A model of risk-sensitive route-choice behaviour and the potential benefit of route guidance

J. Illenberger, G. Flötteröd, K. Nagel

***Abstract*—In this paper, we present simulation studies where we investigate the potential benefit of route guidance information in the context of risk-sensitive travellers. We set up a simple two-routes scenario where travellers are repeatedly faced with risky route-choice decisions. The risk-averseness is implicitly controlled through a generic utility function which evaluates the performance of travellers. We vary both the travellers sensitivity towards risk and the equipment fraction of route guidance devices and show that the benefits of guided travellers increase with their sensitivity towards risk.**

I. Introduction

In the field of transport planning, engineers agree that the problems of transportation are no longer a matter of extending the infrastructure with concrete and steel. The limited funds for road investment and the growing ecological consequences redirected attention to policies that account for an efficient use of existing transport networks [1]. Advanced traveller information systems (ATIS) and advanced traveller management systems (ATMS) are intended to fill in here by providing accurate information through a variety of devices.

An important aspect is the response of drivers to provided information. Although a lot of research has been conducted on this, there is still little knowledge of drivers' reaction to information provision. Since the deployment of ATIS technologies is still in an early state, practical experiences are limited. To gain more insights into travellers' decision making, in-laboratory experiments (for instance [2], [3], [4]) have been proposed. Behavioural models derived from the results of these laboratory experiments can be used in simulations to evaluate ATIS technologies. Travel time savings have been observed in several studies [5], [6], [7], varying from three to 30 percent depending on market penetration and network topology. These studies also indicate that the potential benefit of ATIS is little for regularly occurring congestion patterns, e.g., morning and evening peak-hours, which are in the literature referred to as recurrent congestion. If congestion occurs regularly, travellers start to adapt their travel behaviour. If congestion is caused by external shocks such as accidents or bad weather, the benefit of ATIS is assumed to be much greater. Congestion of this type is referred to as non-recurrent congestion.

Johannes Illenberger (corresponding author) is with the group of Transport Systems Planning and Transport Telematics, Institute for Sea- and Land-Transport, Technical University Berlin, Germany, Salzufer 17-19, 10587 Berlin, +49 30 31478793, illenberger@vsp.tu-berlin.de

Gunnar Flötteröd is with the group of Transport Systems Planning and Transport Telematics, Institute for Sea- and Land-Transport, Technical University Berlin, Germany, Salzufer 17-19, 10587 Berlin, +49 30 31429520, floetteroed@vsp.tu-berlin.de

Kai Nagel is the head of the group of Transport Systems Planning and Transport Telematics, Institute for Sea- and Land-Transport, Technical University Berlin, Germany, Salzufer 17-19, 10587 Berlin, +49 30 31423308, nagel@vsp.tu-berlin.de

The main reason why the benefit of ATIS is greater in case of non-recurrent congestion is that in such a situation the travel time is hard to predict for travellers. The key aspect is that the decision-maker is now faced with uncertainty [8]. There are several sources that can cause uncertainty: (i) variations on the supply-side caused by accidents, road maintenance or failures of transportation systems, (ii) variations on the demand-side, for instance caused by a sporting event, or travellers' misperception of travel times. However, there is little knowledge how travellers perceive uncertainty in travel time. Policies to improve the reliability of travel time can be of different type. For example, quick-response teams to clear up accidents or better organisation of road maintenance may improve reliability on the supply side, while on the demand side ATIS may help travellers to improve trip planning.

Unreliability of travel times results in added costs to the travel costs. Noland et al. [9] distinguish between two types of costs caused by travel time variations. The first, called *expected scheduling costs*, describes the costs caused by the attempt to lower the probability of arriving at an inconvenient time. Travellers may choose to depart earlier in order to be sure to arrive in time or they may choose a longer route with less uncertainty. The second, called *planning costs*, is the "pure nuisance of not being able to plan one's activities precisely because of uncertainty about when a trip will be completed". However, Noland himself shows that the planning costs are negligible.

In this paper, we present simulation studies, where we investigate the potential benefit of route guidance information in a system with risk-averse travellers. We set up a scenario with two routes, where one route has a reliable travel time and the other route an unreliable travel time. Travellers are repeatedly faced with risky route-choice decisions. The fraction of guided drivers and the risk-averseness is varied to investigate the effect of the guidance on risk-neutral and risk-averse systems.

Section II discusses related work on decision-making under uncertainty and travel time reliability. In Sec. III, the concepts of risk-aversion from psychology and economy are introduced and translated to applications in transport. The simulation framework used for the simulation studies is presented in Sec. IV, and in Sec. V the scenario setup is explained. Details of the behavioural model are discussed in Sec. VI, and the results of the simulation studies are presented in Sec. VII. We close the paper with a discussion in Sec. VIII and a summary of the results in Sec. IX.

II. Related Work

Existing behavioural models that account for travel time variability, in terms of departure time choice as well as route choice, can be roughly grouped into three approaches: (i) the "safety-margin" approach, (ii) the "mean-variance" approach and (iii) models that make explicit use of a concave or convex utility function to represent risk-averse or risk-loving behaviour.

Travel time variability can be modelled as an additional cost term in a utility function. For example, the idea of a travel time addend is embodied in the concept of a "safety margin" travellers generate by departing earlier than they would do without travel time variability [10].

An other approach is to capture the disutility of variability by cost terms for early or late arrival. Small [11] originally specifies a model of scheduling choice by $U = \alpha \cdot T + \beta \cdot SDE + \gamma \cdot SDL + \theta \cdot D_L$ where utility U is a function of travel time T, either so-called schedule delay-early SDE, or schedule delay-late SDL, and a fixed penalty for any late arrival. SDE and SDL are defined as the amount of time one arrives at a destination earlier or later than desired, respectively. D_L equals 1 if $SDL > 0$ and 0 otherwise. Coefficients α, β and γ are the costs per minute for travel time, early and late arrival, and θ is a discrete lateness penalty. Small finds that travellers prefer early to late arrival, prefer additional travel time to late arrival, and prefer early arrival to additional travel time. Thus the relative values of the coefficients are $\beta > \alpha > \gamma$, which already captures risk-averse behaviour, i.e. travellers would depart earlier or travel longer in order to avoid the risk of being late.

While in the upper model the travel time is assumed to be known to the decision-maker, Noland and Small [12], [13] extend the above model to a simple expected utility model by explicitly including travel time variability such that $\langle U \rangle = \alpha \langle T \rangle + \beta \langle SDE \rangle + \gamma \langle SDL \rangle + \theta P_L$. Here the expected utility is dependent on expected (or mean) travel time, either expected schedule delay-early, or expected schedule delay-late and the probability of late arrival P_L. Noland and Small use an exponential distribution of travel times to evaluate this model.

Ettema and Timmermans [14] extend the expected utility model of Noland and Small to account for perception errors with respect to the mean and variance of the travel time distribution and the confidence level of travel time information. Beside the misperception of the structural variation in travel time, they also discuss the potential benefit for travellers if they are provided with day-specific travel time information. Ettema and Timmermans conclude that the provision of information leads to a significant reduction of scheduling costs, amounting up to one Euro per trip, whereas the quality of information and the misperception of the quality have only a minor effect.

An essential aspect that determines the benefit of information provision is the travellers' knowledge about regularities in travel conditions such as weather conditions, day of the week or big events. Travellers without any prior knowledge of those traffic conditions benefit the most from information provision while experienced travellers benefit less.

To describe risk-averse and risk-prone behaviour, Chen et al. [15] use the exponential functions $U(t) = -a_1 \left(e^{a_2 t} - 1\right)$ for risk-averse travellers and $U(t) = -b_1 \left(1 - e^{-b_2 t}\right)$ for risk-loving travellers where a_1, a_2, b_1 and b_2 are parameters to be calibrated.

Like Chen, de Palma and Picard [16] and more recently de Palma and Marchal [17] focus on route-choice models dealing with uncertainty. De Palma and Picard [16] collect data via telephone interviews and estimate risk aversion through an ordered probit model. Several utility functions are discussed, such as (i) the "mean-standard formulation" $\langle U \rangle = -\langle t \rangle - \theta^S \sigma$, where t denotes the travel time and σ the travel time variance; (ii) the "mean-variance formulation" $\langle U \rangle = -\langle t \rangle - \theta^V \sigma^2$; (iii) the "constant relative risk aversion" (CRRA) $\langle U \rangle = \left\langle -\left(t^{1+\theta^R}\right)\left(1+\theta^R\right)^{-1}\right\rangle$, and (vi) the "constant absolute risk aversion" (CARA) $\langle U \rangle = \left\langle \left(1 - e^{\theta^A t t}\right)\left(\theta^A\right)^{-1}\right\rangle$. Parameters θ^S, θ^V, θ^R and θ^A represent the risk aversion for the corresponding utility function. De Palma and Picard show that absolute risk aversion θ^A is close to being constant, i.e., independent of the travel time.

The model proposed here uses a generic utility function to evaluate travel legs. The utility is dependent on travel time and on the time a traveller arrives later than desired. The risk-averseness is implicitly controlled by the marginal utility of arriving late. If the marginal utility for arriving late is negative, the behaviour becomes risk-averse, a marginal utility of zero denotes a risk-neutral situation, and a marginal utility greater than zero represents a risk-loving traveller. However, the latter case is not considered in this paper.

III. Concepts of Risk-Aversion

A. Risk-Aversion in Psychology and Economy

The concept of risk-aversion first arose in the 18th century in the context of the St. Petersburg Paradox. The St. Petersburg Paradox describes a lottery game with an infinite expected payoff. Naïve decision theory taking only the expected value into account would recommend to pay an infinite stake to enter the game since the expected outcome is also infinite. A course of action that appears implausible for a rational person.The paradox can be resolved by introducing a non-linear utility function for the payoff. The resolution mainly promoted by Daniel Bernoulli [18] is based on the idea that individuals do not estimate money in proportion to its quantity but rather in proportion to the usage they make of it. His solution involves a utility function of diminishing marginal utility: the logarithmic utility $U(x) = \ln(x)$. Now, the outcome of the lottery game becomes finite. The particular case of a concave utility function translates to the assumption that individuals are averse to risk in that the individual utility of the lottery diminishes when the stakes are increased.

The ideas of Bernoulli have been seized by psychologists and economists in the 20th century and have been formulated in the expected utility hypothesis [19] and later in the prospect theory [20] and the cumulative prospect theory [21]. According to the expected utility hypothesis, the utility perception of an individual facing uncertainty, e.g., a lottery of uncertain outcome, is calculated by considering the utility in each possible state and weighting it with the individual's estimate of the occurrence probability [22], [23], [24]. Figure 1 visualises the problem. Let z be a random variable which can take on the two discrete values z_1 and z_2. Let p be the probability that z_1 happens and $(1-p)$ that z_2 happens. One can regard this as a lottery game with two possible outcomes. Consequently, the expected outcome is $\langle z \rangle = pz_1 + (1-p)\, z_2$. Let $U(z)$ be a strictly concave utility function, which means, the marginal utility of the outcome is diminishing. The expected utility is now $\langle U \rangle = pU(z_1) + (1-p)U(z_2)$.

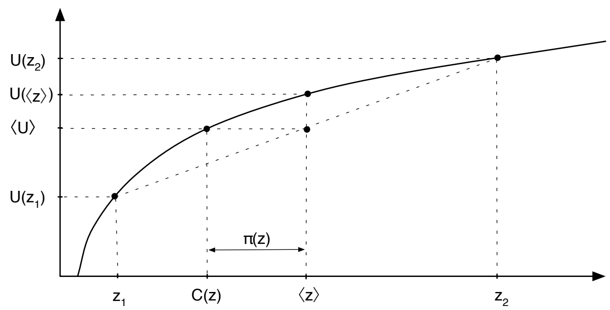

Figure 1. Expected utility theory with log utility function.

As long as the utility function is concave, Jensen's inequality [25] implies that the expected utility is not larger than the utility of the expected outcome:

$$\begin{aligned} \langle U \rangle &= pU(z_1) + (1-p)U(z_2) \\ &\leq U(pz_1 + (1-p)\, z_2) = U(\langle z \rangle) \end{aligned} \quad (1)$$

This represents the utility-decreasing aspect of risk-bearing. One can think of a player facing two lotteries. The risky lottery pays z_1 or z_2 with probabilities p and $1-p$ respectively, while the safe lottery pays $z^* = \langle z \rangle$ for sure. Although the expected outcome in both lotteries is the same, a risk-averse player would prefer z^* with certainty than $\langle z \rangle$ with uncertainty. This is what is captured in the inequality $\langle U \rangle \leq U(\langle z \rangle)$.

Consider now a third lottery which yields in the outcome $C(z)$ with certainty. As depicted in Fig. 1, the utility of this allocation is equal to the expected utility of the random prospect, i.e., $u(C(z)) = \langle U \rangle$. $C(z)$ is known as the certainty equivalent lottery, i.e., the sure-thing lottery which yields in the same utility as the random lottery. Although the outcome $C(z)$ is less than the expected outcome $\langle z \rangle$, a player would be indifferent between $C(z)$ for sure and $\langle z \rangle$ with uncertainty. The difference $\pi(z) = \langle z \rangle - C(z)$ is known as the risk-premium, i.e., the maximum amount of outcome a player is willing to forgo in order to avoid an allocation with risk.

More generally, let $U(z)$ be an elementary utility function, z a random variable, $\langle z \rangle$ the expectation of z and $C_U(z)$ the certainty equivalent, where the subscript U indicates that the certainty equivalent is dependent on the utility function $U(z)$. We define

- *Risk-Aversion* if $C_U(z) < \langle z \rangle$, i.e., $U(z)$ is concave,
- *Risk-Neutrality* if $C_U(z) = \langle z \rangle$, i.e., $U(z)$ is linear, and
- *Risk-Proclivity* if $C_U(z) > \langle z \rangle$, i.e., $U(z)$ is convex.

B. Risk-Aversion in Transport

The concept of risk-aversion can also be applied to transport. The random variable is now the uncertain travel time of a route. For the purpose of our studies we do not model risk-aversion explicitly with a continuous concave utility function but indirectly as depicted in Fig. 2. The disutility for travel is now linear in time. However, the traveller has a desired arrival time or a maximum amount of available travel time budget. If she arrives late or exceeds the maximum amount of travel times she gains an extra penalty.

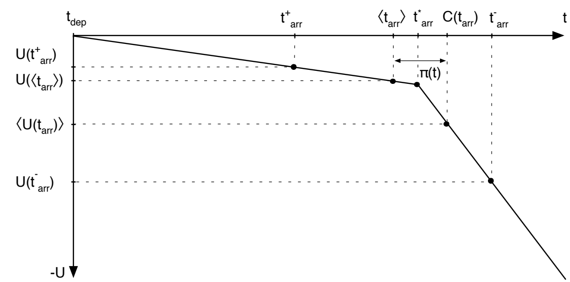

Figure 2. Utility function of a risk-averse traveller.

Consider a driver travelling a route with uncertain travel time. The driver always departs at t_{dep} and arrives on good days at t^+_{arr} and on bad days at t^-_{arr}, while $t^+_{arr} < t^-_{arr}$. The desired arrival time is t^*_{arr}. The expected arrival time is $\langle t_{arr} \rangle < t^*_{arr}$, thus at average the driver can expect to arrive in time. However, since arriving late on bad days causes an extra penalty, the driver will choose to travel an alternative route as long as the alternative route has a guaranteed travel time such that $\langle t_{arr} \rangle < t_{arr} \leq C(t_{arr})$. The absolute difference $\pi(t_{arr}) = \langle t_{arr} \rangle - C(t_{arr})$ is the additional amount of travel time a risk-averse traveller needs to "pay" in order to eliminate risk. At $C(t_{arr})$ the traveller is indifferent between the uncertain or alternative safe route.

In the context of advanced travellers information systems the certainty equivalent $C(t_{arr})$ becomes a relevant value to determine the users' willingness-to-pay for such services. Consider a traffic management centre (TMC) that provides real-time traffic information to drivers. If the TMC can guarantee a certain travel time for the uncertain route, then it can charge a monetary equivalent of the difference of the users' certainty equivalent and the guaranteed travel time. In this particular case a traveller would be indifferent between "paying" the difference in the form of additional travel time by travelling the alternative route or paying the fee for the TMC and having guaranteed travel times on the uncertain route. However, this presumes that the TMC knows the user's aversion to risk which is dependent on individual preferences or trip purpose and might not be easy to capture.

IV. Simulation Framework

A. The MATSim framework

For the studies in this paper we use the MATSim framework [26]. MATSim stands for "Multi-Agent Transport Simulation" and falls into the area of activity-based demand generation.

A complete discussion of the MATSim framework would go beyond the scope of this paper, hence we focus on the packages required for this work. The interested reader is referred to [27] for a detailed description.

In a multi-agent transport simulation each traveller is modelled individually as a so-called agent. An agent is an autonomous microscopic element with its own intentions, preferences, strategies and an explicit model of the decision making process. The last point is rather important, since it is not the vehicle that produces traffic, it is the person who drives it. Furthermore, a person does not only produce traffic, it tries to manage its day and the travel is just one action of the whole day. In MATSim an agent's intention is represented by a so called plan. A plan can be regarded as the agent's intended schedule for a day. It contains activities and travel legs connecting the activities. For the following sections the terms "agent", "traveller", "driver", "user", "individual" and similar are treated in a unified way and always denote the microscopic object that represents the traveller.

MATSim generates demand in form of a synthetic population. The population is a random realisation of census data, i.e., if one in turn takes a census from the synthetic population, it would approximately return the original census. A population can contain up to eight million agents with unique attributes such as age, gender, income or car availability.

The demand generation in MATSim is an iterative process where agents can successively adapt their plans through a evolutionary algorithm. At the beginning of an iteration each agent selects one plan form its plan database. The plan database can hold several plans per agent and represents the memory of a person. Plans in the database can be either previously executed plans or new excogitated options that are to be tested. Which plan will be selected depends on a specific selection rule that is part of the behavioural model and can vary for different scenarios. The rule used here will be described in Sec. VI-B.

The second step of an iteration is the so-called "network loading". All selected plans are run simultaneously through the physical simulation, which is in this case the physical simulation of traffic flow. One can regard this as the interaction of the agent with its environment and other agents where it collects sensory input about its experiences. A queueing model is used for the traffic simulation [28].

A crucial point is the feedback from the physical simulation to the behavioural model. The evolutionary algorithm requires a fitness function to evaluate each plan. The fitness function is realised through a function that calculates the economic utility of each plan measured in monetary units [€]. After each execution of the physical simulation the utility of the selected plans are calculated based on the feedback from the network loading. The feedback is given in form of departure, arrival and link travel times. Selected plans that have been already executed in a previous iteration are re-evaluated and the new calculated utility is add to the previous utility through exponential smoothing. The smoothing factor represents the learning rate with which the agents adapt their plans. Detail of the utility function will be discussed in Sec. VI-A.

With an utility associated to each plan the agent can now drop bad plans from its memory or revise its intentions by creating new mutations from good plans. New mutations incorporate modifications to routes and activity scheduling.

The physical simulation and the evolutionary algorithm are run repeatedly in alternate order for several iterations until the system reaches a desired state. In most cases the behavioural module is setup in such a way that the system preferably converges into Nash equilibrium. However, different setups are possible which do not necessarily show converging behaviour.

To start the entire iterative process an initial set of plans is required. The initial plans can be either manually constructed for smaller test scenarios or in a more complex way by disaggregation census data for large-scale real-world applications.

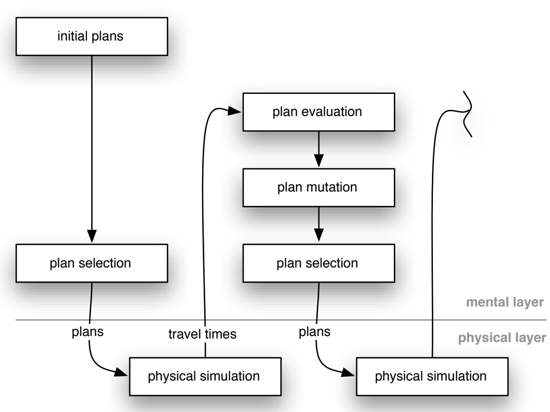

Figure 3. Iterative demand generation process in MATSim.

V. Simulation Scenario

To investigate the reaction of agents facing uncertain situation a system with a random variable is required. In this study the random variable is realised through a variable capacity of a single link in the road network. Consequently, the link travel times and the arrival times which are fed back from the physical simulation to the behavioural model vary from iteration to iteration.

A. Supply Side

For our simulation studies we create a simple road network with one "home" location and one "work" location connected by two different routes. The layout of the network is visualised in Fig. 4, link attributes are listed in Tab. I. The lower route which leads over link 3 and 5 is denoted as the safe route. The upper route which leads over link 2 and 4 is denoted as the risky route, since at has a variable capacity.

- The safe route has a fixed capacity of 7200 vehicles per hour and a free-flow travel time of 435 seconds.
- The risky route has a default capacity of 7200 vehicles per hour, however, an incident is simulated in iterations with even indices by reducing the capacity of link 2 by a factor of 0.3. The free-flow travel time of the risky route is 327 seconds, i.e., less than the safe route.

In the following, iterations with even indices are also referred as the bad day or the bad state of the nature, while iterations with uneven indices are also referred as the good day or the good state of the nature. Although, the capacity reduction of link 2 follows a predefined rule and is by no means random, at the beginning of an iteration the state of the nature is unknown to the agents. Thus the capacity of link 2 and likewise the travel time of the risky route is unpredictable for the travellers.

The storage capacity of link 2 is increased in such a way that even under high demand no spill-back into link 1 can ever occur. Link 7 acts as a return-link that allows agents to do round trips. The return-link is required for consistency but is not relevant for the effects demonstrated in this scenario. Flow and storage capacity of link 6 and 7 are chosen such high, so that agents travelling back to the home location experience always the free-flow travel time.

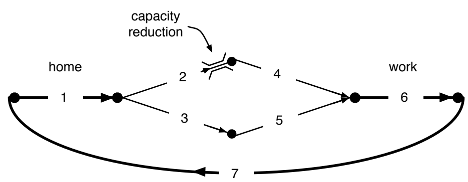

Figure 4. Road network with two routes connecting a home and a work location. The upper route has a variable capacity alternating between 7200 and 2160 veh/h.

no.	length [m]	free-speed [km/h]	capacity [veh/h]	lanes
1, 3, 5	2000	50	7200	4
2	2000	80	7200 / 2160	unlimited
4	2000	80	7200	4
6	2000	50	unlimited	unlimited
7	9656	50	unlimited	unlimited

Table I
LINK ATTRIBUTES.

B. Demand Side

The initial plans are created manually. The population counts 1000 agents, each equipped with two plans. All plans include one round trip, that means, the plan starts with a “home” activity at link 1, followed by a “work” activity at link 6 and again a “home” activity. The two plans of each agent differ in the first travel leg from the “home” to the “work” activity. One plan contains a leg where the agent intends to travel the safe route and one plan where the agent intends to travel the risky route. The travel leg from the “work” activity back to the “home” activity is again only for consistency. According to the route of the first travel leg the plans will be denoted as the risky plan and the safe plan respectively.

Departure times of the first travel leg are chosen in such a way that in every second two agents depart at the “home” location starting around 05:50. Consequently, under unconstrained conditions, i.e., no capacity reduction on link 2, no congestion occurs even if all users travel the same route.

VI. BEHAVIOURAL MODEL

The behavioural model focuses on the selection of plans for execution in the physical simulation and the evaluation of selected plans based on the feedback from the physical simulation. This means an agent’s memory consists always of the safe and risky plan. The agent chooses the safe or risky plan for execution and then re-evaluates the plan with the experienced travel times.

A. Evaluation

In order to compare plans a quantitative dimension is required to be assigned to each plan. A simplified version of the original MATSim utility function [29] is used here and is related to the Vickrey bottleneck model [30], [31]. Basically, the utility function only evaluates travel legs.

After each run of the physical simulation the utility of a selected plan’s execution is calculated. The total utility is composed of the sum of each travel leg’s utility. The utility of a travel leg is again composed of the (negative) utility for travelling U_{trav} and an extra penalty for being late U_{late}:

$$U = \sum_i U_{i,trav} + U_{i,late}. \tag{2}$$

The utility for travelling is assumed to be linear in time

$$U_{trav}\left(t_{trav}\right) = \beta_{trav} t_{trav} \tag{3}$$

where β_{trav} denotes the marginal utility for travel [€/h] and t_{trav} the time spent travelling [s]. The extra penalty for being late equals zero as long as the agent arrives before or at the desired arrival time t^*_{arr}.

$$U_{late} = \begin{cases} \beta_{late}\left(t_{arr} - t^*_{arr}\right) & \text{if } t_{arr} > t^*_{arr} \\ 0 & \text{else} \end{cases} \tag{4}$$

where β_{late} is the marginal utility for being late and t_{arr} the experienced arrival time. Parameter β_{late} controls the risk-aversion of agents. Values less than zero make agents averse to risk, while $\beta_{late} = 0$ represents risk-neutral users. One could even make agents risk-loving by choosing $\beta_{late} > 0$, i.e., agents are rewarded if they are late. However, risk-loving system are not considered here.

The final utility of a plan after iteration i is calculated out of the utility of the recent execution and the utility of the plan in the previous iteration through exponential smoothing:

$$U_i = \alpha U_i + (1-\alpha) U_{i-1}. \tag{5}$$

The smoothing factor α represents the learning rate of agents. If α is set close to 1 agents learn fast since the influence of the current iteration is high. Whereas, if α is set close to 0, agents slowly adapt the utility of the current iteration.

B. Plan selection

Before the run of the physical simulation each agent selects one plan to be executed. The selection rule is modelled through the probability of a plan transition from the currently selected plan to the alternative plan:

$$\frac{p_i}{p_j} = \frac{\kappa + e^{\gamma(U_i - U_j)}}{\kappa + e^{\gamma(U_j - U_i)}} \tag{6}$$

where p_i is the selection probability of the currently selected plan, p_j the selection probability of the alternative plan, κ a non-negative parameter representing explorative behaviour, γ a parameter controlling the rationality of the agent's decision, U_i and U_j the utility for the currently selected and alternative plan respectively. The above formula comprises two aspects: If κ is set to zero Eq. 6 results in a logit model:

$$\frac{p_i}{p_j} = e^{2\gamma(U_i - U_j)}. \tag{7}$$

Coefficient γ controls the rationality, i.e, increasing values for γ represent increasing rationality. Parameter κ introduces an explorative component to the behavioural model. Increasing κ leads to less influence of the logit model, i.e., sufficient great values for κ result in an equal distribution of the selection probabilities. It is required that agents are forced from time to time to select the alternative plan and to "renew" the plan's utility. Otherwise, the danger exists that an agent will be stuck in one plan. If, for instance, the risky plans is executed once on a bad day it is evaluated with a low utility. In the next iteration the safe plan is executed and will gain a better score. If the utility difference between the risky and safe plan is sufficient high, there is a low probability that the logit model will ever select the risky plan again. If the agent is forced to select the risky plan again, controlled by κ, there exists the probability that the plan is executed on a good day and gains a better utility. Consequently, the probability that the logit model induces a plan transition increases.

C. Guidance

A certain fraction f of agents can be equipped with an intelligent in-vehicle device. One can regard this device as a PDA which is supplied with link travel time information from a global traffic management centre and generates route recommendations. If an agent is equipped with such a device it will request it for the fastest route every time it reaches the fork of the safe and risky route, i.e., at the end of link 1. The route recommendations are based on reactive travel times (also denoted as naïve or instantaneous travel times) and thus can be regarded as a defensive estimation. The reactive travel time of every link at any point in time is given by counting the in- and outgoing vehicles and estimating the travel time based on the number if vehicles currently on the link and the current outflow. Travellers equipped with an in-vehicle device are denoted as "guided" agents and always comply with the guidance.

VII. Simulation Results

A. Parameter Setup

In the following simulation studies the effects of the parameters β_{late} and f are investigated. Parameter β_{late}, denoting the penalty for being late, controls the risk-aversion of the agents. The values are varied form 0 €/h (risk-neutral) to –100 €/h (risk-averse). Parameter f represents the fraction of agents equipped with guidance devices and is varied from 0 (no equipped agents) to 1 (all agents equipped). All remaining parameters are set to fixed values, which are

- *Plan evaluation*: The marginal utility for travel β_{trav} is set to -6 €/h and the desired arrival time t^*_{arr} is uniformly set to 6:00 for all agents. The learning rate α (smoothing factor) is set 0.2, i.e., slow learning.
- *Plan selection*: Parameter γ which controls the agent's rationality is set 5 and κ which controlls the explorative behaviour is set to 2.

Simulation runs are conducted with 1000 iterations, while the system requires roughly the first 150 iteration to stabilise. Measurings over a simulation run are averaged over iteration 200 to 1000 so that the transient phase does not influence the averages.

B. Simulations without guidance

Figure 5a depicts the base-case with $\beta_{late} = 0$ €/h and $f = 0$. In the initial plans the risky plan is manually marked as selected, thus all agents travel the risky route in the 0-th iteration. After approximately 100 iterations travellers are distributed almost equally over both routes. For the remaining iterations the distribution of users fluctuates around the equilibrium. Averaged over iteration 200 to 1000 the travel time on the risky route is slightly faster (417 s) compared to the safe route (435 s). Accordingly, there are two users more on the risky route (501:499).

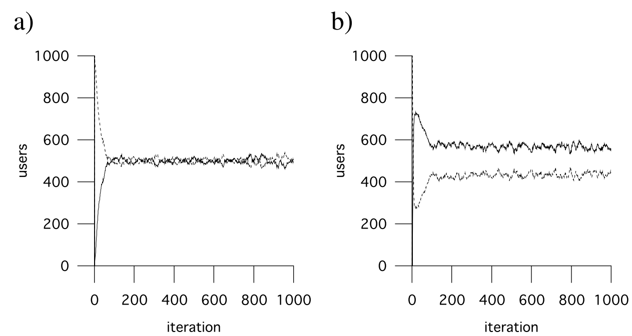

Figure 5. Number of users on the safe route (solid line) and on the risky route (dotted line) over 1000 iterations. a) risk-neutral system with $\beta_{late} = 0$ €/h, b) risk-averse system with $\beta_{late} = -100$ €/h.

In Fig. 5b β_{late} is set to -100 €/h, i.e., agents are now risk-averse. Agents that arrive after their desired arrival time t^*_{arr} = 6:00 gain an extra penalty. As a consequence the distribution of users over both routes shifts to the safe route. Figure 6a shows how the load reacts if the value for β_{late} is successively decreased from 0 €/h to -100 €/h. As the system is turned more risk-averse, more agents switch to the safe route. As a further consequence, the travel time on the risky route on bad days decreases, and averaged over good and bad days the travel time decreases from 417 s to 405 s. On the one hand, the decreasing β_{late} pushes the system towards the safe route, but on the other hand, the decreasing travel time on the risky route partially counteracts the effect.

A switch from the risky route to the safe route is beneficial if the travel time on the safe route is less than the the certainty equivalent travel time $C\,(t_{arr})$ (Fig. 2) on the risky route. However, this situation only arises for agents that arrive early ($t_{arr} < t^*_{arr}$) on good days and arrive late ($t_{arr} > t^*_{arr}$) on bad days if they travel the risky route. Agents that always arrive early, independent of the state of the nature, or always

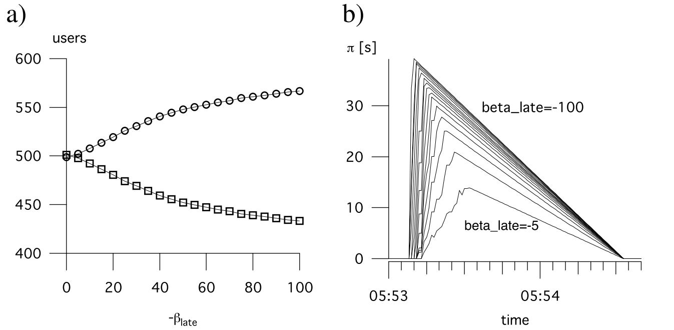

Figure 6. a) Number of users on the safe route (circles) and on the risky route (squares) in dependency of β_{late}, b) Values for $\pi\left(t_{arr}^{risky}\right)$ in dependency of the departure time. The lower curve represents $\beta_{late} = -5$ €/h, the upper curves represent β_{late} increased in -5 €/h steps.

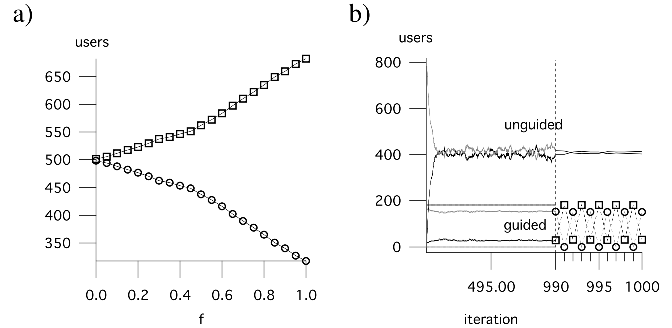

Figure 7. a) Number of users on the safe route (circles) and on the risky route (squares) for $\beta_{late} = 0$ €/h. b) Number of unguided users on the safe route (black solid line) and on the risky route (grey solid line); number of guided users on the safe route (black points/squares) and on the risky route (grey points/black circles). $f = 0.2$. Note the change of the scaling at iteration 990.

arrive late respectively, do not react to a change of β_{late}. For the latter agents $U\left(\langle t_{arr}\rangle\right)$ equals $\langle U\left(t_{arr}\right)\rangle$, while for the former $U\left(\langle t_{arr}\rangle\right)$ is greater than $\langle U\left(t_{arr}\right)\rangle$. On can visualise this effect, if, for instance, one plots the absolute risk-premium $\pi\left(t_{arr}^{risky}\right) = \left|\langle t_{arr}^{risky}\rangle - C\left(t_{arr}^{risky}\right)\right|$ in dependence of the departure time. Figure 6b shows a peak between 5:53:08 and 5:54:34 which height is dependent on β_{late}. Agents departing before 5:53:08 are always early and agents departing after 5:54:34 are always late. For both groups $\pi\left(t_{arr}^{risky}\right)$ is always 0 s. Only the agents departing in between react to variations of β_{late} and travel an additional travel time of $\pi\left(t_{arr}^{risky}\right)$ in order to avoid risk. – This maybe slightly counter-intuitive behavior is due to the fact that all the risk-modelling convexity in the utility function is contained in the bend between early and late arrival. Once an agent operates entirely on one of the two branches, no reaction to risk is left.

C. Simulation with guidance

To investigate the effects of the guidance the fraction f of equipped users is varied from 0 to 1. Figure 7a depicts the distribution of users over the safe and risky route with $\beta_{late} = 0$ €/h. As expected, at average the distribution shifts to the risky route with increased equipment fraction. Details of a simulation run with $f = 0.2$ are shown in Fig. 7b. The upper two curves represent the unguided users, the black line for the unguided users travelling the risky route and the grey line for the unguided users travelling the safe route. The unguided users split up equally over both routes, however, since there are less unguided users compared to the base-case the focal point shifts to about 400 users. The lower points represent the guided users, grey points for those who travel the safe route and black point for those who travel the risky route. The guided users oscillate between both routes. On good days all guided users (approximately 180 users) travel the risky route and on bad days the majority uses the safe route (150 users) and some agents (30 users) remain on the risky route.

Figure 8 compares the average utility for the guided and unguided agents. In case of $\beta_{late} = 0$ €/h the complete system benefits from the guidance, while the guided users benefit more compared to the unguided users. The utility for guided travellers is about -2.03 € and slightly decreases to -2.032 € with increasing equipment fraction. For unguided users the utility increases from -2.11 € at $f = 0$ and levels off at about -2.075 € at and equipment fraction of $f = 0.6$. On the one hand, guided users benefit form the guidance, while the individual utility decreases as more agents are equipped with the guidance. One could say that the value of the travel time information decreases with the number of users that have accesses to this information. On the other hand, unguided users also benefit if other users are guided, since less travellers on the risky route on a bad day is also advantageous for unguided users. However, a saturation is reached at approximately $f = 0.6$ where the slope of utility curve for unguided users is close to zero and even tends to be negative for values $f > 0.8$ respectively.

The picture for $\beta_{late} = -100$ €/h is somewhat different. In contrast to Fig. 8b, guided users do also benefit from an increased equipment fraction, at least for values $f < 0.65$. Since greater values for β_{late} force more (unguided) users to travel the risky route, the utility gain for unguided users is also more pronounced, even for values $f > 0.6$.

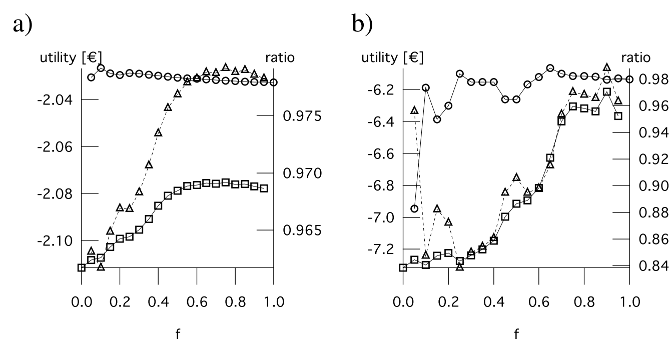

Figure 8. Utility for guided (circles) and unguided (squares) agents. Ratio of guided users' utility over unguided users' utility (triangles). a) $\beta_{late} = 0$ €/h, b) $\beta_{late} = -100$ €/h.

To compare the utility of guided and unguided agents in risk-neutral and risk-averse systems one cannot use the absolute difference, since plans are evaluated with different values for β_{late}. Therefore the ratio of the guided users' utility over the unguided users' utility is used (Fig. 8). For $\beta_{late} = 0$ €/h the ratio lays at 0.963 and increases to 0.978 at the saturation point $f = 0.6$, i.e., the guided users' travel disutility is with a factor of 0.978 less compared to the unguided users' travel disutility. In Fig. 8b the range of the utility ratio is much greater. For equipment fraction of $f < 0.4$ the ratio is at average 0.86 and increases to values around 0.97 for $f > 0.8$. Hence, at low equipment fraction guided users benefit more in a risk-averse system compared to a risk-neutral system.

On can calculate the value $\pi_g(t_{arr}) = \left|t_{arr}^{guided} - C\left(t_{arr}^{risky}\right)\right|$ which resembles the risk-premium, where t_{arr}^{guided} denotes the average arrival time if an agent travels the route recommended by the guidance and $C\left(t_{arr}^{risky}\right)$ denotes the average certainty equivalent of the risky route. The value $\pi_g(t_{arr})$ represents the expected travel time an agent saves if it travels the guided route instead of the risky route. The monetary equivalent of $\pi_g(t_{arr})$ is the maximum amount of money an agent is willing to pay for the guidance, which is $\Delta U = \left|\left\langle U^{risky}\right\rangle - U^{guided}\right|$, where $\left\langle U^{risky}\right\rangle$ is the average utility of the risky route and U^{guided} the average utility of the guided route. Figure 9a depicts π_g as a surface plot over β_{late} and f. For low equipment fraction the expected travel time savings are about 43 s. In the extreme cases with low equipment fraction ($f < 0.1$) and high equipment fraction ($f > 0.8$) the expected travel time savings are nearly independent of the risk-aversion. However, the willingness to pay (Fig. 9b) for low equipment fraction varies between 0.07 € ($\beta_{late} = 0$ €/h) and 1.2 € ($\beta_{late} = -100$ €/h), i.e., even if risk-averse agents gain the same travel time savings they are willing to pay much more compared to risk-neutral travellers. If the equipment fraction f is increased travel time savings and the willingness to pay for the guidance decreases. At high equipment fractions ($f > 0.8$) the situation is similar to the travel time savings, i.e., the willingness to pay increases only slightly with increasing risk-aversion.

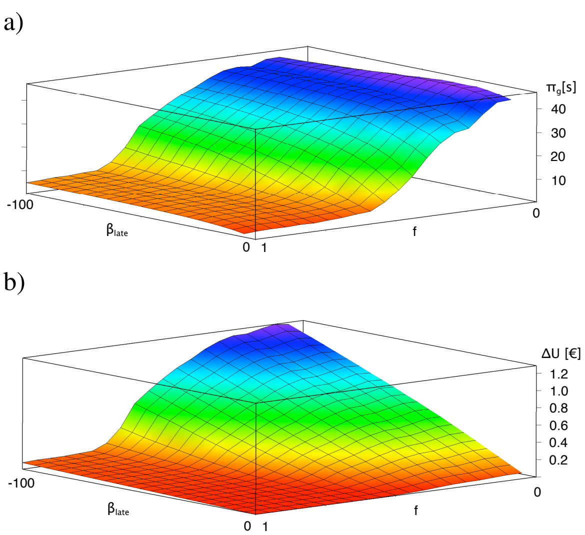

Figure 9. a) The expected travel time an agent saves if it travels the guided route instead of the risky route. b) The willingness to pay for the guidance of a traveller intending to travel the risky route.

VIII. Discussion

In this work only route-choice is addressed. Departure time choice is completely neglected, albeit one may argue that changes in departure time choice occur more frequently than changes in route. A traveller may feel more comfortable to stay with her habitual route and to depart earlier than travelling a new and probably unknown route. Extending the model to account for departure time choice, however, requires also an extension to the utility function. Beside the evaluation of travel time, the performing of activities is required to be monetarised, which introduces a further term to model risk sensitive behaviour. The utility function for activity performing can be either concave, linear or convex translating to risk-averse, risk-neutral or risk-loving behaviour respectively.

It would also be interesting to apply the prospect theory to such route-choice models, since expected utility theory has been criticised to violate state-preferences studies [20], [32]. Avineri and Prashker [8] suggest a paradigm shift concerning the connection between uncertainty in travel time and route-choice preferences. They found evidence of two violations of expected utility theory which are known as the certainty effect and the inflating of small probabilities. The certainty effect describes a situation where the desirability of a prospect is reduced if its character changes from a sure gain to an uncertain gain, even if its expectation is greater. The effect of inflating small probabilities states that individuals tend to underestimate medium outcomes with high occurrence probability and overestimate high outcomes with rare occurrence probability. If the probabilities of winning are substantial most people choose the prospect where winning is more probable. But if the probabilities of winning are miniscule in both prospects people choose the prospect that offers the larger gain.

Avineri and Prashker suggest that prospect theory may better capture the behaviour in route-choice. However, prospect theory was designed for single-choice situations and may fail to predict repeated route-choice decisions with feedback. Furthermore, formalisation and parametrisation needs additional research before prospect theory and its extension the cumulative prospect theory are applicable to choice modelling in transport.

IX. Summary

In this paper we presented simulation studies where travellers are faced with repeated risky route-choice decisions. The sensitivity of drivers towards risk and the equipment rate with route guidance devices is varied to investigate the potential benefit of such devices in a system with uncertainty. For the synthetic scenario of the paper, the following conclusions can be drawn:

- In a risk-neutral system ($\beta_{late} = 0$ €/h) the average disutility of travel for a guided traveller is about 4 % less compared to her disutility in a system without guidance.
- In a risk-averse system ($\beta_{late} = -100$ €/h) the average disutility of travel for a guided traveller is about 19 % less compared to her disutility in a system without guidance.

The model shows that risk-averse agents are willing to pay more for a guidance compared to risk-neutral agents even if the travel time savings are of same magnitude. The framework could be extended to realistic scenarios and realistic numbers for the utility function, and could then be used to perform willingness-to-pay simulations in real world situations.

X. Acknowledgement

This work has been funded by the German Research Foundation (DFG) within the project “State estimation for traffic simulations as coarse grained systems”.

References

[1] R. Arnott, A. de Palma, and R. Lindsey, “Does providing information to drivers reduce traffic congestion?” *Transportation Research A*, vol. 25, no. 5, pp. 309–318, 1991.

[2] J. L. Adler, W. Recker, and M. G. McNally, “Using interactive simulation to model driver behavior under ATIS,” in *Proceedings of the ASCE 4th International Conference on Microcomupters in Transportation*, 1992.

[3] ——, “A conflict model and interactive simulator (FASTCARS) for predicting en-route driver behaviour in response to real-time traffic condition information,” *Transportation*, vol. 20, no. 2, pp. 83–106, 1993.

[4] P. W. Bonsall and T. Parry, “Using an interactive route-choice simulator to investigate drivers’ compliance with route guidance advice,” *Transportation Research Record*, vol. 1306, pp. 59–68, 1991.

[5] J. L. Adler, V. J. Blue, and T. L. Wu, “Assessing network and driver benefits from bi-objective in-vehicle route guidance,” in *78th Transport Research Board Annual Meeting*, 1999.

[6] R. H. M. Emmerink, K. W. Axhausen, P. Nijkamp, and P. Rietveld, “Effects of information in road transport networks with recurrent congestion,” *Transportation*, vol. 22, pp. 21–53, 1995.

[7] K. E. Wunderlich, “An assessment of pre-trip and en-route ATIS benefits in a simulated regional urban network,” in *3rd World Congress in Intelligent Transport Systems, Intelligent Transportation: Realizing the Future*, 1996.

[8] E. Avineri and J. N. Prashker, “Sensitivity of travel time variability: Travelers’ learning perspective,” *Transportation Research C*, vol. 13, pp. 175–183, 2005.

[9] R. B. Noland, K. A. Small, P. M. Koskenoja, and X. Chu, “Simulating travel reliability,” The University of California Transporation Center, University of California at Berkeley, Working Paper UCTC No. 372, 1997.

[10] T. R. Knight, “An approach to the evaluation of changes in travel time unreliability: a "safety margin" hypothesis,” *Transportation*, vol. 3, pp. 393–408, 1974.

[11] K. A. Small, “The scheduling of consumer activities: Work trips,” *The American Economic Review*, vol. 72, no. 3, pp. 467–479, 1982.

[12] R. B. Noland and K. A. Small, “Travel-time uncertainty, departure time choice, and the costs of moring commutes,” *Transportation Research Record*, vol. 1493, pp. 150–158, 1995.

[13] R. B. Noland and J. W. Polak, “Travel time variability: a review of theoretical and empirical issues,” *Transport Reviews*, vol. 22, no. 1, pp. 39–54, 2002.

[14] D. Ettema and H. Timmermans, “Costs of travel time uncertainty and benefits of travel time information: Conceptual model and numerical examples,” *Transportation Research C*, vol. 14, pp. 335–350, 2006.

[15] A. Chen, Z. Ji, and W. Recker, “Travel time reliability with risk-sensitive travelers,” *Transportation Research Record*, vol. 1783, pp. 27–33, 2002.

[16] A. de Palma and N. Picard, “Route-choice decision under travel time uncertainty,” *Transportation Research A*, vol. 39, no. 4, pp. 295–324, 2004.

[17] A. de Palma and F. Marchal, “Measurement of uncertainty costs with dynamic traffic simulations,” Laboratory of Transportation Economics, Lyon, France, Tech. Rep., July 2007.

[18] D. Bernoulli, “Exposition of a new theory on the measurement of risk,” *Econometrica*, vol. 22, no. 1, pp. 23–36, 1954.

[19] J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*. Princton University Press, 1944.

[20] D. Kahneman and A. Tversky, “Prospect theory: An analysis of decision under risk.” *Econometrica*, vol. 47, pp. 263–291, 1979.

[21] A. Tversky and D. Kahneman, “Advances in prospect theory: Cumulative representation of uncertainty,” *Journal of Risk and Uncertainty*, vol. 5, pp. 297–323, 1992.

[22] P. J. H. Schoemaker, “The expected utility model: Its variants, purposes, evidence and limitations,” *Journal of Economic Literature*, vol. 20, no. 2, pp. 529–563, 1982.

[23] M. Friedman and L. J. Savage, “The utility analysis of choices involving risk,” *The Journal of Political Economy*, vol. 56, no. 4, pp. 279–304, 1948.

[24] M. J. Machina, “Choice under uncertainty: Problems solved and unsolved,” *The Journal of Economic Perspectives*, vol. 1, no. 1, pp. 121–154, 1987.

[25] J. L. W. V. Jensen, “Sur les fonctions convexes et les inégalités entre les valeurs moyennes,” *Acta Mathematica*, vol. 30, no. 1, pp. 175–193, 1906.

[26] (2008, November). [Online]. Available: www.matsim.org

[27] M. Balmer, K. W. Axhausen, and K. Nagel, “A demand generation framework for large-scale micro simulations,” in *Transportation Research Board 85th Annual Meeting*, 2006.

[28] N. Cetin, A. Burri, and K. Nagel, “A large-scale agent-based traffic microsimulation based on queue model,” in *Proceedings of the Transportation Research Board Annual Meeting*, no. 03-4272, 2003.

[29] D. Charypar and K. Nagel, “Generating complete all-day activity plans with genetic algorithms,” *Transportation*, vol. 32, no. 4, pp. 369–397, 2005.

[30] W. S. Vickrey, “Congestion theory and transport investment,” *The American Economic Review*, vol. 59, no. 2, pp. 251–260, 1969.

[31] R. Arnott, A. D. Palma, and R. Lindsey, “A structural model of peak-period congestion: A traffic bottleneck with elastic demand,” *The American Economic Review*, vol. 83, no. 1, pp. 161–179, 1993.

[32] E. Avineri and J. N. Prashker, “Violations of the expected utility theory in route-choice stated preferences: the certainty effect and inflating of small probabilities,” *Transportation Research Record*, vol. 1894, pp. 222–229, 2004.