Activity-Based Models: 1993-2012

Berkeley Presentation
September 14, 2012

John L Bowman, Ph.D.
John_L_Bowman@alum.mit.edu
JBowman.net

INTRODUCTION

(page 1) As a second year masters student at MIT in 1993 I was interested in activity-based travel analysis and how it might apply to practical travel forecasting. I had read the research papers of Peter Jones, Eric Pas, Ryuichi Kitamura and others. I had studied the tour-based model development work of the Hague Consulting Group in the Netherlands, Staffan Algers in Stockholm and Ennio Cascetta in Italy. I had gone back to read the foundational works of Hagerstrand, Chapin, Lentorp and others. I was searching for a way to make a contribution in this area that I could use for my masters thesis, but I couldn't get a handle on it.

I remember clearly the day I was discussing these things with Moshe Ben-Akiva in his office, and he said to me something like, "what they need to do is model a person's entire day of activity and travel". For me, that was the big "aha!" moment. The light bulb turned on, setting me on a course that I am still on today, 19 years later.

(page 2) Throughout that time, my research and development efforts have been aimed at bringing so called activity-based travel demand models into the mainstream of urban travel forecasting.

Along the way many other people have made significant contributions toward the same objective.

Today I will provide highlights from that experience and give a picture of where things stand in 2012.

EARLY HISTORY IN THE USDAY ACTIVITY SCHEDULE APPROACH

(page 3) My comments will fall roughly into three categories

(page 4) The first few slides that I will show you today come from presentations I gave back in the mid 90s.

This simple graphic, which comes directly from the first TRB presentation I ever gave, shows how Moshe Ben-Akiva, Dinesh Gopinath and I proposed to model a person's day activity pattern that overarches and ties together tours in a nested structure of discrete choice models spanning a day.

Tours are modeled conditional on the activity pattern, with lower priority tours modeled conditional on higher priority tours

And the activity pattern is affected by the ease of carrying out the activities, as measured in the tour models.

(page 5) The day activity schedule can be viewed as an evolutionary step in the use of disaggregate discrete choice methods to model travel demand.

--Trip-based models were developed in the early 70s

--And these were followed, especially in Europe, by tour based models.

The essence of our contribution at the time was to integrate the representation of travel demand across an entire day.

(page 6) Although the exact representation of the activity pattern has varied from time to time, this slide from that first presentation still captures the essence.

- The activity pattern identifies the number and purposes of the tours that a person conducts in their day
- and the number and purposes of additional stops on those tours.

(page 7) For each tour, the timing, destination and travel mode are modeled

- for the tour
- and for each stop on the tour

(page 8) The day activity schedule model operates iteratively with equilibrium assignment models to generate its predictions.

It essentially substitutes for trip generation, distribution and mode split in a traditional 4-step model system.

This makes it manageable for a planning agency using a traditional 4-step model to advance to the activity schedule approach.

(page 9) By the end of 1994 I had implemented the MIT prototype.

This slide from my 1995 TRB presentation gave a report on how well we had achieved our objectives.

For the most part we were very pleased with the results

- although we hadn't yet actually integrated the model with the traffic assignment models
- and we considered our use of five time periods for time-of-day modeling inadequate.

(page 10) Our MIT results caught the attention of Keith Lawton, the director of modeling at Portland Metro,

and he decided to implement the approach there.

(page 11) I led a design effort to fill in details missing from the prototype, and Mark Bradley subsequently developed an operational model system for Metro while I continued my studies at MIT.

It was the first activity-based model used for policy analysis.

(page 12) Here are a couple results from the earliest documented application, reported before the equilibration with traffic assignment was implemented.

At the time we were apologetic about the enormous size of the price increases in the test scenarios (changing variable costs from 8 cents per mile to 16)

They showed several important aspects of the model sensitivity:

(page 13) the decrease in tours shows that activity generation is sensitive to LOS, but less so than VMT

(page 14) The decrease in tours and miles came primarily from auto travel

(page 15) Elasticity of demand was greater for discretionary and maintenance than for work purposes, when the costs increased for all times of day

(page 16) and when the costs changed only in the peak period, there was a shift from the peak periods to the off-peak periods (right hand column, off-peak row)

(page 17) Unfortunately, the Portland model was mothballed because Metro's money was short and they decided to focus their attention on TranSIMS, where Federal money was available.

But fortunately, by then San Francisco County was interested in developing an advanced model, and it was decided to implement a version of the activity-schedule approach, with somewhat simplified models and integration.

The SFCTA model was developed by PB, Mark Bradley and Cambridge Systematics.

(page 18) It was the first activity-based model to be used extensively and on an ongoing basis for policy analysis and decision-making.

And it has been enhanced over time. For example, in 2007 it was enhanced to

- include toll vs free alternatives in mode choice
- and to use randomly distributed values of time in the models, where values of time are drawn from lognormal distributions as part of the simulation.

And recently they have been incorporating bicycle route choice and its logsums into the model system and using it for bicycle project evaluation.

(page 19) Several years after the SFCTA model went live, the New York and Columbus models were implemented by PB (page 20) These models represented the first attempts to explicitly model household interactions.

The Columbus model also

…represented time in one-hour periods,

…introduced multi-threaded software

…and added additional model components related to parking.

(page 21) The next major implementation that I was directly involved with occurred in 2005-2006 in Sacramento.

(page 22) Modeling detailed outcomes was important to SACOG. So we substantially improved model resolution with regard to…

(page 23) Purpose…

(page 24) Time…

(page 25) And especially space.

We modeled all location choices

- as choice of parcel
- instead of the customary traffic analysis zone

(page 26) We also emphasized integration of the components of the AB model (which I will address further in a couple minutes), and adequate equilibration between the AB demand model and network assignment.

(page 27) The Lake Tahoe project was the first implementation for a small MPO, and also the first to transfer and recalibrate a model built for another region, rather than building the model system and software from scratch with local data.

So by late 2007 there were five AB model systems in use in the United States. From my perspective, things had moved at a snail's pace. I hadn't realized how entrenched the 4-step model was and how reluctant most agencies would be to try this promising new approach.

(page 28) But the pioneering agencies were meeting with success, and there were pressing policy issues for which the 4-step modeling methods were not working well. And in 2007 and 2008 two events occurred that raised the profile of AB modeling and provided impetus for its adoption.

(page 29) The first was the publication of TRB Special Report 288.

(review the bullet points)

(page 30) The report issued a strong condemnation of the current state of the practice

(read the bullets)

(page 31) And it made some strong recommendations.... (read the slide)

The federal sponsors have taken these recommendations seriously and have begun to act on them.

And the report seems to catch the attention of metropolitan agencies.

(page 32) Independent of the federal study and report, activity-based modeling made its way into the California state law and regulation

(read the slides)

Interestingly, nowhere do these documents actually define what activity-based models are, even though they refer to them explicitly and say that they should be developed and used.

(page 33) But it seemed clear that they were referring to the model systems I already mentioned, plus others that were by then in the pipeline.

(page 34) Jumping forward to today, this map shows the installed model systems and those under development. Not surprisingly, the heaviest activity has occurred in California.

(page 35) ...and new developments continue to start in the U.S., and have begun to occur internationally as well.

BASICS OF THE “ACTIVITY-BASED” MODEL SYSTEMS

(page 36) Now let’s switch gears and take a look at the nature of these so called activity-based model systems

(page 37) To do this I start by showing the framework of a trip-based model system, since

--that is what these agencies are familiar with

--AND much of the framework and components are the same

(page 38) The transport model system takes inputs from the urban forecasting and planning process:

(page 39) --transport networks representing potential future scenarios

(page 40) --zonal attributes representing forecasts of size and distribution of employment and population in the region

(page 41) --and socioeconomic attributes of the population.

(page 42) It then predicts zone-to-zone trip flows

(page 43) --and assigns those trips to the network

(page 44) This occurs iteratively to achieve a final result where

--inputs to the trip demand models are consistent with

--predicted link flows and travel speeds coming from assignment of the demand they generate

(page 45) In some cases, accessibility measures from the transport model system serve as input to the land use forecasting models

(page 46) The trip demand includes: --trips associated with special generators, such as airports

(page 47) --trips carried out by persons who live outside the region

(page 48) --trips conducted by residents of the region

(page 49) --and commercial movements

(page 50) These trips are usually modeled by separate demand models and the results are combined into trip matrices for highway and transit assignment.

Here is where we can start talking of the difference introduced by an activity-based model….

(page 51) In the existing so-called activity-based model systems, only the trips of residents are modeled using the activitybased model.

(page 52) These are produced by a household travel demand simulator.

--It generates a synthetic population representing the future population of the region,

--and predicts activities and trips for every member of each synthetic household.

(page 53) In some cases, the activity-based model uses a large database of parcel attributes instead of only zonal attributes.

--It takes a lot of resources to generate the parcel database

--But it is very desirable because it provides substantially better information about

--transit accessibility,   --travel times for non-motorized and short trips, --and destination attractiveness in heterogeneous zones.

In places where high quality parcel data is not available, we have begun to define microzones that are about the size of census blocks, to get much of the benefit of parcels with a lot less work (page 54) The main output of an activity-based travel demand simulator is a detailed itinerary for every person in every household

(page 55) It comes in the form of tables or lists that identify attributes

--of each household

--each person in that household

--each person's day

--each travel tour in their day

--and each trip on that tour

(page 56) These trips are then aggregated and combined with the other predicted trips into matrices for assignment.

(page 57) The activity-based model output can also be aggregated in a great variety of ways,

- in order to extract information that is not available from trip-based models,
- ESPECIALLY the impact of a policy on various population subsets

So, at this point we see all the components of a so-called activity-based model system needed for transport forecasting.

(page 58) The travel demand simulator itself consists of a lot of interconnected model components, usually discrete choice models. This abstract diagram shows four basic categories in a hierarchy.

(page 59) And this diagram from an existing model system shows specific choice models within each category.

The models in the system are applied for each household one at a time according to the hierarchy of the system, eventually simulating all the elements of the one-day itinerary for each member of the household.

(page 60) Most of the model components in the model system are logit models representing the choice of one alternative from a discrete set of alternatives.

Each alternative has a utility function comprised of a vector of variables multiplied by a vector of estimated coefficients.

Using these the probability of each alternative is calculated

And a Monte Carlo draw is made to simulate the choice of one of the alternatives.

(page 61) Including a lot of related model components makes it possible to simulate itineraries that appear to be realistic.

But in order for them to behave realistically, they must be appropriately integrated.

(page 62) Given a hierarchical structure of connections among the models, we can speak of integration in two directions….down the hierarchy, and up the hierarchy.

(page 63) Speaking first about downward integration, models lower in the hierarchy need to treat the outcomes of higher models as given.

For example:

--If a person is a worker with a non-home usual work location,

(page 64) then their day is likely to be a work day.

If their day is a workday,

(page 65) then a work tour is likely to occur.

If there is a work tour, then it is more likely to go to their usual workplace than elsewhere,

And non-work tours cannot be made at times of the day that have already been claimed by a work tour.

(page 66) And Intermediate stops are likely to occur near their work destination.

In order to properly achieve this downward integration, each model lower in the hierarchy

--must be constrained by the outcomes of models higher in the hierarchy

--and its utility functions must also realistically represent the effects of those higher outcomes.

(page 67) Models higher in the hierarchy need to be sensitive to conditions affecting lower level models. This is upward integration.

(page 68) For example,

- If transit service in a certain corridor is significantly improved for trips occurring during the peak period,

(page 69) then

- tours by transit at that time of day should increase among persons living and working in that corridor
- and their tours in the corridor by other modes should probably decrease

(page 70) more people will arrange their day to include a transit commute,

- but non-work auto tours might increase
  - as workers run errands after work that they used to do by car on their work commute.
  - and more non-workers have a car available during the day

(page 71) Workers in the area may be more likely to buy a transit pass,

--and more households in the area should eventually own less cars

The preferred method of capturing these kinds of affects is through the use of logsums.

- These measure the expected utility across lower level alternatives
- and are used as explanatory variables in the upper level models

A big problem is that Logsum calculations take a lot of time

- so compromises are made in order to speed up run times,
- and creative techniques are developed to reduce the bad impact of such compromises.

I can't overemphasize the importance of upward integration. It is what makes the upper level decision models sensitive to changes in transport level of service.

Trip-based models can't achieve upward integration in this way because they don't explicitly model days and tours.

And if an activity-based model lacks upward integration, or does it poorly, then it has lost one of its main reasons for existence.

ADVANCED FEATURES

(page 72) Trip-based models have been around for a long time. Over the years they have been enhanced and tweaked repeatedly. By now, developers have gotten out of them about all that their underlying approach allows.

Activity-based models, on the other hand, are early in their life-cycle. The disaggregate microsimulation framework provides a good environment for the incorporation of advanced features over time as research and development continue.

(page 73) Let's look at a few advanced features that were not in the first activity-based models, but have been implemented in at least one of the current operational models.

(page 74) The use of fine spatial scale for destination choice, first implemented using block-face destinations in Portland, and then using parcels in Sacramento, opens the door for several advanced features that use the added information that the fine scale provides.

(page 75) Destination choice models typically use 'size variables' to measure the attractiveness of aggregate destination zones. Typical size variables include employment of various types, school enrollment and population or households in the zone. Zones with more size attract more activities.

However, when the destination zones are defined as parcels or very tiny zones, it is possible to distinguish between the attraction of a destination itself, and the effect that the destination's neighborhood has on its attractiveness. We measure, for example, the employment of each type in a buffer area surrounding the parcel. Employment that is very near the parcel centroid is weighted more heavily than employment further away. We do this also with school enrollment and housing, and also attributes such as street intersections of various types.

(page 76) These variables are then used as explanatory variables in the location choice models.

In this model of intermediate stop location choice for escort purpose, we see in the center column that these trips are attracted to parcels with school enrollment and, to a lesser extent, employment. But the righthand column shows us that parcels near schools are more attractive than parcels far from schools, and parcels in industrial zones are much less attractive.

(page 77) Fine spatial scale also enables much improved measurement of impedance for transit trips because walk access distance can be measured much more accurately.

To do this, although the transit in-vehicle time is measured zone-to-zone, the distance of each parcel to the nearest transit stop is used to estimate the walk access and egress time at the trip ends.

In this figure, the pentagon on the left represents the origin zone. In it two different parcels are depicted, one next to the transit stop, with short walk access time, and another in the far corner with long access time.

(page 78) Improving the impedance measure in this way provided dramatic improvements in the quality of the parameter estimates in the Portland model.

Using the traditional zone-based transit walk access time measure, the walk access value of time was quite low and statistically insignificant. But when walk access was measured as a block-face attribute, the value and significance of walk access time surpassed that of all other time components.

(page 79) A further improvement to the measurement of transit impedance made possible by fine spatial scale involves the use of transit assignment and skimming procedures that can measure transit in-vehicle time and other attributes between transit stops, rather than between zones.

In situations where a person can choose among several transit boarding locations, the AB model models the choice of transit stop pair taking into consideration total path impedance, including the stop-to-stop transit impedance as well as the access and egress times to and from the particular transit stops.

(page 80) The last advanced feature I will mention made possible by fine spatial scale is the accurate measurement of trip distances for all modes that use the road network.

For each parcel, the nearest node on an all-streets network is identified. Then for short trips the distance between any two parcels is measured as the shortest path on the all-streets network.

This technique isn't feasible for long distance trips because there would be too many parcel-to-parcel distances. But it IS feasible for short distances, where measurement error in zone-to-zone distance is large compared to the total distance.

(page 81) In recent years the implementation of various forms of road pricing and area pricing has led to an interest in modeling the effects of pricing schemes.

(page 82) It has been shown that sensitivity to price, or the value of time, is not linear in the population, but it has a long tail, much like a lognormal distribution.

Furthermore, for a given person, price sensitivity varies a lot depending on trip purpose and other factors that are difficult to model explicitly, such as penalties incurred for being late.

Given the traditional diurnal peaking patterns in congested urban areas, policy-makers want to test policies that vary by time of day.

Activity-based microsimulation models are well-suited for modeling such effects

- They can use values of time that vary from trip to trip systematically by income and purpose, and randomly according to a lognormal distribution
  This enables the models to realistically represent the small portion of trips for which willingness to pay a toll to save time is actually quite high
- They also include explicit time-of-day choice where prices are higher at some times than at others, enabling them to model the shift away from congested time periods induced by peak period tolls.

(page 83) The way this is done is by modeling the choice between the best tolled path and the best untolled path for auto trips, using skimmed path attributes that vary by time period, path type, VOT and auto occupancy.

(page 84) These methods and the associated model parameters have been developed in federally funded research projects, and are being incorporated in all the latest activity-based model implementations.

(page 85) We know that in households with more than one person, activities and travel are often coordinated and members of the household sometimes travel together.

The first activity-based models modeled activity schedules of individuals, not households. They dealt with household interactions by using household characteristics as explanatory variables.

A more desirable way to account for household interactions is to model household outcomes explicitly.

(page 86) These include modeling of a household day pattern, where the tendencies in some households to stay home from work or school together, or for one or the other of two parents to stay home with children, are model explicitly.

They also include the modeling of joint tours for non-work purposes.

And the modeling of situations where one person drops off or picks up one or more others at work or school.

(page 87) This is the schematic of the individual model system that I showed earlier

(page 88) And here is the more complicated version with explicit household models incorporated in the day pattern

This is a particularly challenging area of research and development for a couple reasons:

- First, the options available to multi-person households are almost limitless. Specifying these models presents a combinatorial challenge.

- Second, it is difficult to incorporate intra-household outcomes without compromising the integration among the model components that is essential for the models to perform realistically.

(page 89) The final topic I'll mention, in which there is much more to be done than has already been done, deals with vehicles and parking.

So far, activity-based models have primarily been focused on the activities and travel of persons.

But from a policy standpoint, the concerns have been centered around traffic congestion and air quality, and these deal primarily with vehicles.

There have been some steps taken to incorporate the vehicles into the AB models, but there is a lot of room for improvement in this area

(page 90) One thing that has been done in Sacramento and elsewhere is the modeling of park and ride lot choice.

Given a trip that has already been determined to be park and ride, the model simulates the choice of a particular park and ride lot to make the transfer between auto and transit.

The utility of each park and ride lot depends on the impedance of travel from origin to destination via that lot.

The simulator keeps track of the fill level of each lot at each minute of the day.

The availability of a lot for a given trip depends on whether there are spaces left available at the time the vehicle arrives at the lot.

A similar method can be used for "park and walk" in urban locations where many people use off-site parking facilities and walk to their final destination.

(page 91) There are additional desirable enhancements that have not been incorporated into any operational models, but which would enable better modeling of environmental effects.

The first is to enhance the auto ownership models so that the type of vehicles in the household's fleet is known.

Then when tours by auto are modeled, the specific household vehicle would be chosen.

Each vehicle's itinerary would be accounted for.

- In tour mode choice, drivers could not choose to drive a car if it was not available at the time of the tour.
- And the cold and hot starts and fuel consumption for each vehicle could be accounted for by time of day and location.

(page 92) And with this one last slide I list a few areas of improvement and extension of the activity-based approach where there is active research and development occurring.

These include… (read from slide)

(END)