Travel Modeling 101

Boston Region MPO
Prepared by MPO Staff
2012

TRAVEL MODELING 101

Introduction

This discussion is designed to give the layperson a sense of what a travel model is and how it is used. It is important that this description be provided for those who are interested, because travel modeling underlies so much of the work done by the MPO. An attempt was made to keep the discussion as straightforward and jargon-free as possible.

Following this introduction, there is a brief section that describes the essence of the MPO´s set of travel models, and then another section that describes in more detail each step of the travel-modeling process. The final section discusses how the model set is maintained and applied in projects. The reader is invited to contact Scott Peterson, Director of Technical Services, (speterson@ctps.org), at the MPO´s Central Transportation Planning Staff (CTPS), for additional information.

A set of travel models is nothing more than a tool, albeit a computationally complex tool. A travel model set predicts how many trips will be made by people in a given region on a typical day, where those trips will go, and what modes and routes those trips will utilize.

Travel models are built on the basis of people´s observed behavior. That behavior is typically measured from the results of a household travel survey. Statistical analysis of the survey data yields the models that describe this behavior. The models are then used to forecast future travel under certain assumed conditions: assumed population and employment patterns, assumed costs of the different travel modes, and an assumed transportation system.

Travel models exist to provide information about projects, programs, and policies for decision makers. In the context of projects, they are used to predict how many vehicles will use a new or modified roadway, or how many people will board a new or modified transit line. In the context of programs, such as that of the Regional Transportation Plan, they are used to predict total system-wide travel patterns. In a policy context, the travel model set has been used, for example, to inform MBTA decisions about new and modified fare policies.

The Boston Region MPO has always had a travel model set, and the staff maintains and applies (runs) it. It has undergone major improvements over the years: it was last completely rebuilt in the mid-1990s on the basis of new survey information. The model set is a so-called four-step one. It is so named because, historically, models of this sort were composed of four main models or model steps. The MPO´s travel-modeling process, as will be seen, is actually made up of six main steps.

Essential Features of the Model Set

The travel model set spans 164 communities in eastern Massachusetts—63 more than are included in the Boston Region MPO area. The reason for this greater expanse is that travel patterns in the MPO area are intricately linked to those in the adjoining regions. Currently, the model set represents weekday travel only. It is an aggregate model set, in application, in the sense that it represents and forecasts travel for groups of travelers, rather than for individual travelers. It is a planning-level tool, not an operations-analysis-level one.

The model set focuses on passenger travel within the region. There is, however, an adjunct commercial-vehicle model that is also run when necessary. Another adjunct model forecasts ground-access trips to Logan Airport made by air travelers. In addition, after the travel model set is run, the outputs from it are often entered into a model developed by the staff that, with the use of EPA-supplied emissions factors, computes regional pollutant emissions associated with the region´s vehicles. These adjunct models will not be mentioned further. More information about them can be obtained by contacting the Director of Technical Services at CTPS.

The model set is multimodal: it represents and forecasts travel by all of the region´s transit modes, by all types of single-occupant and multiple-occupant vehicles, and by walking. It represents travel in the region for all purposes, including work, shopping, recreation, and personal business. The forecasts that it yields are for four different time periods during a weekday: morning peak, midday, afternoon peak, and evening/night. In the discussion that follows, trip purpose and time periods (along with many other model details and nuances) will be largely ignored for simplicity. The reader is reminded, however, that much of the process described below is followed for each of eight trip purposes and each of four time periods.

Major Steps in the Travel-Modeling Process

The term four-step model refers to the third through sixth steps described below, all of which deal with the demand side of the process. We have added another step—a vehicle ownership model—and we also choose to refer to network and zone system development as a step, given its central role in the process. We therefore have six steps, as described below.

Network and Zone System Development

Prior to evaluating the demand for travel in the region, as is done in the model steps after this one, the region´s supply of transportation must be represented. This is done by building a computer model of the region´s roadways and transit systems. Once built, this transportation network can be used to calculate the distances, times, and costs of traveling from one place to another in the region. These measurements are then used in the demand-modeling steps described below, and in the final step, the modeled demand is loaded onto or assigned to the network to yield forecasts of traffic volumes on specific roads and transit ridership on specific lines. The network includes all transit lines in the region, all expressways and arterials, and many, but not all, collector and local streets.

This computerized representation of the region´s network has to be anchored spatially: that is, it must be tied into some representation of real points on the ground so that the travel demand calculated in the steps below can be related to it. The way this is done is by dividing the region into so-called transportation analysis zones (TAZs). These are loosely analogous to census tracts, but they are typically smaller (one square mile as opposed to about three for a census tract), and they are crafted to conform to the transportation network and natural features. Each TAZ is connected to the network so that the trips generated on the demand side of the model set can be loaded onto the network to obtain forecast traffic volumes and transit ridership. There are 2,727 TAZs in the region. (There are also several loosely analogous points along the periphery of the modeled region that allow simulated trips to pass into and out of it. The additional complexity these impose on the modeling process will be ignored here.)

Vehicle Ownership Model

This model predicts how many vehicles households will own, based on various different characteristics of the households. Vehicle ownership is important because it is associated with how many trips household members make and their probability of choosing transit. The household characteristics used to predict vehicle ownership are income, household size, and number of workers. In addition, household location has some influence on auto ownership. Therefore, the population and employment density surrounding the household´s location, as well as its accessibility to transit, are accounted for as well.

Trip Generation

This step predicts how many trips, for each of several trip purposes, a household will make on a typical day. Along with vehicle ownership, supplied by the vehicle ownership model just described, this model step uses household size and number of workers per household as inputs. As can be imagined, all other things being equal, households with more people, more workers, or more automobiles will make more trips. These relationships are all represented in the models. The generated trips are aggregated by TAZ.

While households produce trips, other places attract them. (Households themselves also attract trips, such as when a friend visits. For simplicity, this discussion ignores this.) Offices, schools stores—these are examples of the many kinds of places that attract the trips that households produce. To forecast these "trip attractions," another set of equations is used. Each equation relates the number of employees, by employment type (basic, retail, or service), to the number of trips attracted to that kind of place per day. All else being equal, larger establishments of a given kind have more employees, and therefore attract more trips per day. Retail establishments generally attract more trips per employee than office establishments, etc. As with the household ends of trips, these are aggregated by TAZ.

Some trips have neither end at a home, of course. A trip made from an office to a nearby deli at lunchtime is an example of such a trip. We have separate equations that forecast the numbers of these trips into and out of each TAZ.

Trip Distribution

All that is ascertained in the trip generation step is that each TAZ has so many trips emanating from its households and so many trips attracted to the various places within it. It is not known where the trips produced by the households are going to, or where the trips attracted to the various places are coming from. In other words, trip generation tells us how many trips begin and end in each TAZ. What it does not tell us is how all these “trip ends” link to form whole trips from one TAZ to another.

The trip distribution step provides this linking function: it creates trips from one place to another, from one TAZ to another. This is this accomplished by making two simplifying assumptions. The first assumption is that trips produced by households in a given TAZ will be attracted to other TAZs in direct proportion to the relative sizes of those other TAZs. This is known as the Law of Gravitational Attraction, and for that reason, trip distribution models of this sort are typically called gravity models. The size of a TAZ, in this case, is measured by how many trips are attracted to it in total (as has been estimated in the trip generation step). So, all other things being equal, in forecasting the number of trips produced by households in one TAZ that will be sent to another TAZ, the larger the other TAZ, the larger the number of trips that will be sent to it.

The second simplifying assumption is that trips produced in one TAZ will be attracted to other TAZs in inverse proportion to the travel time separating the two TAZs: the longer it takes to get to a TAZ, the fewer will be the number of trips sent to it from a producing TAZ. The travel times are computed from the transportation network, and they consider both the time to drive and the time to take transit, if available, from one TAZ to another.

Of course, most TAZs produce some trips that are attracted to that same TAZ, such as trips from home to the local elementary school. The trip distribution step is able to estimate these trips in the same way that it estimates trips between TAZs.

Mode Choice

From the previous two steps, the number of trips from every TAZ to every other TAZ is obtained. This information is organized into a matrix, often referred to by travel modelers as a trip table. In the case of the Boston Region MPO, a 2,727-by-2,727 matrix is produced, because there are 2,727 TAZs in the region.

One thing that is not known at this point is what travel modes will be utilized in making the trips from each TAZ to every other TAZ, and that is where the mode choice model comes into play. The mode choice model splits the trips in every TAZ-to-TAZ interchange by mode, based on the relative times and costs associated with each available mode. All else being equal, a given mode´s share of trips in a particular TAZ-to-TAZ interchange will increase as the time and/or cost of using that mode decreases relative to the other available modes. Other factors considered are travelers´ household size relative to household vehicle ownership, and TAZ population density (all else being equal, higher density indicates that more of the transit trips will have walking as the access mode).

The modes considered are auto drive alone, auto shared-ride, walk-access transit, drive-access transit, and walk all the way. As with the trip distribution model step, the times and costs associated with each mode for each TAZ-to-TAZ interchange are computed using the transportation networks. In the case of the automobile modes, in addition to time spent driving on the road network, any time spent getting into and out of parking facilities is included. The costs include operating costs (fuel, oil, maintenance), and tolls and parking charges, if any. The costs of owning a vehicle (insurance, depreciation, interest) are not included, as they are considered "sunk" costs: they do not factor into the day-to-day decisions of whether to use the vehicle for a particular trip.

For the transit modes, in addition to time spent onboard the transit vehicle, we calculate time spent walking to and from stations/stops, time spent waiting for the initial transit vehicle and any others transferred to, and any time spent driving or being driven to a transit station. Costs include fares, of course, as well as auto-operating costs and parking lot charges if transit is accessed by automobile.

The walk-all-the-way mode is free, and the time to walk from one TAZ to another is computed by applying 2.5 mph to the distance between them.

Route Assignment

The previous model step yields forecasts, for each TAZ-to-TAZ interchange, of the numbers of trips that will likely be made by each available mode. What has not been determined by the mode choice step, however, is how many trips will utilize which routes through the network. Most zonal interchanges have more than one roadway routing connecting them, and many in the core part of the region also have more than one transit routing connecting them. In the route assignment step, therefore, the auto mode trips are assigned to different roadway routes, and the transit trips to the available transit routes. Walk trips are not generally assigned. Instead, their numbers and spatial patterns, as output from the mode choice step, are summarized and examined.

The route assignment process is somewhat analogous to mode choice. A particular routing along the roadway system will be assigned more of a zonal interchange´s trips the cheaper and faster it is compared to other available routes. The same goes for transit. All of the time and cost components that are considered in trip distribution and mode choice are also considered here, the difference being that, in this step, the routes are considered as distinct from one another, not lumped together into TAZ-to-TAZ measurements of time and cost as in the previous two steps.

Congestion is accounted for in this process. Roadways that are congested will be able to process only so many vehicles in a given period of time. The excess vehicles are assigned to other routes that have available capacity. On the transit part of the network, crowding on buses and trains is likewise accounted for, as are the capacities of station parking lots.

The result of this model step is a representation of the region´s total trip making on the transportation network. Vehicle volumes on specific roads and transit boardings on specific bus and train lines are obtained from this model step and are used in project evaluation.

How Modeling Works in Practice

Travel-modeling practice at the Boston Region MPO can be divided into two major areas of effort: development and application. The staff is responsible for both areas.

Model Development

Model development is an ongoing task. With a set of models and supporting routines as complex and data intensive as these are, there is always the need to update and improve them. To be precise, there are actually two parallel versions of the same model that are maintained. One represents current travel conditions. The other is the same essential model, but its inputs are for a future year, typically a year two or more decades away. For simplicity, this distinction will be ignored, and the model set will be discussed here as if there is just a single one that represents both current and future conditions. Model development can be divided into three general subtasks as follows.

Minor Improvements

These consist of things like updating the speeds on various roads in the network and making the computer routines that run the models more efficient. This is an ongoing activity.

Major Improvements and Revalidation

This is more involved. Periodically, some major factor in the model needs to be changed. Perhaps the MBTA has changed its fares, the MPO has obtained an updated set of traffic and transit boarding counts, or we have decided to adjust the form of one of the models a little. When such a thing happens, the model set must be revalidated. This simply means that, after the change is made or new counts are obtained, the model set must be run and adjusted until it replicates observed travel data: it must be shown to be valid in its representation of what we see around us before it can be used to forecast. (When the model set was first created, it was validated in this way before it was used.)

Rebuilding

Rebuilding the model set from scratch is a major effort, and it only occurs every several years. At the Boston Region MPO it last occurred in the mid-1990s, as has been mentioned. It occurred then because the MPO had conducted a new household travel survey. The data from such a survey allows modelers to completely rebuild models because the data, when analyzed statistically, yield the form and numerical contents of models. (This is formally known as model estimation.) These data yield the parameters of the models—the numerical factors or coefficients, roughly speaking.

Examples of these parameters are the number of shopping trips that a household of three people with one worker make in a day, and the average time that a household member living in a particular part of the region takes to travel to work. Many of these parameters are fairly straightforward and easy to understand. Others are arcane to non-modelers. They exist, for example, in what are called utility equations in the mode choice models, where they measure the impact that changes in a particular mode´s travel times or costs would have on the probability of a household with certain characteristics choosing that mode. In all, there are scores of these straightforward and more arcane parameters scattered throughout the model set.

It is extremely important that the models be rebuilt periodically, because the underlying behavior of travelers that they are supposed to represent mathematically can change over time; without doing a new household survey, modelers would not know about those changes, and the model forecasts could therefore be in error. After a major rebuilding, the models are revalidated, as described above.

Regardless of what, exactly, has been done to the model set, once the changes have been made and it has been revalidated, it is ready to be used in the next model-application project.

Model Application

In any given project, several different future scenarios are typically analyzed. The scenarios may be, for example, variations of a transit line being planned—say, a bus-oriented scenario, a bus rapid-transit-oriented one, and a rail-oriented one. There are often variants of these main scenarios as well. In other projects, the alternative scenarios may consist of different assumed forecasts of household and employment patterns. In some projects, the alternative scenarios consist of combinations of transportation and demographic possibilities.

In each case, the relevant characteristics of the transportation services or demographic assumptions are incorporated into the model set, and the model set is run. (When the demographics change, the work is not trivial: all of the demand models have to be rerun to produce new trip tables.) A variety of outputs will be obtained from the models, and summarized and reported. In addition to yielding traffic volumes and transit ridership on the region´s network, a host of other outputs are available. The possible outputs are too numerous to mention here, but they include anything from mode shares by TAZ or collection of TAZs, through travel times from any point to any other point under congested conditions, to vehicle-hours traveled. As mentioned earlier, air quality impacts can be estimated by entering outputs into an adjunct model.

The travel model set is always run for the entire region, even if the analysis at hand only concerns a single corridor. In the latter case, there might be some detail added to the network and TAZ system in the corridor of interest, and there might be some refinement of the demographic assumptions in that corridor, but the model set will still be run for the entire region. That is done, in part, because we want to be sure to capture travel impacts outside of the corridor that may result from changes within the corridor; this is particularly a concern in transit project analysis. Another reason that is done is that it is very time consuming to extract only the portion of modeled area that is of immediate interest.

It should be noted, finally, that the travel model described herein is a planning-level tool. For projects requiring highly precise traffic forecasts that take into account all of the complex, underlying dynamics of traffic flow, and for highly precise transit-vehicle-operations analyses, a different tool is required. In those cases, a micro-simulation tool would be used. CTPS has several such tools. More information can be obtained about them by contacting the Director of Technical Services at CTPS.