Discussion-piece #4
Ridership and User Benefits from Variable Trip Tables
Federal Transit Administration
June 6, 2006

1. Motivations. FTA has long required that project sponsors maintain fixed person-trip tables in the development of ridership forecasts for baseline and build alternatives that FTA will use to evaluate proposed New Starts projects. Historically, this requirement has reflected both the absence of methods to predict changes in trip tables caused by major transit projects and the need to avoid yet another source of potential over-estimates of transit ridership on New Starts projects. One consequence of this policy is that it precludes the estimation of project-specific mobility benefits associated with changes in travel patterns caused by major new transit facilities. Over the past several years, however, the problem of overestimates may have moderated and methods to predict transit-caused changes in trip tables have become more common. Therefore, given FTA's ongoing efforts to credit proposed projects with their full range of likely benefits, a reconsideration of the fixed-trip-table policy may be in order.

2. Variable trip-tables in travel models. Trip distribution models predict the attraction locations of the trips produced in each zone. These models have historically adopted the "gravity" formulation and have used highway travel time as the measure of separation between zones producing and attracting trips. Often, these models have been "doubly constrained" with iterative methods that adjust the predicted trip tables to ensure that the total number of trips attracted to each zone matches the number of trip-attractions predicted (in trip generation) for that zone. Over the past several years, an increasing number of models have added transit service levels to the measures of separation between zones. This trend has been most evident in new trip distribution models that use logit formulations. Consequently, some current model sets produce person-trip tables that, at least for some trip purposes, are sensitive to transit service levels. Forecasts for an alternative that introduces a major transit facility will include person-trip tables that have more travel in the zone-to-zone cells well served by the new facility than in the baseline alternative – and will yield higher overall transit ridership, more riders on the new facility, and more mobility benefits compared to a fixed-trip-table forecast.

3. User benefits from variable trip-tables. FTA's current measure of user benefits is derived from the denominator of the (logit or nested-logit) mode choice model. The denominator captures all attributes of all modes considered by the model, and is therefore an ideal descriptor of the overall mobility – and changes in the overall mobility – of travelers. A similar measure can be computed from the denominator of (logit) trip distribution (or "destination choice") models. In the easiest computational case, the destination choice model will employ a logsum term from the mode choice model as the measure of separation between zones:

$$(1)\ \text{Probability of travel from i to all possible j} = \frac{\exp(C_{ls} \times \text{logsum}_{ij}) \times \text{size}_j}{\text{sum}_j[\exp(C_{ls} \times \text{logsum}_{ij}) \times \text{size}_j]}$$

where i and j are the production and attraction locations, logsum is the natural logarithm of the denominator of the logit mode choice model, C_{ls} is the coefficient on the logsum variable, and $size_j$ is some measure of the size or attractiveness of j. The logsum variable extracted from the mode choice model for use in this destination choice equation is exactly the same measure from mode choice used in FTA's current calculation of user benefits. The denominator of this destination choice model measures the total mobility of travelers producing trips from zone i, choosing among all of the zones in the region as possible destinations. This structure provides for the easiest user-benefits computations because it effectively combines the mode- and destination-choice models into a large nested-logit model that is internally consistent and precise in its accounting of the effects of changes in any attribute of any mode. The denominator converts easily into a generalized price of travel to possible destinations through its natural logarithm, divided by the mode-choice coefficient on in-vehicle time scaled by the logsum coefficient:

$$(2)\ \text{Price of travel from i to all possible j} = \frac{\ln\{sum_j[\exp(C_{ls} \text{ x } logsum_{ij}) \text{ x } size]\}}{Civt \text{ x } Cls}$$

This calculation is exactly analogous to the calculation of the price of travel from i to each j (using the $logsum_{ij}$) that is done in the current FTA user benefits computations. The change in price caused by a major transit investment will appear in both its mode-choice component and this more inclusive destination-choice component. User benefits derived from both mode-choice and destination-choice effects are simply the product of this more inclusive price change times the number of travelers producing trips from zone i. Consequently, with local models structured along the lines sketched above, an internally consistent way is available for computation of user benefits that include the benefits derived from changes in travel patterns in response to transit investments.

4. <u>An example.</u> The attached Excel file illustrates the approach with home-based-work-trip models structured as those sketched above. The first worksheet provides a map of the 10-zone system for which the models are applied. The second worksheet summarizes the model parameters (highlighted in green), productions and attractions, and zone-to-zone impedances for highway and transit travel; it also provides entry locations (highlighted in yellow) that can be used to change these zone-to-zone impedances to represent highway and/or transit improvements or changes in development patterns. The third worksheet computes transit shares for both the baseline and build alternatives, and computes the logsum terms from the denominators of these logit calculations.

   The fourth worksheet applies the destination-choice model for both alternatives. It first uses the logsum terms from the mode choice worksheet to compute the singly constrained results, and then adjusts the "attractiveness" terms for each destination zone j to iterate the model to a doubly constrained result. The differences between the two results are evident in the deltas they produce. The singly constrained model produces a baseline trip table whose column sums do not match the predicted trip attractions in each zone. More importantly for this example, the column sums change between the baseline and build alternatives. In contrast, the doubly constrained model (after three iterations) matches the column sums to the trip attractions; it therefore has identical column sums for both baseline and build.

The fifth worksheet "UBfromMC(1)" computes user benefits from the mode choice prices (using the current FTA approach) for the singly constrained trip distribution model. In each zone-to-zone cell, the worksheet converts changes in the mode-choice logsum term between the alternatives into price changes and multiplies them by the (fixed baseline) number of trips. The sixth worksheet "UBfromMC(2)" attempts to compute user benefits from the same mode choice prices but bases the calculations on the different person-trip tables for the baseline and build alternatives that emerge from the doubly constrained trip distribution model. This worksheet uses the price and trips in each cell to compute travel "expenditures" in the cell, and computes user benefits as the differences in these expenditures.

The sixth worksheet "UBfromTD" computes user benefits from the trip-distribution model itself using equation (2) above to calculate trip-distribution prices and price changes, for both the singly and doubly constrained trip distribution models. For the build alternative included in the example, the results for the singly constrained model are 113 hours at the trip distribution, 2.5 hours more than from the standard approach using mode choice alone. Consequently, the standard FTA approach with fixed trip tables captures 98 percent of the benefits. The results for the doubly constrained model are more complex. While the net benefits at the trip distribution level are the same as for the singly constrained model, both benefits and disbenefits occur for individual production zones. Benefits for trips produced in zone 2 increase to 138 hours – 25 hours and 23 percent more than the zone 2 benefits in the singly constrained model. This increase is the direct result of the iterations in the doubly constrained model: they increased the number of trips from zone 2 to zone 1, the cell with the 20-minute improvement – as part of the effort to reconcile the shortfall of in the zone-1 column-sum in the singly constrained model. This increase is exactly offset, however, by the disbenefits incurred by travelers from all other zones. These disbenefits occur because the iterations in the doubly constrained model decreased the relative attractiveness of zones 2 through 10 to shift some trips produced in all zones to zone-1 destinations. Consequently, the iterations make the effective price of travel increase to zones 2 through 10, and travel to those zones incurs negative user benefits. A real-world interpretation is that the improved mobility from zone 2 to zone 1 increases the competitiveness of residents of zone 2 in the competition for jobs in zone 1. Because there are a fixed number of jobs in zone 1 (in the models, anyway), some residents of other zones lose out, and have to take jobs in second-best locations. The reverse effect happens with a decrease in travel time to a zone that has a surplus in its column sum compared to its trip attractions in the singly constrained model (zone 9, for example): the iterations increase the attractiveness of zones 1-8 and 10, so the secondary user benefits to these zones would be positive.

An important insight is evident in the result that the TD-level benefits are smaller that the computed MC-level benefits for the doubly constrained model. This outcome demonstrates the infeasibility of computing benefits from changes in trip tables with mode-choice-level prices of travel. The computation is infeasible because the mode choice model is entirely unaware of the adjustments made to the attractiveness of each destination j during the iterations to doubly constrain the distribution results. This inconsistency confounds the calculation of trip-distribution effects with mode-choice-level information.

The last worksheet applies the same mode-choice and trip-distribution models to a very simple problem with one production zone, two attraction zones, and two travel modes.

5. Potential problems. While it appears that plausible user benefits can be computed for trip-distribution effects under ideal conditions, a number of barriers remain to the implementation of this idea nationally. Probably no model set in the United States employs trip distribution and mode choice models precisely consistent with the structure above for all trip purposes, socio-economic markets, and times of day. Some use the above structure for work travel or for travel generated by car-less households, but rely on highway travel time alone to represent impedance in for other purposes and markets. Some models use a harmonic mean formula to compute composite highway-and-transit impedances weighted by mode shares from the mode choice model; that formulation might support workable user benefits calculations based on the trip distribution model in spite of inconsistencies between the composite impedance measures in mode choice and trip distribution. Most models partition zones by walk-accessibility and apply the mode choice model separately for the individual markets; but no trip distribution models recognize the different levels of mobility that result for the various walk-access markets. More significantly, of course, is the reality that most model sets continue to use highway travel time alone to represent impedance in trip distribution – and therefore provide no basis for computing transit-caused user benefits from trip distribution.

6. Conclusions. At least one way exists in conventional practice to compute user benefits that include the effects of variable trip tables in addition to the mode-choice-level benefits currently recognized by FTA from fixed-trip-table calculations. However, relatively few model sets in the United States have the properties necessary to support these calculations. Further, the level of effort to maintain such model sets is higher – perhaps significantly higher – than to maintain models that use highway-only impedances in trip distribution and thereby ignore transit influences on overall travel patterns. Finally, experiments with the attached spreadsheet suggest that the additional benefits from variable trip tables are small – apparently well less than 10 percent. Given the absence of a ready basis nationally for computation of these benefits and their apparently modest size, FTA continues to view the project-specific computation of variable-trip-table benefits as a low priority. FTA has no current plans to expand the user benefits measure to capture these benefits and will, instead, continue to rely on the allowances for missing benefits (outlined in discussion-piece #1) to cover their absence.