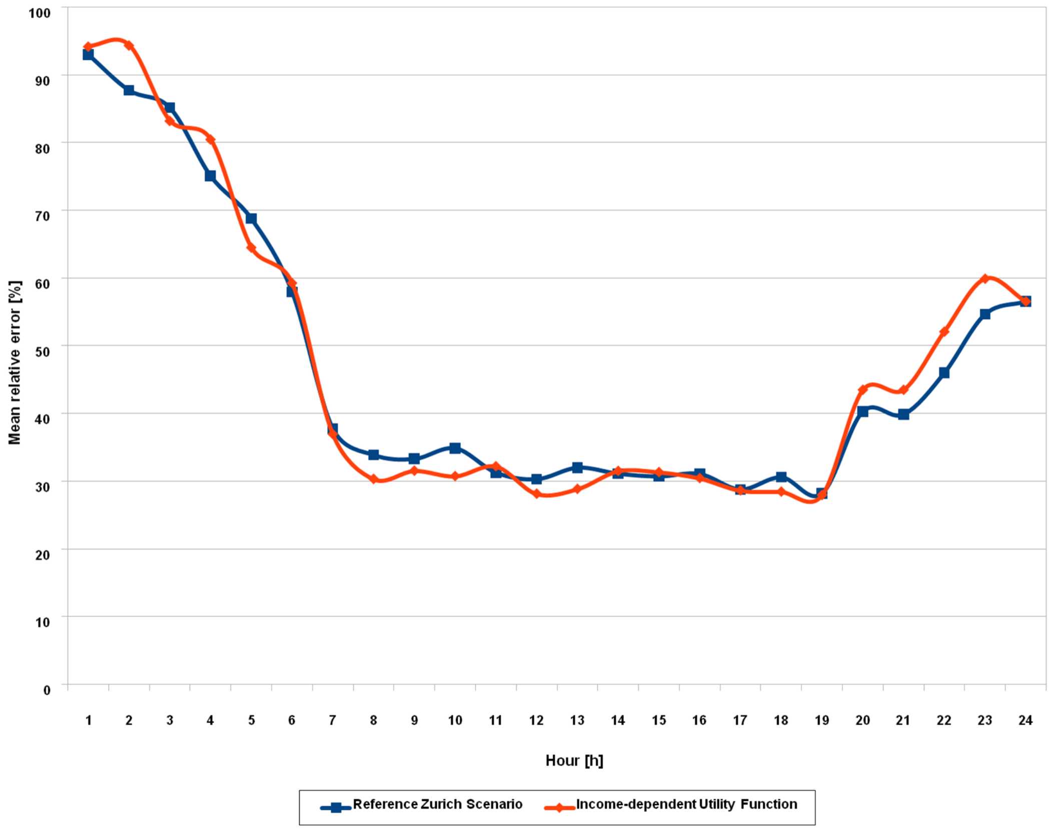

Figure 5: Realism of the two simulations. 159 traffic counting stations provide real traffic counts for the Zurich area. The blue curve shows the mean relative error when comparing the simulation traffic volumes of the reference Zurich scenario with real values, the red curve the comparison using the income-contingent utility function of this paper.

5.5 Results

The base case of the Zurich scenario exhihits a modal split of 60.9%:39.1% (car:pt). Fig. 6a depicts the modal split in the income deciles of the population. In contrast to the base case of the test scenario shown in Fig. 2a, the distribution here is more homogeneous: Both modes are used across all deciles.

Fig. 6b presents changes to the modal split in the income decils of the population compared to the base case. One can observe that with increasing income, more persons switch from car to pt.

Increasing utility gains of agents with higher income can also be seen in Fig. 7a that depicts the average utility change of each income decile. Each dot is located in the middle of the decile and represents the average utility change per decile. For representation purposes the dots are connected with lines. Obviously, one recognizes rising utility gains with increasing income. In terms of utils, the slope of the curve is slightly positive. The subsequent section will show, however, that this increase has even stronger effects when converting utils into money.

Fig. 7b breaks the average utility gains of every income decile from Fig. 7a down to two different user groups. As expected, synthetic persons using pt both before and after the pt improvement obtain most of the benefits (in green). Synthetic persons using car both before and after the pt

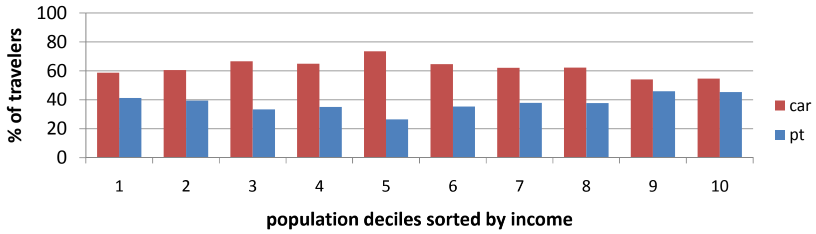

(a) Base Case

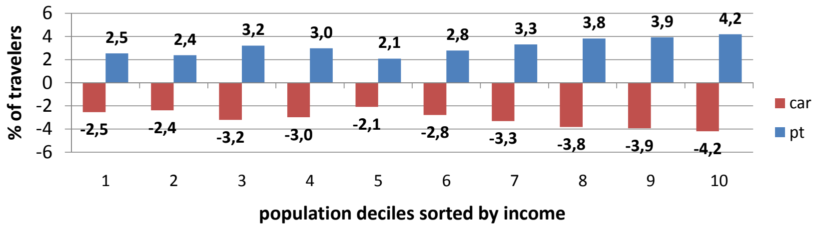

(b) Pt Speed Improvement

Figure 6: Modal Split over deciles of the population sorted by income. Absolute values for the base case, changes in percentage points for the pt speed increase. Red bars depict car drivers, blue bars public transit users.

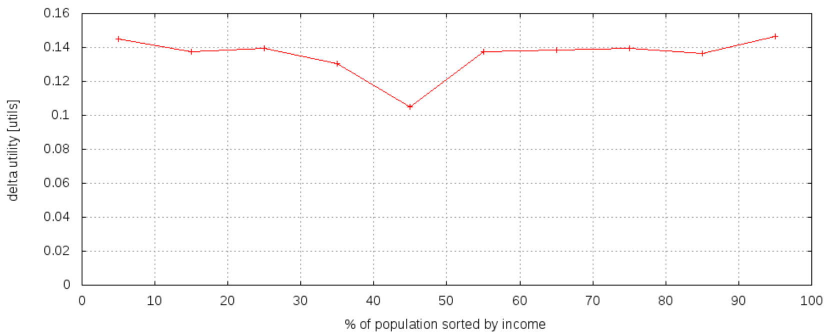

(a) Average utility changes

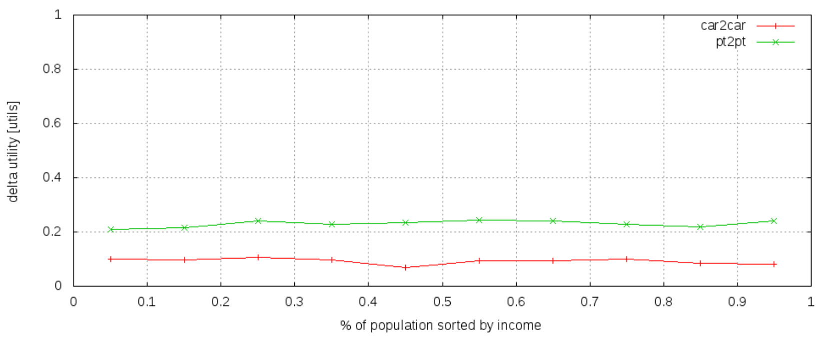

(b) Average utility changes per user group

Figure 7: Average utility changes per income decile

improvements also gain, but considerably less than the pt users (in red). Results for the mode switchers are not shown because the stochastic fluctuations overwhelm the signal.

6 Discussion

This paper starts from a estimates of individual utility functions based on the estimation of logit models. A possible interpretation of this approach is that the utility function is simply a device that helps to construct quantitatively descriptive behavioral models of individuals. In this way, one first constructs and estimates the behavioral model, and then runs a simulation model populated with entities (synthetic persons) using this behavioral model.

The interpretation of the utility as an indicator of individual gains or losses is essentially an afterthought, with no meaning to the simulation results except that random utility modeling has something to do with the individual opimization of utility. Still, the identification of winners and loosers seems to be inherently quite robust since it results directly from the behavioral model that is based on surveys or observations. It was also shown that it is possible to quantify individual gains and losses with respect to an individual utility level. For a possible aggregation of these quantitative numbers in order to obtain some indicator of welfare change and a discussion of problems, see below.

General critical issues of this approach enter from different angles, such as:

- Since the logit model is stochastic, our agents are stochastic as well, and some may *de*crease the systematic part of their utility function (what was called “logit switchers” above). This means that the results cannot be interpreted at the single-agent, single-plan level; one either needs to aggregate over subpopulations or over all the plans of every agent. The latter is less problematic in terms of the comparability of utility functions between different persons (see below).
- One may argue that it is invalid to cast human behavior into an optimization problem at all, and one should rather resort to simple procedural rules (e.g. Moss and Sent, 1999; Simon, 1997).
- Even if individuals’ behavior can be cast as an individual optimization problem, it is by no means clear that ex-post happiness is optimized by ex-ante optimization of this descriptive function.

Still, to re-iterate: If one takes into account that one has some random “logit” winners and losers, then the conceptual path leading to the identification of individual winners and losers seems quite straightforward.

Concerning the question whether individual utility changes can somehow be aggregated into some indicator of welfare, three obvious methods come to mind:

- Just add up all individual utility changes, possible by subpopulation (Fig. **??**).
- Convert all individual utility changes to individual willingness-to-pay/willingness-to-accept values (Fig. 8). In the context here, the most straightforward (albeit not the only) way to do this seems to convert the utility changes according to the utility of money $4.58/y_j$ in Eq. (9) into monetary changes:

  $$\Delta m_j = 4.58 \cdot y_j \cdot \Delta V_j \ , \qquad (10)$$

  where y_j is once more the daily income. Given a utility gain by the policy measure, this would be the amount of money one would need to take away to revert the person back to her original level of utility. For a discussion of income effects see below.

- Convert all individual utility changes to individual equivalent "hours of leisure time". Given a utility gain by the policy measure, this would be the amount of leisure time one would need to take away to revert the person back to her original level of utility. One has

$$\Delta V_j \approx \frac{\partial V_j}{\partial t_{perf,leisure,j}} \cdot \Delta t_{perf,leisure,j} \tag{11}$$

and therefore

$$\Delta t_{perf,leisure,j} \approx \left(\frac{\partial V_j}{\partial t_{perf,leisure,j}} \right)^{-1} \cdot \Delta V_j \ . \tag{12}$$

That is, the (linearized) gain/loss of every individual's equivalent hours of leisure time would be computed by multiplying every individual's utility gain/loss by their inverse marginal utility of leisure time. As a more robust alternative to a leisure activity, the "home" activity could be used.

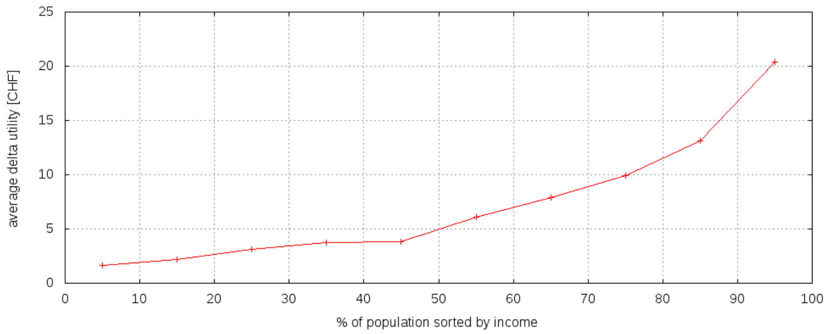

Figure 8: Average willingness-to-pay for the pt speed increase; values per income decile

The last item brings to light that maybe more effort needs to be made to build quantitative models that include time pressure. Clearly, the vastly different values of time by trip purpose (often up to a factor of four, see Mackie et al., 2003) indicate that this needs to be considered. Nevertheless, concentrating on time pressure alone means a different focus than values of time: While (monetary) VoT are also coupled to income, time pressure expressed in equivalent hours of leisure time may be the same for a single mother or for a busy executive.

At any rate, the comparison of Fig. 7b to Fig. 8 makes clear that the selection of the aggregation procedure can lead to quite different equity interpretations: While the simple but conceptually problematic summation of utilities leads to similar gains across all income groups (Fig. 7b), the willingness-to-pay, sorted by income group, implies that high income groups would have a disporportionately high willingness-to-pay for the measure considered (here the pt speed increase; Fig. **??**). This is, in fact, quite intuitive: Higher income groups have a disproportionately high willingness to pay for good schools, a good health system or a good transport system. In this sense, a progressive tax may not even be redistributive with respect to such types of government expenses, since it just reflects the willingness-to-pay for improving the corresponding services. Clearly, this is only true if the higher income groups are not operating in a separate system, such as private schools, private health insurance, or private transport systems. Let us stress once more that these different attempts of measuring welfare rely on exactly the same description of human behavior. This highlights that the model predicting the system's reaction to a policy change is independent from the different interpretations of measuring welfare.

Unfortunately, the above observations interact with economic appraisal. The aggregated willingness-to-pay may seem the most consistent approach, but is quite different from approaches to-date where time gains of all individuals are valued equally. Conversely, using aggregated equivalent hours of leisure time may seem more equitable, but it is a bit unclear how to convert them into money for cost-benefit analysis.[13] Given, however, that at this point one has already left the ground of individual monetized consumer surplus, it might be more honest to quantify the benefit no longer in monetary terms, but directly in equivalent hours of leisure time. Projects could be assessed by the equivalent hours of leisure time they would generate per unit of money invested.

Eq. (10) assumed that the conversion from utility into willingness-to-pay or willingness-to-accept, respectively, could be based on the income from before the introduction of the policy measure. Since the measure discussed in this paper does not change the income, this can be justified. In situations where the policy measure changes the available income, the issue is not so clear any more, and income effects (Herriges and Kling, 1999; Jara-Díaz and Videla, 1989) need to be taken into account. Since this is not relevant for the present paper, let it suffice to say that one would need to go back to Eq. (5) or Eq. (6) and trace the effect of the interaction between cost c and income y and how that affects Eq. (10).

Two other issues should be addressed in the future:

- The survey should be designed in a way that all neecessary parameters can be estimated independently.
- For the present paper, public transit is assumed to be 100% reliable, and no fluctuations due to geographic location or line cycles are considered. In principle, multi-agent transport simulations make it possible to combine multiple demographic attributes of the population of interest, e.g. by viewing the geospatial distribution of winners and losers of a measure (see Grether et al., 2008). Neither the measure of this paper nor the public transit simulation features geospatial diversity. Thus analysis in the geographic dimension is strongly homogeneous and a spatial pattern is not visible. In case of a policy that is targeted on some geospatial impact (Rieser and Nagel, 2009), the multi-agent approach could give interesting insights into geospatial distribution of gains and losses.

7 Conclusion

Standard economic policy evaluation recommends the realisation of projects if the monetized aggregated economic benefit outweighs their monetary costs. In democratically organized societies, however, it is quite difficult to realize projects when they have a negative impact on some part of the population even if this is a minority. The microscopic simulation approach presented in this paper is capable to help designing better solutions in such situations.

In addition, it is shown that income can and needs to be included in utility calculations for a better understanding of problems linked to acceptability. Furthermore, and going beyond Franklin (2006), it is shown that the approach works in a large-scale real world example.

Acknowledgments

This work was funded in part by the "Bundesministerium für Bildung und Forschung" (BMBF) within the research project "Adaptive Verkehrssteuerung" (Advest), and in part by the "German Research Foundation" (DFG) within the research project "Detailed evaluation of transport policies using microsimulation". Our computer cluster is maintained by the Department of Mathematics at Berlin Institute of Technology. The paper benefited from discussions with G. Sammer.

[13] A possible way may be to use some average, i.e. income-*in*dependent, conversion.

References

R. Arnott, A. de Palma, and R. Lindsey. Economics of a bottleneck. *Journal of Urban Economics*, 27(1):111–130, 1990.

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.

J. Bates. Measuring travel time values with a discrete choice model: A note. *Economic Journal*, 97(386):493–98, June 1987. URL `http://ideas.repec.org/a/ecj/econjl/v97y1987i386p493-98.html`.

J. Bates. Economic evaluation and transport modelling: Theory and practice. In K.W. Axhausen, editor, *Moving through nets: The physical and social dimensions of travel*, chapter 10, pages 279–351. Elsevier, 2006.

M. Ben-Akiva and S. R. Lerman. *Discrete choice analysis*. The MIT Press, Cambridge, MA, 1985.

J. Bottom. *Consistent anticipatory route guidance*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, 2000.

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.

Y. Chen, M. Rieser, D. Grether, and K. Nagel. Improving a large-scale agent-based simulation scenario. VSP Working Paper 08-15, TU Berlin, Transport Systems Planning and Transport Telematics, 2008. See www.vsp.tu-berlin.de/publications.

J.P. Franklin. Decomposing the distributional effects of roadway tolls. Technical report, Washington D.C., 2007.

J.P. Franklin. *The distributional effects of transportation policies: The case of a bridge toll for Seattle*. PhD thesis, University of Washington, Seattle, 2006.

C. Gawron. *Simulation-based traffic assignment*. PhD thesis, University of Cologne, Cologne, Germany, 1998. URL `www.zaik.uni-koeln.de/AFS/publications/theses.html`.

D. Grether, Y. Chen, M. Rieser, U. Beuck, and K. Nagel. Emergent effects in multi-agent simulations of road pricing. In *Annual meeting of the European Regional Science Association ERSA '08*, 2008.

D. Grether, Y. Chen, M. Rieser, and K. Nagel. Effects of a simple mode choice model in a large-scale agent-based transport simulation. In A. Reggiani and P. Nijkamp, editors, *Complexity and Spatial Networks. In Search of Simplicity*, Advances in Spatial Science, chapter 13, pages 167–186. Springer, 2009a. doi: 10.1007/978-3-642-01554-0.

D. Grether, B. Kickhöfer, and K. Nagel. Policy evaluation in multi-agent transport simulations considering income-dependent user preferences. In *Proceedings of the Meeting of the International Association for Travel Behavior Research (IATBR)*, Jaipur, India, December 2009b. Also VSP WP 09-13; see www.vsp.tu-berlin.de/publications.

D. Grether, B. Kickhöfer, and K. Nagel. Policy evaluation in multi-agent transport simulations. *Transportation Research Record: Journal of the Transportation Research Board*, 2175:10–18, 2010. ISSN 0361-1981. doi: 10.3141/2175-02.

J.A. Herriges and C.L. Kling. Nonlinear income effects in random utility models. *The Review of Economics and Statistics*, 81(1):62–72, 1999. URL `http://www.jstor.org/stable/2646786`.

Horni, A. et al. in preparation.

S. Jara-Díaz. *Transport Economic Theory*. Elsevier, 2007. ISBN 9780080450285.

S. Jara-Díaz and J. Videla. Detection of income effect in mode choice: Theory and application. *Transportation Research Part B: Methodological*, 23(6):393–400, 1989. doi: 10.1016/0191-2615(89)90040-4. URL `http://www.sciencedirect.com/science/article/B6V99-466M41S-J/2/13e7436b7b974449ec62759d430d720d`.

S. Jara-Díaz, M. Munizaga, P. Greeven, R. Guerra, and K. W. Axhausen. Estimating the value of leisure from a time allocation model. *Transportation Research B*, 42(10):946–957, 2008. doi: 10.1016/j.trb.2008.03.001.

Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–91, March 1979.

T. Kämpke. The use of mean values vs. medians in inequality analysis. 2008. Forschungsinstitut für anwendungsorientierte Wissensverarbeitung.

Anders Karlstrom and Edward R. Morey. Calculating the exact compensating variation in logit and nested-logit mode with income effects: theory, intuition, implementation, and application. Elsevier Science, May 2004.

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, May 1991.

B. Kickhöfer. Die Methodik der ökonomischen Bewertung von Verkehrsmaßnahmen in Multiagentensimulationen. Master's thesis, TU Berlin, June 2009. Also VSP WP 09-10; see www.vsp.tu-berlin.de/publications.

K.M. Kockelman. A model for time- and budget-constrained activity demand analysis. *Transportation Research Part B: Methodological*, 35(3):255–269, 2001.

N. Lefebvre and M. Balmer. Fast shortest path computation in time-dependent traffic networks. In *Proceedings of Swiss Transport Research Conference (STRC)*, Monte Verita, CH, September 2007. See www.strc.ch.

P.J. Mackie, S. Jara-Díaz, and A.S. Fowkes. The value of travel time savings in evaluation. *Transportation Research E*, 37(2-3):91–106, 2001.

P.J. Mackie, M. Wardmann, A.S. Fowkes, G. Whelan, J. Nellthorp, and J. Bates. Valuation of travel time savings in the uk. summary report to teh department for transport, london. Technical report, carried out under contract PPAD 9/65/79 by the Institute for Transport Studies, University of Leeds, and John Bates Services, 2003.

K. Meister, M. Rieser, F. Ciari, A. Horni, M. Balmer, and K.W. Axhausen. Anwendung eines agentenbasierten Modells der Verkehrsnachfrage auf die Schweiz. In *Proceedings of Heureka '08*, Stuttgart, Germany, March 2008.

S. Moss and E.-M. Sent. Boundedly versus procedurally rational expectations. In A.H. Hallett and P. McAdam, editors, *Analyses in macroeconomic modelling*, Advances in computational economics, pages 115–146. Springer, Heidelberg, 1999.

K. Nagel, D. Grether, U. Beuck, Y. Chen, M. Rieser, and K.W. Axhausen. Multi-agent transport simulations and economic evaluation. *Journal of Economics and Statistics (Jahrbücher für Nationalökonomie und Statistik)*, 228(2+3):173–194, 2008. Special issue on agent based models for economic policy advice, edited by B. LeBaron and P. Winker.

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. An improved framework for large-scale multi-agent simulations of travel behaviour. In P. Rietveld, B. Jourquin, and K. Westin, editors, *Towards better performing European Transportation Systems*. Routledge, London, 2006.

M. Rieser and K. Nagel. Combined agent-based simulation of private car traffic and transit. In *Proceedings of the Meeting of the International Association for Travel Behavior Research (IATBR)*, Jaipur, India, December 2009. Also VSP WP 09-11; see www.vsp.tu-berlin.de/publications.

M. Rieser, D. Grether, and K. Nagel. Adding mode choice to a multi-agent transport simulation. *Transportation Research Record: Travel Demand Forecasting 2009*, 2132:50–58, 2009. doi: 10.3141/2132-06.

SFSO. Eidgenössische Volkszählung. Swiss Federal Statistical Office, Neuchatel, 2000.

SFSO. Ergebnisse des Mikrozensus 2005 zum Verkehr. Swiss Federal Statistical Office, Neuchatel, 2006.

H.A. Simon. *An empirically based microeconomics*. Cambridge University Press, 1997.

K.A. Small. The incidence of congestion tolls on urban highways. *Journal of Urban Economics*, 13(1):90–111, January 1983. URL `http://ideas.repec.org/a/eee/juecon/v13y1983i1p90-111.html`.

K. Train. *Discrete choice methods with simulation*. Cambridge University Press, 2003.

W.S. Vickrey. Congestion theory and transport investment. *The American Economic Review*, 59 (2):251–260, 1969.

M. Vrtic, N. Schüssler, A. Erath, K. Meister, and K. Axhausen. Tageszeitliche Fahrtenmatrizen im Personenverkehr an Werktagen im Jahr 2000. Research report, Swiss Federal Department for Environment, Transport, Energy and Communication, Swiss Federal Office for Spatial Development, Swiss Federal Roads Authority and Swiss Federal Office for Transport, IVT, ETH Zürich, Zürich, 2007.

M. Vrtic, N. Schüssler, A. Erath, M. Bürgle, K.W. Axhausen, E. Frejinger, M. Bierlaire, R. Rudel, S. Scagnolari, and R. Maggi. Einbezug der Reisekosten bei der Modellierung des Mobilitätsverhaltens. Schriftenreihe 1191, Bundesamt für Strassen, UVEK, Bern, CH, 2008. final report for project SVI 2005/004.