*Transportmetrica*, Vol. 1, No. 1 (2005), 13-22

# MODELLING LEARNING AND ADAPTATION IN TRANSPORTATION CONTEXTS

THEO ARENTZE[1] AND HARRY TIMMERMANS[2]

Received 9 February 2004; received in revised form 3 May 2004; accepted 25 May 2004

This paper reviews previous research on modelling learning and adaptation in transportation contexts. In addition, it provides a framework for modelling these processes in future research, addressing some issues and problems that to date have received not much attention in transportation research.

KEYWORDS: Activity analysis, learning and adaptation, beliefs, information gain

## 1. INTRODUCTION

Predicting transport demand has traditionally been one of the key concerns in transportation research. Over the years, many different mathematical models have been proposed in a variety of disciplines concerned with transportation to predict the demand for transport, the feasibility of new projects and plans, and the impact of a variety of transport policy measures. Although by its very nature, prediction concerns the future and the topic area at large requires an understanding of consumer response to changing transport conditions, the vast majority of transport demand models are cross-sectional in nature. The relationships established in one point in time are typically used to predict future behaviour. Relatively few models are concerned with the dynamics of behaviour. Moreover, unlike cross-sectional models, the dynamic models still have quite different theoretical underpinnings; an overriding conceptual framework and modelling approach still has not evolved in the literature.

The aim of this paper is to review the existing literature on dynamic models of learning and adaptation. As a framework, we use the distinction between models under assumptions of static knowledge, models under assumptions of learning, and models under assumptions of learning and perception of uncertainty. Models under assumptions of static knowledge do not explicitly take into account the fact that interaction with the environment generally leads to adaptation of rules of behaviour through learning. Instead, they typically assume either time-invariant behaviour or involve some component of structural change. In contrast, models under assumptions of learning do assume that individuals learn based on experiences. These models, however, may not take into account the impact of uncertainty on decision-making and, vice versa, the impact of learning on uncertainty.

In the remainder of this paper, we will discuss in more detail examples of models belonging to each of these classes and their dominant theoretical underpinnings. The paper will be completed by discussing some potentially relevant areas of future research.

[1] Eindhoven University of Technology, PO Box 513, 5600 MB Eindhoven, The Netherlands.

[2] Eindhoven University of Technology, PO Box 513, 5600 MB Eindhoven, The Netherlands. Corresponding author (E-mail: eirass@bwk.tue.nl).

## 2. ADAPTATION UNDER ASSUMPTIONS OF CERTAIN, STATIC KNOWLEDGE

The field of predicting transport demand has evolved from the four-stage modelling approach, through tour-based discrete choice models, to several activity-based approaches (e.g. Timmermans et al., 2002). These activity-based approaches predict which activities are conducted where, when, for how long, the transport mode and route involved, and sometimes some additional facets. Regardless of the specific approach, these activity-based models depict structural relationships in the data. When used for prediction or impact assessment, the policy concerned is defined in terms of the explanatory variables of the model, and keeping the estimated parameters or elicited rules fixed, the structural relationships are then used to predict changes in behaviour. There are only a few exceptions to this general approach.

One of the first exceptions is AMOS (Pendyla et al., 1995, 1997, 1998). Consumer response to policies is explicitly observed and these outcomes are then combined with data on activity-travel patterns using neural networks. A second example is an extension of the Albatross model (Arentze and Timmermans, 2000, 2003b). Albatross is the only currently fully operational computational process model of transport demand and consists of a large number of decision heuristics (rules) to predict activity-travel patterns and their characteristics. Conceptually similar to AMOS, it is assumed that individuals will exhibit a primary response, which refers to the choice of a specific strategy aimed at reducing the impact of the policy. In their case study of congestion pricing (Arentze et al., 2004), this is modelled as a choice between alternative ways of fully or partly reducing the increase in costs. The following set of response alternatives was identified: (i) no change; (ii) eliminate the trip by conducting the activity at home; (ii) eliminate the trip by skipping the activity; (iii) reduce the distance of the trip by conducting the activity closer to home; (iv) change of transport mode; (v) change the departure time of the trip, and (vi) change the route of the trip. The purpose of the model was to predict the probability of these strategies to reduce travel costs for an activity, using a multinomial logit model. A stated adaptation experiment was designed to measure the relative strengths of the primary response. The primary responses capture the way individuals adjust the basic features of their activity-travel pattern program. Essential to the approach is that the rules, embedded in the Albatross model, were used apart from re-scheduling the activities, to predict adapted activity-travel patterns under specific policy scenarios.

Although at the present stage of development still based on cross-sectional data, the Aurora model was explicitly developed to model short-term adjustment of activity-travel schedules. The original conceptualisation of the model is described in Timmermans et al. (2001). The model is concerned with the dynamics of activity scheduling and rescheduling decisions as a function of unexpected events during the execution of activity programs. The model allows for several behavioural principles and decision styles, including risk-avoiding and opportunistic behaviour. Results of numerical simulations suggested that the model predicts theoretically expected emergent behaviour (Joh et al., 2001a, b). Later, the original theory was elaborated and extended to include many different facets of activity rescheduling behaviour, including destination choice (Joh et al., 2002). Again, numerical simulations supported the face validity of the model. A recent empirical estimation of part of the model provided further support to its face validity (Joh et al., 2003).

# 3. ADAPTATION UNDER ASSUMPTIONS OF LEARNING

## *3.1 Non-stationary, uncertain environments*

An examination of the above models indicates that the estimated utility functions or the derived choice rules used to schedule activities are not based on process data of adaptation behaviour, but rather on outcome data of observed activity-travel patterns. All these models assume, either implicitly or explicitly, that individuals have static preferences about the facets of the transportation system, the urban environment and the institutional context that influence their choice behaviour. The error component of random utility models is typically interpreted in terms of heterogeneity or stochastic preferences and not in terms of non-stationarity. Moreover, there are no explicit assumptions of how individuals adapt the rules of their behaviour through learning. When the focus of the model is on predicting short-term change, this may not be an important issue, but in the context of predicting mid or long-term adjustment behaviour, such as in the context of transport demand management, it may be.

Arentze and Timmermans (2003a) have recently put forward a conceptual framework for modelling learning and adaptation. This framework is especially relevant for non-stationary, and therefore uncertain, environments. They argue that in a dynamic sense choices are contingent upon the outcomes of previous choices. By repeatedly making decisions, an individual acquires knowledge (learns) about his environment and thereby forms expectations about attributes of the environment. Individuals may make different choices over time and thus learn which of these choices is most instrumental in realising particular goals. Individuals may forget about their experience in particular situations as a function of time and the characteristics of the event itself. Although their conceptual framework, based on reinforcement learning, focused on rewards and hence the valuation of choice alternatives, a similar structure can be assumed for updating the perception of transportation system attributes. Thus, in this context, adaptation and learning processes may involve two levels. One concerns individuals' perception of the real environment. Through repeated experiences, individuals will update their perception of attributes of the real world, which serves as a base for their choices. A second level concerns the strategies that individuals apply in their choice behaviour to find the best available action in a specific context.

Central to their conceptualisation is that an individual's choice leads to action $a$. Each action results in an outcome that can be positive or negative and these outcomes are incorporated in future choices. The expected outcome of an action will thus be based on previous experiences. Various assumptions can be formulated for estimating these expected outcomes. A natural way of estimating the expected outcome $Q_t(a)$ of action $a$ at time $t$ is by averaging the outcomes $r$ when the action was selected at previous choice occasions. Thus,

$$Q_t(a) = \frac{r_1 + r_2 + ... + r_{n_a}}{n_a} \tag{1}$$

where $n_a$ is the total number of times action a was selected before.

From a conceptual perspective, equation (1) would imply that individuals keep a memory of the outcomes of previous experiences, and that a memory trace is activated, and perhaps reinitialised at the next choice occasion. Alternatively, it can be assumed that individuals cognitively keep some summarising measure of past experiences, which is incrementally updated[3]. A general form of such an incremental update rule is:

$$Q_t(a) = Q_{t-1}(a) + \lambda[r_t - Q_{t-1}(a)] \tag{2}$$

$$Q_t(a) = Q_{t-1}(a) + \lambda r_t - \lambda Q_{t-1}(a) \tag{3}$$

$$Q_t(a) = \lambda r_t + (1-\lambda)Q_{t-1}(a) \tag{4}$$

This equation has become dominant in the transportation literature (Ben-Akiva et al., 1991; Koutsopoulos and Xu, 1993; Axhausen et al., 1995; Nakayama and Kitamura, 2000). Parameter $\lambda$ is typically interpreted as a measure of habit strength. If it is equal to zero, there is no further updating. Alternatively, if one assumes that the updating parameter is a constant, it follows that:

$$\lambda r_t + (1-\lambda)Q_{t-1}(a) = \lambda r_t + (1-\lambda)[\lambda(r_{t-1}) + (1-\lambda)Q_{t-2}(a)] \tag{5}$$

$$= \lambda r_t + (1-\lambda)\lambda r_{t-1} + (1-\lambda)^2 Q_{t-2}(a) \tag{6}$$

$$= \lambda r_t + (1-\lambda)\lambda r_{t-1} + (1-\lambda)^2 \lambda r_{t-2} + \ldots + (1-\lambda)^{t-1}\lambda r_1 + (1-\lambda)^t Q_0(a) \tag{7}$$

$$= (1-\lambda)^t Q_0(a) + \sum_{i=1}^{t} \lambda(1-\lambda)^{t-i} r_i \tag{8}$$

Thus:

$$Q_t(a) = (1-\lambda)^t Q_0(a) + \sum_{i=1}^{t} \lambda(1-\lambda)^{t-i} r_i \tag{9}$$

Hence, the expected outcome is a weighted average of the initial estimated outcome and the past outcomes. The weight given to past outcomes depends on how long ago it was experienced. More specifically, the weight given to outcome *i* decreases as the number of intervening outcomes increases. The weight decays exponentially and for this reason equation (9) is sometimes called an exponential recency-weighted average.

At least two different theoretical foundations can be given for this specification. An argument typically given in the reinforcement learning literature (Kaelbing et al., 1996) is while the averaging rule may be appropriate in a stationary environment, it makes more sense to weight more recent outcomes more heavily in a non-stationary environment as more recent experience may provide more reliable information. In addition, the specification can be linked to theory of memory retention (e.g. Anderson, 1983), in which case $\lambda$ is a memory decay parameter. Note that strictly speaking the meaning of time differs between these two lines of research. In reinforcement learning, time relates to the previous experiences, whereas in the theory of memory decay it depends on clock time.

It is assumed that individuals choose an action based on expected outcomes. This means that additional assumptions are required to link expected outcomes to the choice of actions. A simple assumption is to choose the action that maximises the expected outcome (assuming it is positively scaled, otherwise the assumption should be that the expected outcomes is minimised). Thus,

$$\Pr(a \mid t) = \begin{cases} 1 & \text{if } Q_t(a) = \max_{a'}[Q_t(a')] \\ 0 & \text{otherwise} \end{cases} \tag{10}$$

---

[3] The question whether individuals cognitively update their expectations (mean values) or store and retrieve information in their memory is a matter of research and may depend on the choice context and choice problem. This discussion goes beyond the scope of the present paper.

This so-called greedy rule, however, would imply that if the first choice is sub-optimal, an individual will never learn whether other actions will yield a better outcome. Alternative actions need to be sampled at least once or every once in a while. A simple alternative is to choose with small probability ε, an action at random, uniformly and independently of the expected outcomes. An advantage of this ε-greedy method is its simplicity. However, it may not represent the most valid choice strategy. The chance of choosing an action may be some monotonic function of their expected outcomes. For example a Gibbs distribution may be assumed, implying that action $a$ is chosen with probability:

$$\Pr(a \mid t) = \frac{e^{Q_t(a)/\tau}}{\sum_{a'} e^{Q_{t(a')}/\tau}} \tag{11}$$

where τ is a positive parameter. If τ approximates zero, equation (11) is equal to the greedy approach. If it approaches infinity, all actions become nearly equiprobable. Note that equation (11) is mathematically equivalent to the multinomial logit model. It means that τ can also be interpreted as the scale of the error term of the utility function. This overall framework was applied by Ettema et al. (2003a, 2003b) in studies of departure time and route choice behaviour.

## 4. ADAPTATION UNDER ASSUMPTIONS OF LEARNING AND PERCEPTION OF UNCERTAINTY

Although the above conceptual framework is potentially powerful in modelling reinforcement learning, the uncertainty itself is not taken explicitly into account in action selection. It is assumed that individuals make choices based on expected outcomes, without considering the relative amount of uncertainty related to these outcomes, although such extensions are possible (Ettema et al., 2003a, 2003b). There are several other approaches however to decision making under risk. The best known of these and also applied mostly to transportation problems is Kahneman's and Tversky's (cumulative) prospect theory.

Kahneman and Tversky (1979) showed that estimated utility functions and choice probabilities depend on how the choice alternatives are framed. Individual choice behavior differs, dependent on the framing of the outcomes of actions in terms of losses or gains. Moreover, they showed that the utility of risky lotteries is non-linear in outcome probabilities. Asymmetries cannot be explained in terms of classical utility functions and these anomalies led Kahneman and Tversky (1979) and Tversky and Kahneman (1992) to abandon the utility maximization model in favor of the more behaviorally realistic prospect theory.

Prospect theory assumes that lotteries are evaluated in a two-step process: an initial phase of editing and a subsequent phase of evaluation. In the editing phase, the lottery outcomes are coded as gains or losses relative to some reference point, which is usually the current asset position, but may be influenced by the presentation of the lottery or expectations of the decision-maker.

Prospect utility assumes that the uncertainties or expected gains and expected losses are defined. In an experimental setting, the researcher can frame these, but in the real world the uncertainties are influenced by an individuals' previous behavior. Moreover, individuals can become involved in strategic information search to reduce uncertainty and these processes are not captured by Kahneman's and Tversky's (cumulative) prospect theory. Arentze and Timmermans (2003c, 2004) therefore formulated a more general theory and model for this case of learning and adaptation. The approach they suggest can be seen as an extension of the reinforcement-learning model. As described above, in reinforcement learning, a subject's current beliefs about attributes of choice alternatives are represented by an expected value for each attribute. In contrast, in the Bayesian-learning framework that they propose, a subject's current beliefs are modeled as a probability distribution for each attribute. In addition to the expected value, a distribution provides information about the degree of uncertainty. Thus, the expected utility gains and losses associated with each alternative can be computed. The Bayesian method of belief updating is used as a model of cognitive learning during the implementation of activities and travel. Based on this method, the expected information gain of each alternative can be computed before implementing the choice. In the model they propose, subjects take into account the risks (as in prospect theory) as well as the expected information gain in making decisions. So, whereas decisions to engage in exploration are random in a reinforcement-learning framework (e.g., equation (11)), they are a function of expected information gain and expected utility losses and gains in the new model.

Their operationalization of the model is related to learning about destinations in an activity-based perspective. They assumed that the environment is represented as a regular grid of cells. Each cell represents a location where an activity can be conducted and is described by a vector of variables, influencing spatial choice behavior. Further, they assumed that a trip from an origin to a destination is modeled as a path through the transport network between the nearest nodes from the cells.

In addition, they assumed that the variables describing locations are discrete and that individuals hold some beliefs about these discrete values. These beliefs reflect an individual's uncertainty about the state of the environment. An individual's awareness space or mental map is then represented by the full set of probability distributions.

An individual's beliefs about the state of the environment are updated (cognitive learning) based on observations during the implementation of trips and activities. Because the observations are not necessarily error-free, a Bayesian updating mechanism is assumed. Let $X$ be a particular variable in a certain cell and $Y = y_s$ denote the outcome that $X = x_s$ in that cell. Possible error of observation means that the probability of $Y = y_s$ given that $X = x_s$ is not necessarily equal to one and the probability that $Y = y_s$ given that $X \neq x_s$ is not necessarily equal to zero. Given some observation outcome $y_u$, the belief in $x_s$ is then updated according to the well-known Bayesian method:

$$P(x_s \mid y_u) = \frac{P(y_u \mid x_s)P(x_s)}{\sum_{s'} P(y_u \mid x_{s'})P(x_{s'})}, \forall \mathrm{s} \tag{12}$$

where:

$P(x_s)$ is the prior belief in $x_s$;
$P(x_s \mid y_u)$ is the updated belief after observation $y_u$;
$P(y_u \mid x_s)$ is the probability of observation $Y = y_u$ given $X = x_s$;
The updated belief, $P(x_s \mid y_u)$, is the prior belief in a next observation, implying an incremental learning process.

Observations can be made during travelling on a link or conducting an activity in a cell. To reflect the notion that error may vary across time and space, a sensitivity parameter is assumed and defined as:

$$\theta = f(S, D, X) \tag{13}$$

where $S$ is the state of the individual, $X$ is the variable being observed and $D$ is the distance between the individual and the object of the observation. The state of the individual may include many things, such as the speed of travelling, the transport mode used, the motivational state, etc. A multinomial logit model was used to predict conditional observation outcome probabilities:

$$P(y_u \mid x_s) = \frac{\exp(\theta\beta_{us})}{\sum_{s'=1}^{n} \exp(\theta\beta_{us'})}, \ \forall u,s \tag{14}$$

where $\beta_{us}$ are observation-bias parameters and $\theta$ is an observation-sensitivity parameter.

Apart from observation, their model of mental maps takes into account two additional aspects. First, belief updating takes place in the context of a Bayesian belief network (BBN). In a BBN, the links between beliefs represent causal or statistical dependencies between variables allowing the individual to make inferences about locations. As a consequence of reasoning, an observation tends to have a wider impact than just the belief to which it is related. Second, the difference between the current belief and the a-priori belief in a certain value of a cell represents specific knowledge about that cell. As a consequence of limited memory retention, it is assumed that specific knowledge is subject to decay over time and this is modelled by letting posterior probabilities return with a given step size to their corresponding a-priori values. The step size is a parameter determining the speed with which the results of belief updating decay over time.

Expected utility, expected information gain and uncertainty are assumed to be the three main criteria in location choice under uncertainty. Information gain is defined as the decrease in uncertainty caused by making a trip and conducting an activity at the destination. The subjective value of information gain has three components. First, the information obtained allows the subject to make better-informed decisions in the future. Second, information gain is proportional to the degree of novelty of a choice alternative and, as such, measures the extent to which the choice of the alternative can satisfy curiosity or pleasure of exploring. Third, on the negative side, information gain is inversely proportional to the familiarity of a choice alternative.

Uncertainty is proportional to the difficulty of predicting the consequences of a choice alternative (in terms of its utility). As prospect utility theory states, the influence of uncertainty on the perceived value of a course of action is complex. In determining a choice, subjects consider a base alternative as a reference point and evaluate alternatives on the expected loss and gain they would yield. The size of the distortion of the symmetry can be conceptualised as the ratio between the weights of expected loss and expected gain. The ratio depends on the position of the person on a risk aversion – risk-taking scale. Risk aversive persons have a high (absolute) value of the ratio compared to risk takers and consequently avoid choice alternatives with a higher expected loss more strongly than risk takers do.

Based on these concepts, and assuming the origin location, route, transport mode and purpose of the trip as given, the utility of a location choice alternative $l$ is defined as:

$$U_l = U_l^{\text{trip}} + U_l^{\text{dest}} \tag{15}$$

where $U_l$ is the utility of cell $l$, $U_l^{\text{trip}}$ is the utility of the trip and $U_l^{\text{dest}}$ is the utility of the destination. Conceptually, expected utility, expected information gain and uncertainty are all included in both the trip and destination component. In a simulation, they assumed that there is no uncertainty related to the transport network so that information gained by travelling arises exclusively from observations of land characteristics. Then, the terms on the right hand side of equation (15) can be decomposed as:

$$U_l = V_l^{\text{trip}} + \mu^g E(G_l^{\text{trip}}) + E(V_l^{\text{dest}}) + \mu^g E(G_l^{\text{dest}}) + \mu^- E(V_l^{-,\text{dest}}) + \mu^+ E(V_l^{+,\text{dest}}) \quad (16)$$

where $E(G)$ is an expected information gain, $E(V)$ is an expected utility, $E(V^-)$ is an expected utility loss, $E(V^+)$ is an expected utility gain and μ are weights. The weights are influenced by the positions of the subject on the familiarity – novelty scale and the risk aversion – risk-taking scale. The utility of a trip, $V_l^{\text{trip}}$, is defined in the usual way as a function of distance and possibly other attributes of the route.

The expected loss and gain are defined based on the notion that, given a reference utility level $V^0$, the expected utility level can be decomposed as:

$$E(V) = P(V < V^0)E(V - V^0 \mid V < V^0) + P(V > V^0)E(V - V^0 \mid V > V^0) + V^0 \quad (17)$$

where $P(X)$ is the probability of event $X$, $E(V \mid X)$ is the expected utility given event $X$ and $V^0$ is the utility of a base alternative. The first and the second product on the right hand side of equation (17) correspond to the expected utility loss and expected utility gain respectively.

The expected gain and loss at the utility level are elements that the model shares with prospect theory. However, the expected information gain is a new element defined as:

$$E(G_l) = \sum_{l' \in L} \sum_{k \in K} E(G_{l'k}) \quad (18)$$

where $L$ is the set of cells that can be observed on a trip to $l$ and $K$ is the set of relevant attributes. They conceptualise the expected gain from a (future) observation as:

$$E(G_{lk}) = H(X_{lk}) - E\{H(X_{lk})\} \quad (19)$$

where $H(X_{lk})$ is the entropy of $X_{lk}$ before and $E\{H(X_{lk})\}$ the expected entropy after the observation. In determining the expected value, the model takes the subject's observation-sensitivity function into account so that it predicts a low value for cells that are more distant from the route or, otherwise, less well observable.

## 5. CONCLUSIONS AND DISCUSSION

The goal of this paper has been to briefly discuss some recent and still ongoing work on modelling learning and adaptation in transport settings. Learning occurs because the environment is non-stationary, creating uncertainty. Several approaches were briefly discussed. Elements of a more general conceptual framework were identified, and it was shown how different models fit into this general framework.

Due to space limitations, this state-of-the-art paper was necessarily very selective. First, many empirical studies on uncertainty, a rapidly growing area of research, were not discussed. Secondly, not all forms of uncertainty were discussed. Of special importance here is the uncertainty of the choice behaviour of other travellers as this strongly influences the status of the transportation system over time. Han et al. (2004) have recently explored the possibilities of game theory to address this problem.

Finally, the lack of space prevented us to develop a mega framework, which would entail the various forms of uncertainty and related learning and adaptation processes. Our intention however is to develop such a framework in the near future and conduct numerical and empirical studies into the dynamic of activity-travel choice behaviour.

# REFERENCES


Anderson, J.R. (1983) *The Architecture of Cognition*. Cambridge, Harvard University Press.

Arentze, T.A., Hofman, F. and Timmermans, H.J.P. (2004) Predicting multi-faceted activity-travel adjustment strategies in response to possible congestion pricing scenarios using an Internet-based stated adaptation experiment. *Transport Policy*, **11**, 31-41.

Arentze, T.A. and Timmermans, H.J.P. (2000) Albatross, A Learning-Based Transportation Oriented Simulation System. EIRASS, Eindhoven University of Technology.

Arentze, T.A. and Timmermans, H.J.P. (2003a) Modeling learning and adaptation processes in activity-travel choice: A framework and numerical experiment. *Transportation*, **30**, 37-62.

Arentze, T.A. and Timmermans, H.J.P. (2003b) Albatross 2.0, A Learning-Based Transportation Oriented Simulation System. EIRASS, Eindhoven University of Technology.

Arentze, T.A. and Timmermans, H.J.P. (2003c) Representing mental maps and cognitive learning in micro-simulation models of activity-travel choice dynamics. *Proceedings of the IATBR Conference*, 10-16 August, Lucerne.

Arentze, T.A. and Timmermans, H.J.P. (2004) Capturing the role of awareness and information search processes on choice set formation in models of activity-travel behaviour. Accepted for presentation at the Annual Meeting of the Transportation Research Board.

Axhausen, K., Dimitrakopoulou, E. and Dimitripoulos, I. (1995) Adapting to change: Some evidence from a simple learning model. *Proceedings PTRC*, P392, pp. 191-203.

Ben-Akiva, M.E., De Palma, A. and Kaysi, I. (1991) Dynamic network models and driver information systems. *Transportation Research*, **25A**, 251-266.

Ettema, D., Tamminga, G. and Timmermans, H.J.P. (2003a) A micro-simulation model system of departure time and route choice under travel time uncertainty. Paper presented at the 10th International Conference on Travel Behaviour Research, August, Lucerne.

Ettema, D., Timmermans, H.J.P. and Arentze, T.A. (2003b) Modelling perception updating of travel times in the context of departure time choice under ITS. Paper presented at the Euro-workshop “Behavioural Responses to ITS”, 1-3 April, Eindhoven University.

Han, Q., Dellaert, B., van Raaij, F. and Timmermans, H.J.P. (2004) Modeling the impact of emerging, uncertain, time-dependent, aggregate activity-travel patterns on individual activity participation and timing decisions. Accepted for presentation at the Annual Meeting of the Transportation Research Board, Washington D.C., USA.

Joh, C.-H., Arentze, T.A. and Timmermans, H.J.P. (2001a) A theory and simulation model of activity-travel rescheduling behaviour. Paper presented at the 9th World Conference on Transportation Research (WCTR), July, Seoul, to appear in Proceedings.

Joh, C.-H., Arentze, T.A. and Timmermans, H.J.P. (2001b) Understanding activity scheduling and rescheduling behaviour: Theory and numerical illustration. *GeoJournal*, **53**, 359-372.

Joh, C.-H., Arentze, T.A. and Timmermans, H.J.P. (2002) Modeling individuals' activity-travel rescheduling heuristics: Theory and numerical experiments. *Transportation Research Record*, **1807**, 16-25.

Joh, C.-H., Arentze, T.A. and Timmermans, H.J.P. (2003) Estimating non-linear utility functions of time use in the context of an activity schedule adaptation model. Paper presented at the 10th International Conference on Travel Behavior Research (IATBR), August, Lucerne.

Joh, C.-H., Arentze, T.A. and Timmermans, H.J.P. (2004) Activity-travel (re)scheduling decision processes: Empirical estimation of the Aurora model. Annual meeting of the Transportation Research Board, Washington D.C., USA.

Kaebling, L.P., Littman, M. and Moore, A.W. (1996) Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, **4**, 237-285.

Kahneman, D. and Tversky, A. (1979) Prospect theory: An analysis of decisions under risk. *Econometrica*, **47**, 263–291.

Koutsopoulos, H.N. and Xu, H. (1993) An information discounting routing strategy for advanced traveler information systems. *Transportation Research*, **3C**, 249-263.

Nakayama, S. and Kitamura, R. (2000) A route choice model with inductive learning. Paper presented at the 79th Annual Meeting of Transportation Research Board, Washington D.C., USA.

Pendyala, R.M., Kitamura, R., Chen, C. and Pas, E.I. (1997) An activity-based microsimulation analysis of transportation control measures. *Transport Policy*, **4**, 183-192.

Pendyala, R.M., Kitamura, R. and Reddy, D.V.G.P. (1995) A rule-based activity-travel scheduling algorithm integrating neural networks of behavioral adaptation. Paper presented at the EIRASS Conference on Activity-Based Approaches, May, Eindhoven, The Netherlands.

Pendyala, R.M., Kitamura, R. and Reddy, D.V.G.P. (1998) Application of an activity-based travel demand model incorporating a rule-based algorithm. *Environment and Planning B*, **25**, 753-772.

Polak, J. and Hazelton, M. (1998) The influence of alternative traveller learning mechanisms on the dynamics of transport systems. *Transportation Planning Methods* **1**, 83-95.

Timmermans, H.J.P., Arentze, T.A. and Joh, C.-H. (2002) Analyzing space-time behavior: New approaches to old problems. *Progress in Human Geography*, **26**, 175-190.

Timmermans, H.J.P., Arentze, T.A. and Joh, C.-H. (2001) Modeling the effects of anticipated time pressure on the execution of activity programs. *Transportation Research Record*, **1752**, 8-15.

Tversky, A. and Kahneman, D. (1992) Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, **9**, 195–230.