Ideas for Incorporating Behavioral Dynamics in Activity-Based Travel Demand Models

Paper submitted to the 3rd TRB Conference on Innovations in Travel Modeling, May, 2010

Mark Bradley

Introduction

This is not a new topic, but it is one that needs to be revived in the field of travel demand modeling. In recent years, we have seen many advances and innovations, in both academia and practice, in the areas of land use modeling, activity-based demand modeling, dynamic network traffic modeling, discrete choice model estimation, and others. What has not been seriously addressed, however, is one of the key remaining methodological shortcomings of travel demand models: the models are designed to predict changes in behavior over time, but they are specified as static and instantaneous. Virtually all of the travel forecasting models in use today are based on cross-sectional data from a single point in time, and travel choices are modeled solely as a function of circumstances at that point in time. In reality, however, travel behavior is often conditioned by habits and past circumstances, and changes in behavior are often triggered by changes in circumstances over time. In other words, our models are based on a snapshot, when we really need to use a video camera.

The apparent recent lack of interest in dynamic models of travel demand has not always been the case. In the late 1980's and early 1990's, the Dutch Mobility Panel Survey was fielded in the Netherlands. Participating households completed seven-day travel and activity diaries at intervals of six months, and many households remained in the sample for a number of years. This type of panel data provided a new and unique opportunity to model travel behavior as it unfolded through time for individual households and persons. Based on this data, Meurs (1991) estimated a panel-based model of mode choice behavior, using analysis methods developed by Heckman (1981) and others in fields where panel survey data was more prevalent. Based on the same data Kitamura and Bunch (1989) developed a dynamic model of vehicle transactions within households, and Goulias and Kitamura (1997) developed the MIDAS model for application in the Netherlands. The MIDAS model has some of the features of today's activity-based microsimulation models, simulating travel choices for individuals and households. Its unique feature is that it also simulates changes in those households over time, and simulates changes in travel choice behavior that are influenced by explicit changes in household and travel conditions between time periods.

In retrospect, the interest in panel surveys and dynamic travel demand models appears to have peaked at the time of the First (and only) US Conference on Panels for Transportation Planning, held at Lake Arrowhead, CA in 1992. At that time, the Puget Sound Regional Transportation Panel Survey was getting underway, and much of the attention was understandably on how to administer such ambitious multi-wave, multi-day travel surveys efficiently (Murakami and

Ulberg, 1997), and on how to make the data representative in spite of the sample attrition that inevitably occurs (Pendyala et al., 1993). Relatively little attention was paid to specific ways in which the data could be analyzed and used to improve travel demand models as applied in practice. For a variety of possible reasons, that situation has continued, and no further major transportation panel surveys have been carried out, with the exception of the German Mobility Panel, which has been carried out continuously since 1994 (Zumkeller and Chlond, 2009). In the US, although some research has continued in academia, dynamic, panel-based travel demand forecasting models have yet to be developed for practical use.

In the spirit of the late Prof. Ryuichi Kitamura, who has been our profession's most articulate and prolific advocate for dynamic models, I believe that the time is ripe to bring behavioral dynamics into practical travel demand forecasting models. In the remainder of this paper, I briefly summarize why this development is so important. I then give reasons why it now seems much more feasible than it did in the past. I close by providing some specific ideas for how dynamics of travel choice behavior could be incorporated into activity-based model systems.

Why Dynamic Models?

When people talk about transportation policy and how it influences behavior, we usually use words such as "change", "shift", or "move". Similarly, when we judge the performance of our forecasting models, we typically look at elasticities--how the dependent variables change relative to changes in the input variables. In both cases, what is implied is change <u>over time</u>, and it is well known empirically that longer term elasticities tend to be different than short term elasticities. Yet, our travel demand models are not estimated or implemented in that way. There is no explicit link between the forecasts for one year and the forecasts for a future year, and no possible way that the models, as they are specified, can predict different short-term and longer term impacts from a single policy change. To do so, the model specification would need to include one or more of the following:

- <u>State-dependence</u>: Choice behavior in one period is dependent on the choice made in a previous period. (For example, the decision of which mode to use to regularly commute to work depends on the mode that the person has regularly used in the past, since that influences the level of information and experience that the person has regarding the modes.) This is a type of heterogeneity that needs to be modeled in a way that can distinguish it from other types of heterogeneity or taste variation, related to income, age, personality type, and so forth. Such an analysis requires segmenting the model depending on the choice made in the previous period, and/or adding explicit variables to represent habit or "inertia" effects. This in turn requires panel data or before-and-after data from multiple time periods in order to statistically sort out the different types of heterogeneity.

- Lagged variables: It is often the case that it takes people time to adjust to new travel circumstances. For example, if fuel prices increase dramatically, some people may eventually buy a more efficient vehicle or change their work or home location to reduce driving distances. The longer that fuel prices remain high, the more common such behavior becomes. So, people are not reacting only to the current fuel price, but, in a sense, to past fuel prices as well. Such phenomena can be modeled by using exogenous variables from multiple points in time.

- Change/Rate of change variables: People often change behavior when they are spurred by a significant change in their choice environment. Changes in travel habits tend to occur when people get a new job, retire, move house, get married, get a divorce, have children, or become "empty-nesters". Drivers tend to react more strongly to a sudden, sharp jump in fuel prices or travel congestion than they do to a slow, steady rise. So, it is not only the magnitude of a change that influences behavior, but also the rate at which the change occurs, and whether they occur in "quantum leaps" or continuously over time. This is represented by using as variables the difference in the input variables from one time period to the next.

When variables of the type listed above are included in choice models, they can show a wider variety of predicted behavior over time, including asymmetric effects (for example, different magnitude of changes resulting from fuel price increases as opposed to fuel price decreases), lagged effects, different short- and longer-term elasticities, and even oscillating behavior. Since these different types of variables can all represent similar-seeming behavioral patterns over time, it is important to have theoretical hypotheses about which dynamic processes are most important, and to try several different model specifications to test those hypotheses. For example, Bradley (1997) tried several different model specifications on a before-and-after mode choice data set from the Netherlands, including static, one-period multinomial (MNL) and nested logit models, as well as multi-period models using lagged variables, inertia variables, endogenous segmentation, and exogenous change variables. The models were then applied over time to determine the short term and longer term choice elasticities implied by the various model specifications. The results from that paper are shown in Figure 1, Clearly, the behavior predicted by the models depends a great deal on the types of dynamic variables included in them. One of the conclusions of that paper is that data from a single before and after survey from only two or three points in time will generally not be adequate to allow conclusive testing of different dynamic model structures. That requires data from more points in time and a wider variety of choice contexts, both of which a regional panel survey can provide.

The argument is sometimes raised that these types of dynamic effects are important for short-term forecasting, but for longer term forecasts for twenty or thirty years out, static cross-sectional models will be just as accurate. It is indeed tempting to think that this is true, as it would increase our confidence in our current state-of-the-practice models. There is, however, no logical reason why it should be true. It is the same as saying that the specific events in history have not been important in determining where we stand today. For example, one can question whether travel behavior would be the same in 2030 if fuel prices rise gradually by a few percent per year each year, versus another scenario where they arrive at the same 2030 level, but with many very sharp spikes and drops in between. I would not be willing to bet much money that peoples' overall travel habits in 2030 would be very similar in the two scenarios. But, we need dynamic models to predict how they would be different.

Figure 1: Summary of elasticities from example model specifications, from Bradley (1997)

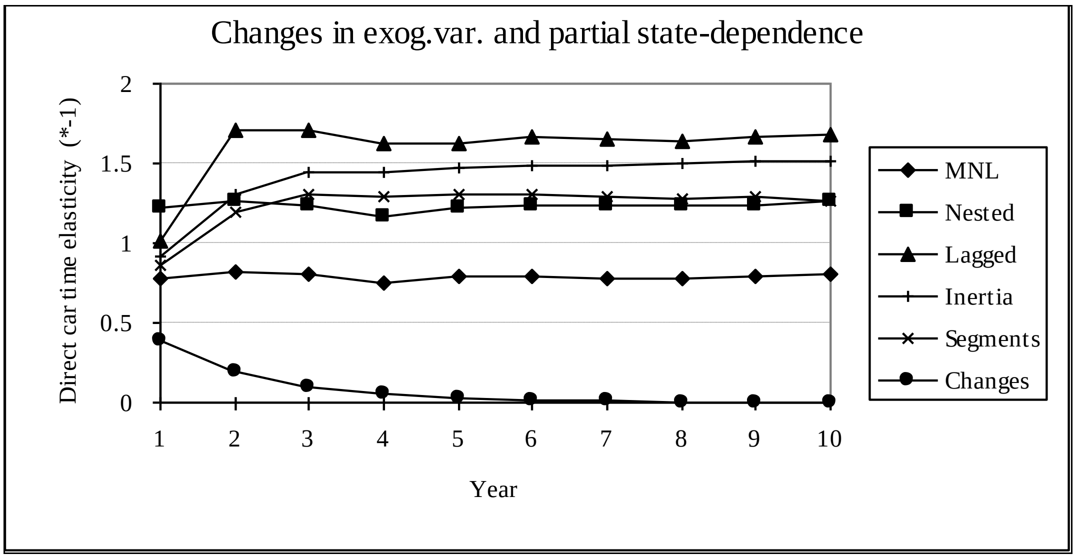

Why Now?

Until now, the major roadblocks holding back dynamic travel demand modeling have been largely practical—the lack of sufficient disaggregate longitudinal data, the difficulty of estimating complex dynamic discrete choice models, and the lack of flexibility in model application software packages. All of these obstacles appear to be disappearing in recent years:

- Longitudinal data: Ideally, dynamic models will be based on panel data with at least 3 waves of data per household, and at least 7 days of travel and activity data per wave (and preferably even more waves and days). The considerable respondent burden and administrative burden involved in carrying out panel surveys with typical phone-plus-mail travel diary survey methods has made such surveys seem impractical, for the most part. Currently, however, we are seeing the beginning of a major shift towards passive, GPS-based methods for collecting household travel data. When time and location data from GPS devices are merged with detailed land use data and other exogenous data sources, it is not necessary to ask respondents any direct questions in order to get a quite complete picture of their travel and activities. In the future, we can expect such GPS data to be provided by mobile phones instead of specialty devices, and, once permission is received from respondents, it may be possible to interact with those devices remotely with respondents hardly even aware that they are being "surveyed". In such a scenario, it will be possible to collect data more frequently and for longer periods of time from individual households, providing the type of longitudinal data needed for dynamic models, but with few of the traditional problems associated with panel surveys.

- Mode estimation software: In the past, it has typically been necessary to program the likelihood functions for dynamic, panel-based models in complex software packages such as GAUSS or MATLAB. Now, however, there are various commercial and free software packages available that include mixed logit and other GEV methods, making it much easier to estimate models that include various types of error components structures and types of heterogeneity. In recent years, such methods have mainly been applied to static models, but could be used for dynamic models as well.

- Model application software: Typical software packages used to apply travel demand models in the past, such as TP+, TransCAD and EMME/2, have used a zone-based looping structure that is not very flexible in terms of explicitly linking travel demand models across time periods, simulating individual households, etc. Now, with the newer generation of activity-based models, travel demand modeling is moving towards a micro-simulation approach that is much more suitable and adaptable for applying dynamic models. Specifically, such a model application framework can create and store a "running history" of the predicted choice situations and choice behavior for each simulated household, across days, weeks and years.

How could dynamic models fit into existing activity-based model frameworks?

There are several different time scales that need to be captured in travel demand forecasting, as depicted in Table 1. In the very long term, land use changes due to land development and the location decisions of businesses and households take place over years, and may remain unchanged for as long as decades. Because land use models deal with physical entities such as land and buildings, it has been quite natural to use a dynamic, state-based microsimulation framework in models such as UrbanSim and PECAS, where each year's simulation starts with the physical parcels and buildings from the previous year's simulation, and then simulates discrete changes that persist through time. As shown in the intermediate shading in the table, some land use models also deal with behavior such as choice of workplace and school, vehicle ownership, or demographic shifts within households, while, in other instances, those aspects of behavior are left to be dealt with in travel demand models and/or synthetic population generators.

At the other extreme, decisions that influence how traffic flows on a network, such as driving behavior and route choice, take place on the time scale of seconds; minutes at the most. As the computation has become more feasible, models of traffic flow are also moving towards dynamic microsimulation and dynamic traffic assignment (DTA). In some cases, those dynamic models can also reflect changes in trip departure time, while that choice dimension is also predicted as part of most activity-based model systems as well.

The primary time scale dealt with in all activity-based travel demand models is the hours in a single, representative day. The day's activity schedule is predicted and chained into tours, and the mode, destination and timing of each tour and trip is also predicted.

Table 1: Time scales of travel behavior and modeling domains

Time Scale	Relevant Behavior	Modeling Domain
Seconds	Driving behavior	Traffic
	Route choice	Microsimulation
Minutes	Trip departure time choice	
Hours	Tour/trip mode choice	Activity-Based
	Tour/trip destination choice	Travel Demand
Days	Day's activity schedule	Microsimulation
Weeks	Activity scheduling "habits"	
	Location choice "habits"	
Months	Mode choice "habits"	
	Buying/selling vehicles	
Years	Choice of school	
	Choice of workplace	
	Household transitions	
	Choice of residence	Land
	Business location	Use
Decades	Land development	Microsimulation

What is lacking in current activity-based travel demand models are the dynamic shifts in travel behavior patterns that manifest not within a day, but over a periods of weeks, months, and even years. These are depicted in the table as "habits" related to activity scheduling, locations visited, and modes used, but they may not be what we normally think of as habits. They can include the places that people choose to go grocery shopping or to the dentist, how early people choose to get up to go to work in the morning, or how often someone goes to the gym or walks the dog. Many of these things may seem to the casual eye as constraints more than "habits" (work starting and ending times, for example), but through analysis of longitudinal data, one finds that people often actually change things that they have labeled as constraints in the past.

Another finding from past analysis of longitudinal data is that it is very hard to find relationships between a single day's travel and activities in one year and a single day's travel activity in the next year, even if it is the same day of the week and time of the year. There is too much day-to-day variability in detailed travel and activity patterns to expect much consistency between one day and the next day, never mind two days that are a year apart. To begin to find relatable patterns across longer time intervals, one needs to aggregate the period of observed behavior from a day to at least a week, and preferably longer, such as two weeks or a month. Also, one needs to aggregate the measures of observed behavior over those longer periods into categories of similar behavior. For example, if somebody takes transit to work 18 days out of 20 in November one year, and then takes transit to work 17 days out of 20 in November the next year, one might conclude that those numbers are close enough to consider them essentially the same type of usual or typical behavior.

Figure 2 provides a schematic diagram of how "higher level" or aggregated travel and activity patterns of the type just described could be used in a dynamic activity-based model structure (in this example, also integrated with a land use microsimulation model). After land use and longer term choices such as choice of work and school location are predicted, these condition a model that predicts one of a discrete typology of higher level travel and activity patterns. This higher level pattern in turn strongly conditions the choice of a travel and activity pattern for a single, representative day that is used to provide forecasts of specific trips for use in traffic and transit modeling. This diagram strongly resembles the structure of current activity-based model systems, but with two main additions: (1) the introduction of a new model of the higher level, composite travel and activity pattern, and (2) explicit linkages between the predicted behavior of a person/household between one year and the next, both for the longer term choices and the higher level activity patterns (but <u>not</u> for the predicted single day patterns, for the reasons described in the previous paragraph). Note that models of auto ownership transactions should also be linked explicitly from one period to the next, although it is not obvious whether that should be modeled "above" or "below" the models of higher level travel and activity patterns.

Once we have sufficient longitudinal data over periods of weeks or months to do the analysis, the key challenge to create a useful model with this approach will be in defining an appropriate typology for the "higher level travel and activity pattern", across a period of a week or longer. Some useful dimensions of this typology would be:

- The frequency of participating in specific types of activities.
- The regularity with which specific destinations are visited.
- The typical mode used and average trip length for travel for specific purposes.
- The average number of trips per day and trips per tour (trip chaining)
- The average amount of out-of-home activity time versus in-home activity time
- The regularity of trip departure times and scheduling for specific purposes.

A variety of pattern recognition techniques and variance reduction approaches could be used for such an analysis. Although some examples of this type of analysis exist in the literature, there have been relatively few because of the difficulty of collecting appropriate data. I am optimistic that new GPS data collection techniques and survey programs will help to address the lack of data. To create useful dynamic models, however, it will also require a great deal of careful and creative analysis and, most importantly, the willingness to view travel behavior through a different lens—one that can be set to focus on changes over time.

Figure 2: General Schematic Diagram of Dynamic Activity-Based Model Structure

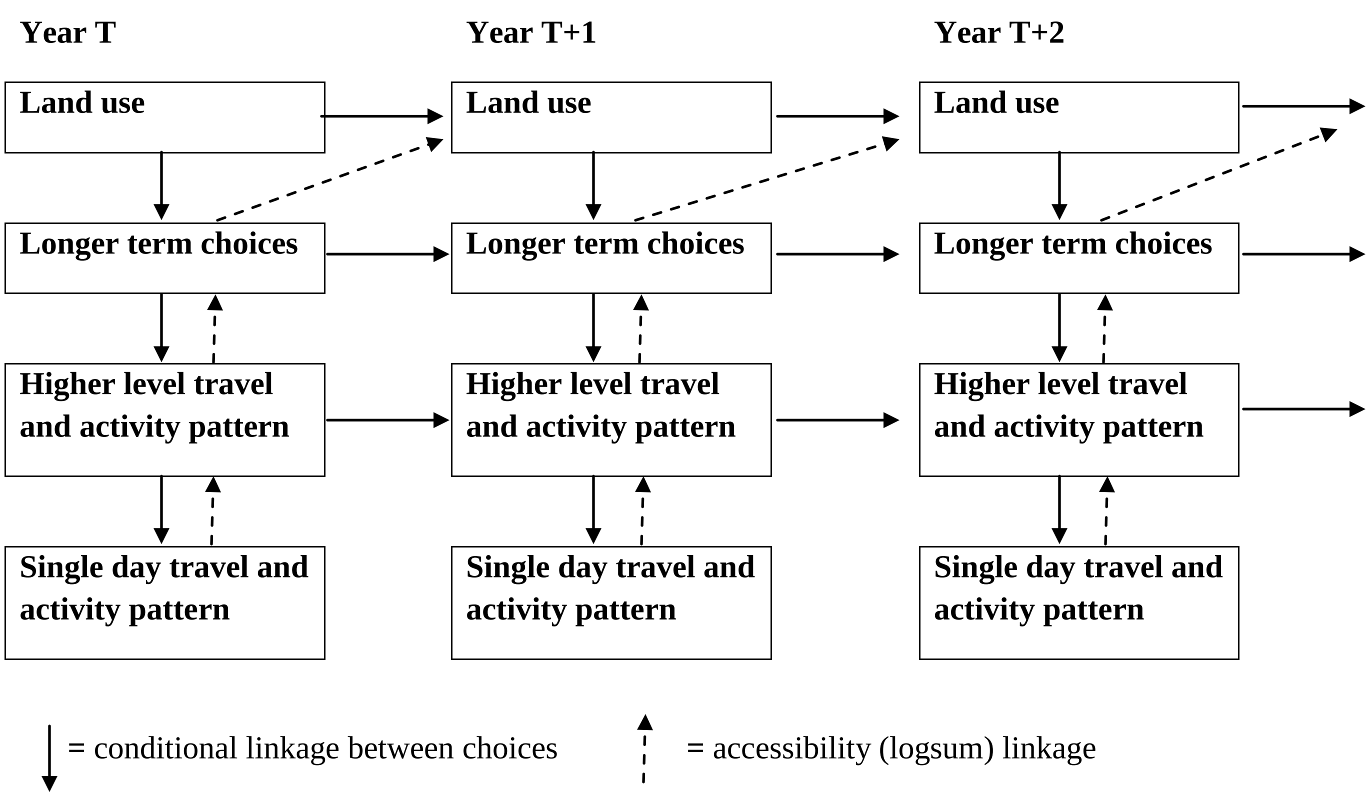

References

Bradley, M. (1997). A practical comparison of modeling approaches for panel data. In T.F. Golob, R. Kitamura and L. Long (eds.), *Panels for Transportation Planning: Methods and Applications*, T. Golob, et al. ed., Kluwer , Dordrecht, NL, pp. 281-304.

Goulias, K.G. and R. Kitamura (1997) Regional travel demand forecasting with dynamic microsimulation models. In T.F. Golob, R. Kitamura and L. Long (eds.), *Panels for Transportation Planning: Methods and Applications*, T. Golob, et al. ed., Kluwer , Dordrecht, NL, pp. 321-348.

Heckman, J.J. (1981). Statistical models for discrete panel data. in Manski and McFadden (eds.): *Structural Analysis of Discrete Data with Econometric Applications*; MIT Press, Cambridge, MA.

Kitamura, R. and Bunch, D. (1989). Heterogeneity and state dependence in household car ownership: A panel analysis using ordered-response probit models with error components. Research Report of Transportation Research Group, University of California at Davis.

Meurs, H. (1991). *A Panel Data Analysis of Travel Demand*. Doctoral dissertation. University of Groningen, Netherlands.

Murakami, E., & Ulberg, C. (1997). The Puget Sound Transportation Panel. In T.F. Golob, R. Kitamura and L. Long (eds.), *Panels for Transportation Planning: Methods and Applications*, T. Golob, et al. ed., Kluwer , Dordrecht, NL, pp. 159-191.

Pendyala, R.M., K.G. Goulias, R. Kitamura and E. Murakami (1993) Development of Weights for a Choice-Based Panel Survey Sample with Attrition. *Transportation Research Part A* 27 (6), 477 - 492

Zumkeller, D. and B. Chlond. (2009). Dynamics of Change: Fifteen-Year German Mobility Panel. Paper presented at the Transportation Research Board 2009 Annual Meeting. Paper #09-0463. Washington, DC.