Task 13.1 Preliminary Analysis of the Effects of Accumulation of Model Errors

We need a philosophical position here. One position is that the matrices extracted from the observed data are the best estimate and that our model is attempting to synthesise these in a way in which we can then forecast them. This is a sensible position when we have high sampling rates (such as in the rail survey) and we can expect the observed matrices to be a close fit to reality. But it is not so sound for the low sample household survey where our expansion factors are subject to error as are the average travel behaviour characteristics obtained from subsets of the household data.

So our main aim in this task is to understand where our fully synthetic estimates differ significantly from the original calibrated model and the observed data. If in the independent validations in the subsequent task it appears that the differences introduced by the fully synthetic model are crucial then we may need to correct them. But, given the household survey sampling rate it is also conceivable that some of these differences may improve model fit.

In general, the aim is to see how the model varies from calibration results and the original household survey data. Given the sampling statistics, the checks should be done at TLA and sector levels only. In the following 'calibrated model' refers to the results of the model during the estimation process, while 'fully synthetic model' refers to the final implementation.

(1) Planning data: the model is calibrated primarily on the household survey data, this being expanded to planning data household totals. Thus the population and number of households by area should be compatible between the expanded household data and the planning data. This should be verified – no action required unless something unexpected turns up.

(2) Family structure model errors: differences from the household survey, on which the trip end and car ownership models were calibrated, will arise mainly from the forecast distribution of households by type, as the population type distribution is constrained to the census. The synthetic household distributions should be compared with the household survey. I would think that it is not worth trying to verify the cross-distributions – we presumably have some idea of the errors from the calibration report. No action required unless something unexpected turns up.

(3) Car ownership model errors: compare observed, calibrated model and synthetic car ownership model values of p0, p1, p2+, cars per household and total number of cars. The errors/differences arise from:

   – the process of fitting the model to census data, whose car ownership will be different from the household survey, and

   – running it on the planning data/family structure model which, given that the planning data is the same as used for household survey expansion, should only introduce errors associated with the family structure allocation of household types;

   – no action required unless something unexpected turns up.

(4) Trip end model errors (by purpose, P and A): the production and attraction model fits by sector and TLA will be changed by: the planning data, family structure and car ownership model inputs and, possibly, by the balancing process. In the calibration of these models we sought to reproduce sub-regional (territorial local authority - TLA) totals, but did not constrain sector values. At this stage we need to confirm that TLA totals are in good agreement with observed data and we need to understand the spread of sector errors (by comparison with the observed data and calibrated model sector values). I do not think we can consider any sector-based corrections, but the possibility of re-tuning the TLA totals will need to be borne in mind. So keep a record of the factors needed to correct any significant TLA errors (>5%, say); we may need them.

(5) Networks. The public transport network should be identical to that used in model calibration so will introduce no new errors[1]. But the generalised costs extracted from the am and interpeak road networks will have changed through the iteration processes as the link speeds are tuned to the synthetic road matrices. We should produce comparisons of synthetic road generalised costs with those used in calibration. Perhaps the easiest would be to aggregate these to sectors and TLAs by weighting using the observed matrix. See Task 13.4 for actions.

(6) Distribution & mode choice errors: these should be analysed in 2 steps:

- the first step is the first iteration, ie the matrix forecasts using the networks used in calibration, before feedback to the road network; this will tell us how the sub-model error accumulation has affected the matrices;
- the second step is to analyse the matrices after the iteration process has converged; this will tell us how these matrices have changed as a result of using the converged road speed estimates.

Diagnostics/comparison against observed and calibrated matrices should be by purpose and mode:

- average trip length
- overall mode shares
- total trips
- and matrices at TLA (and possibly sector?) level

(7) Time period factors: we need to repeat the matrix validations in 6 for the time period matrices (all purposes should be aggregated together) – comparing observed, calibrated and synthesised. As before this should be done in 2 steps. See Task 13.4 for actions.

(8) Assignment: see Task 13.2.

[1] In fact this expectation was wrong - we were provided with an updated, different public transport fares matrix late in the project.