National Transport Model - Working Paper 1

Contents

Chapter 1 - Overview

Introduction

This is the first in a series of Working Papers describing the National Transport Model. The model is subject to continuous development; this paper describes version 2.1, the model as it existed in November 2002 and as used for the First Progress Report on the Ten Year Plan.

A higher-level paper, "*Overview of the NTM*" presents an overview of the model, something of its history, a summary of forecasting assumptions, and illustrative results.

This Working Paper describes the structure of the model and how the different components interact.

Working Paper 2 - model inputs gives more details of forecasting assumptions, representation of policies within the model, and the data that went into it.

Working Paper 3 - model performance describes the model calibration, gives examples of the range of disaggregate results, presents key elasticities, and sensitivities of model results.

Working Paper 4 - welfare module looks at how model results are used to derive estimates of economic welfare.

Together these Working Papers are intended to cover all model integration issues. More details of the development of individual model components are contained in a series of research reports, written by the Departmental staff or external consultants who did the original development work.

The main purpose of the NTM is to indicate to policy-makers the likely impacts of national policies, and the cumulative national impact of policies that can only be applied locally. For some aspects of policy appraisal, it is important to be able to consider variation by time of day and day of week. For some aspects, a high level of market segmentation is appropriate. For some aspects, geographic detail is required. Combining these different dimensions within a simple model structure was judged to be computationally infeasible. Instead, the approach adopted was a modular structure in which different modules have different dimensions.

The following diagram shows this modular structure, which was adopted for a number of reasons:

- The model is multi-modal, including all significant surface modes of travel
- It relates total travel to land use and behaviour
- Because the Department had no tradition of national modelling, it was developed incrementally, linking together a number of existing components with new modules
- It makes good use of available data sources, and didn't require any new data collection
- It allows representation of a wide range of policies combined into a single structure
- It is broadly speaking at the limit of what can be done with PC technology and model runs that take no longer than overnight
- It addresses the Department's principal transport targets, which are formulated in terms of road congestion and travel by different modes

Figure 1: structure

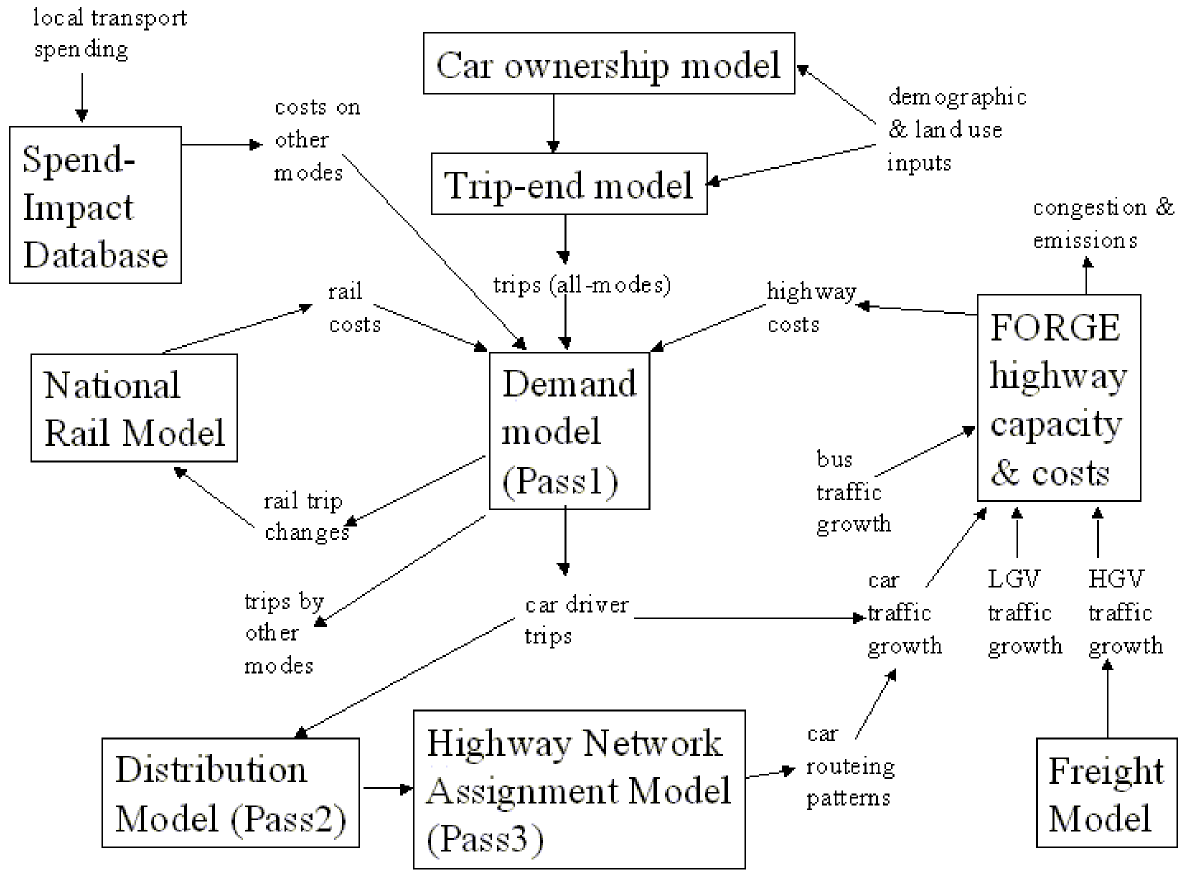

Structure & index to chapters

Figure 1 presents an overview of the modelling system, illustrating how the different components relate to each other. (Any such diagram is a compromise between being comprehensive and being comprehensible - different reports present this structure in different levels of detail).

In the centre of the system is the Demand model (sometimes referred to as "Pass1"). This is where mode choice is modelled. The inputs are total numbers of trips (which are taken to be invariant to cost) and the costs by each mode. The outputs are numbers of trips by each mode, highly segmented by trip length, trip purpose, and person type. This model is described in chapter 3.

At the top of the diagram, the car ownership and trip end models estimate the number of trips in each future year, as a function of demographic and land use inputs, and various economic forecasting assumptions. These models are presented in chapter 2.

At the bottom of the diagram, there are two possible paths through the system. The Pass2 and Pass3 models take a long time to run, and are currently run only once for each future year. A synthetic matrix of car driver trips is estimated and assigned to a GB highway network, and data from this run is stored as a set of routeing patterns (referred to as "mileage profiles"). These relate car trips between different origins and destinations to car traffic on different types of road within different parts of the country. For policy tests, these fixed routeing patterns are used to convert trips into traffic without re-running the Pass2 and Pass3 models. More details are given in chapter 4.

In the bottom-right hand corner of the diagram, a freight model forecasts HGV traffic growth. Modelling of HGV and LGV traffic is described in chapter 5.

FORGE is the model which gives estimates of highway congestion and pollution, and models the traffic response to charges and changes in speed on particular road types. This component has been developed from the "fitting-on" module in the 1997 National Road Traffic Forecasts, and is presented in chapter 6.

On the opposite side of the diagram, the National Rail Model (NRM) assigns the rail passenger trips from the Demand model to a detailed geographic network of rail services. The resulting journey time, crowding, and rail fare outputs are sent to the Demand model, and the assignment results are used for detailed analysis of policy impacts on rail. The NRM is the subject of chapter 7.

Chapter 8 deals with the Spend Impact Database (SID for short) and how the impacts of future year local transport spending are estimated.

The elements of the model having been covered in some detail, chapter 9 sets out a summary of the data flows between model components, and the levels of disaggregation involved (both geographical detail and segmentation of the travel market). Chapter 10 briefly describes the processes of calibrating and running the model, and how it is used to appraise the impact of different transport policies, including the estimation of economic welfare.

Chapter 2 - Car Ownership and Trip End Models

The car ownership module is based on that used in NRTF97, but has been changed to be responsive to car purchase and ownership costs, and includes an estimate of the impact of company cars. The trip end model is now considerably more detailed spatially than the Department's previous National Trip End Model, as well as including trips by all modes (rather than just by car) and having greater segmentation of traveller types.

Demographic / land use data

The main input to both car ownership and trip end models is the planning database. This contains projections of

- population by age band,
- one-person and multi-person households,
- employment by sector,
- proportion of "working age" people in employment.

The figures go back to the 1991 Census and forward to 2036, although the later years are inevitably based less on information and more on extrapolation.

The figures at county or district level are based on published projections (ONS figures for population, former DTLR figures for numbers of households, Cambridge Econometrics employment growth projections applied to an ONS base).

The geographical level below districts is a system of around 1200 zones covering GB. Zones nest within local authority districts, and are designed to reflect urban/rural development patterns. Within a district, a set of parameters determines the proportion of growth in households and employment that occurs in the urban and rural zones of the district. These parameters could be used to reflect the impact of planning policy as it affects urban concentration, but this facility has not been used in the current published results. The default values reflect ONS trend-based projections.

The level below zones is wards - each zone is a contiguous group of Census wards. Some parts of NTM use ward level data, but we have no information on likely future patterns of development at this level of detail. The assumption made is that zonal growth is spread uniformly over all wards within the zone.

Inputs to car ownership model

The car ownership model is run at zonal level, and thus picks up differences in car ownership between different types of area. Car ownership is highly correlated with income, although for a given income, rural households are more likely to own a car than urban households, other things being equal. Smaller urban areas are intermediate between rural areas and larger urban areas, as might be expected. Other than the planning data, the inputs to the model are:

GDP - growth in GDP per household is taken as an overall growth factor for household income. No differences in income growth between different types of household or different areas of the country are modelled, although the implications of planning data changes are followed through. (E.g. if an area has above-average growth in employment, then to the extent that the average employed household is richer than the average non-employed one, the model results will be consistent with faster income growth for that area).

Licence-holding - a sub-model predicts total national growth in the proportion of adults holding driving licences. This uses a cohort approach that models males and females separately (succeeding generations reach higher levels of licence-holding, reach it sooner in life, and reduce their ownership more slowly in old age). This national aggregate projection is used to moderate the observed time trend in licence-holding statistics.

Company car ownership -changes to taxation of company cars are seen as unlikely to have a large direct effect on car ownership. Most company car owners either need the car for work or have sufficient income for their household to own at least one car anyway. But statistically, presence of a company car makes a household more likely to own a second vehicle than if the first car is a non-company vehicle, even allowing for income effects, and this effect is modelled. Given an estimate of change in numbers of company cars resulting from some policy, the model will predict the resulting change in second car ownership.

Ownership and running costs - terms representing the impacts of car ownership cost and running cost were introduced into the model, based on elasticities derived from aggregate time series work. Car ownership cost is currently projected to continue its trend decline, although this is a possible policy variable.

Car running cost is assumed constant at this stage. To avoid introducing a more complex model looping structure, the car ownership model predicts car ownership at constant running costs, and the elasticities of car travel to running costs in later stages of the model include impacts on the car stock.

Car ownership model - general

Given planning and economic inputs, the main function of this model is to predict in each zone the numbers of persons in households owning no car, one car, or two or more cars. These figures are then passed on to the trip end model.

The approach used is prototypical sampling. It starts from a file of survey data, each record referring to one household, giving various characteristics of the household - income, household structure, area type - and the number of cars owned. For each future year and each zone of the country, the program calculates a weight to be given to each household record that will cause the weighted sample of households to match zonal totals of population by age, gender and employment status. The household-based car ownership model is then applied to these weighted samples, ensuring that the overall income growth matches that which is input.

The household-based model was calibrated on FES/NTS data from a number of years 1971 to 1996, and thus picks up both time series and cross-sectional effects.

Trip end model

National Travel Survey data suggests that total trip rates are largely stable over time. The trip end model thus concentrates on cross-sectional variation. The model assumes that, once all trips (including short walks) are incorporated, trip rates for each category of person are constant. So that, whilst car owners make more trips than non-car-owners, and working age people more trips than elderly people, behaviour within each category will remain the same.

The trip end model operates at a high level of disaggregation, as analysis of NTS data showed significant variation in trip rates across a number of different dimensions. Each combination of eleven person types and eight household size / car ownership categories - 88 combinations - has its own trip rate for each of fifteen purposes. Variation of trip rates by area type was examined, but proved to be non-significant in most cases. Having incorporated variation by car ownership and employment status, both of which are highly correlated with income, any remaining variation in trip rates with income was found to be not significant.

The trip rates are derived from NTS, and are not normally varied to represent policy impacts, although we might conceivably want to do model sensitivity tests where a proportion of trips for particular purposes are assumed to be replaced by telecommunications. NTS includes information on where each respondent's home is, but very little on the location of their trip origins and destinations.

Applying the trip rates to estimate total trips is straightforward; most of the work done by the model is involved with allocating trip-ends to wards.

Each trip has two ends. The attraction end is the end where the activity that is the reason for the trip is located; the production end is the origin of the outward journey and the destination of a return journey. Around four-fifths of all trips are home-based. The number of trip attractions is constrained to be equal to the number of trip productions, for each of a series of "balancing areas" based on journey-to-work areas from the 1991 Census.

Home-based trip productions are located in the zones where people live, which is known from the planning data. All trip attractions are distributed across wards according to some land-use indicator statistic. For example, for commuting trips, the statistic used is total employment. For shopping trips, it is total retail employment. For trips to visit friends and relatives at their home, it is total households. Note that one retail job might attract less than one commuting trip per day (on average people travel to their usual place of work on around 70% of all working days) but will also attract a much larger number of shopping trips per day.

Finally, the distribution of non-home-based productions follows the distribution of home-based attractions, so that for example if people go to work or shop in an area, it will become the production zone for a number of onward non-home-based trips.

Whilst the trip production figures are segmented by different household types, person types etc, the attractions are not. The trip end model knows nothing about which people travel to where; linking the two ends of the journey is done at the following stage. Having produced figures for trip-ends by purpose at ward level, these can be aggregated as required by different models within the NTM framework.

For more detailed information, see

Research report #1 - car ownership model

Research report #2 - trip end model

Chapter 3 - the Demand Model (Pass1)

The Pass1 component of NTM models the main demand responses - distribution and mode choice. It has a low level of geographical detail, and a high degree of market segmentation, distinguishing the different costs faced by different segments of the population and their different sensitivities to different elements of travel cost.

Geography

The following 9 area types are distinguished:

- Central London
- Inner London
- Outer London
- Inner Metropolitan areas (main cities of conurbations)
- Outer Metropolitan areas (remainder of former metropolitan counties)
- Urban big (population >250K)
- Urban large (population 100K to 250K)
- Urban medium (population 25K to 100K)
- Urban small and Rural (population <25K)

A limited regional dimension has been introduced, dividing some area types into three broad areas of the country. This is to better represent policy options involving investment in particular rail corridors. With this split in place, there are 15 area types. Each of the 1200 model zones is allocated to a single area type.

Travel options and travel costs

Given the number of trip productions in each area type, the model represents a structured set of three choices faced by travellers:

- how far to travel - which of 13 distance bands in NTS (from under 1 mile at the shortest to over 300 miles at the longest)
- within that distance band, what area type to travel to (the choice will vary by area - e.g. starting from Central or Inner London, only London destinations are available within the shortest distance bands)
- having chosen a destination area type, which of six modes of travel (walk, cycle, car driver, car passenger, bus, rail) to use to get there.

The proportion of travellers choosing each option is based on "generalised cost" - a linear combination of time, money and other factors, expressed in time units. E.g. for car drivers, the elements of generalised cost include:

- Travel time
- Time to find a parking space
- Money cost of parking
- Money cost of fuel for the journey

- Modal constant term representing comfort and convenience of car travel.

The costs of travel at the third level - mode choice - are either averages from more detailed models within the NTM structure, or broadly representative (e.g. bus fares for short trips between certain area types are represented as a cost of 50 pence plus 12.5 pence per mile).

The costs of each option at the second level - destination area type - are an appropriate weighted average of the third level costs, plus a constant term calibrated to ensure that the number of trips choosing to travel to each area type is consistent with the trip attraction totals from the trip end model.

The costs of each option at the first level - choice of trip length - are a similar composite of second level costs, plus a constant calibrated to ensure that the model reproduces NTS trip length distributions in the base year.

These travel costs can be changed in different ways to represent the impact of policies - see chapter 8.

Market segmentation and values of time

The same structured set of choices is modelled separately for different segments of the total travel market - each a combination of person type, household car ownership, and trip purpose. Pensioners and children get discounted fares on public transport modes. High earners, travellers on Employers Business, and households with high car ownership are treated as having higher values of time.

The value of time affects the relative weight of money costs and time costs within the generalised cost formulation. As incomes rise, it is generally believed that values of time increase (but not necessarily proportionally). Within the model this is represented as money costs becoming less significant to the traveller, so there is a shift towards more expensive modes and longer travel distances. This is the mechanism which the model uses to explain increases in trip length over time.

With generalised cost in time units, the model assumption is the conventional one in transport modelling - that one minute of time remains equally significant to the traveller in the base and future years. As a sensitivity test, and as a way of generating a range about model results, we have also undertaken model runs where the value of time is assumed constant over time. This assumption would also give a constant impact of a £1 change in money cost over time, other things being equal.

Inputs and Outputs

The inputs to the Pass1 model are the trip-ends, elements of travel cost, and changes in values of time over time. The output is a database containing the numbers of person-trips

- of each segment (person type / household car ownership / trip purpose)
- in each production area type
- choosing each mode / destination area type / distance band combination.

Using an average trip length for each distance band, traffic figures (person-miles) can be produced as well as results in terms of trips.

For more detailed information, see

Research report #4 - demand model

Research report #5 - demand calibration

Chapter 4 - Highway distribution and assignment

The aim of this part of the process is to convert numbers of car driver trips between origin and destination area types into traffic (vehicle kilometres) growth on different types of road in different area types.

One of the observable realities that the model seeks to explain is faster traffic growth on motorways and trunk roads than elsewhere. In order to model a choice between road types, a fine-zoned model is required. (A coarse-zoned model would tend to load all traffic onto the major roads).

One cause of this differential growth is trip lengthening - a greater proportion of trip mileage is on major roads for long trips than for short trips. Another cause is increasing values of time over time, making longer but faster routes more attractive than slower direct routes on minor roads, for the same journey. The model is set up to represent both of these effects. It cannot yet deal with the other main factor which is land use effects - zones that are well served by motorways and trunk roads being more attractive to development and thus gaining trips faster.

The full process involves:

- Disaggregating the 15x15 car trip matrix emerging from Pass1 into a 10,000x10,000 matrix
- Assigning this matrix to a GB road network containing >100,000 links
- Comparing with base year traffic in order to get a file of traffic growth to use in subsequent processing.

Because the full process of matrix disaggregation and assignment is slow and cumbersome, it is normally run only once for each future year, and not re-run for each policy test.

In order to set up a quicker process, from the results of each future-year assignment, we tabulate "mileage profiles" - information on car routeing patterns, embodying knowledge of which area types are traversed by journeys between which other area types, and what mix of road types is used for different lengths of journeys.

For each policy run of the model, we derive an estimate of future year traffic on each road type as the product of the Pass1 results and the mileage profiles, e.g.

Traffic on rural Motorways in East Midlands	=	**Σ** distance bands O&D area types	**trips**	**X**	average mileage on rural motorways in East Midlands driven in a trip of given length

In this example, trips between two zones in London would be most unlikely to do any mileage on East Midlands motorways. But an average trip with one end in London and a length of 200-300 miles might spend a significant proportion of its total mileage on East Midlands motorways.

The following sections describe very briefly how the distribution and assignment processes work.

Distribution model

The Department does not have a national observed trip matrix for highway trips. The model synthesizes such a matrix, using a set of distribution models. These relate the estimated number of car trips between each pair of zones to the journey cost, derived from the detailed highway network. The estimates of trips are constrained to match trip-end totals at the level of detailed zones, and Pass 1 results in terms of trip length distributions and numbers of trips between each pair of area types.

This distribution modelling is undertaken in two stages, moving from the 15x15 zone results in Pass1 to an intermediate zoning system of area types within each county, roughly 250x250 zones. These are then further disaggregated to 10,000x10,000 in a set of regional distribution models, as the PC software available at the time of development couldn't cope with a single distribution model that big.

In each case trip-ends are used to weight the individual destinations appropriately, and a travel cost (deterrence function) is derived from highway travel times in the base year assignment model run.

Assignment model

The model network currently represents all A-roads and motorways, roughly three-quarters of B-roads, and one-quarter of C-roads and unclassified roads. All-or-nothing assignment is used throughout, with a spread of traffic across different routes resulting only from the wide range of origins and destinations. Base year link speeds are used, derived from speed-flow relationships and traffic counts where available, and broad area-type averages for most of the minor roads where traffic count data is more sparse.

For calibration, assignment runs were undertaken with a number of different values of time, and the value selected was that which most closely reproduced the base year distribution of traffic between road types at the national (GB) level.

The assignment process is not intended to give accurate traffic levels on individual roads. The main purpose of the assignment model is to set up correspondence between trips (by distance band and origin/destination area types) and traffic (by road type, region, area type) for use in policy runs of the model.

For more detailed information, see

Research report #4 - demand model

Chapter 5 - Freight Modelling

The freight model within NTM was developed by consultants MDS Transmodal, and is known as GBFM (Great Britain Freight Model). This has been developed from two pieces of earlier work - one used in the original Ten Year Plan analysis, looking at forecasting national trends, and one for the SRA, which used detailed geographic networks.

The earlier forecasting model had six stages, modelling:

- Tonnes lifted (freight volume for each of 15 commodities)
- Length of haul (tonne-kilometres)
- Modal split
- Categories of HGV (choice between six size categories)
- Choice of route & HGV type
- Loading factors (converting to vehicle-kilometres)

To determine tonnes lifted, we looked at trends and relationships with GDP for each of the major commodities. For each commodity, there was also a trend in length of haul. This trend was used to factor a base year county-level trip matrix (built from CSRGT data), typically increasing the numbers of long trips and reducing the numbers of short trips.

This model could divert traffic to rail, but didn't have a base matrix of rail traffic. Combining the earlier model with the work for SRA remedied this to give a true multi-modal model. The other major changes are that road costs are derived from a highway network that is a simplified version of that used in the Pass3 model, and there is more segmentation of HGV types.

For each cell of the county-to-county matrix, a logit model calculates the modal split, based on detailed models of the costs of operation for road and rail. Because rail has lower per-mile costs but high access costs, the model tends to give a higher rail share of the freight market for counties which are further apart.

For each commodity/county-pair combination, the model applies a fixed split between six categories of HGV (based on size and number of axles), so that the fleet mix changes in response to trip lengthening, or to changes in the mix of commodities carried, or to changes in costs.

There is then a more detailed split into 18 HGV types, each of which has its own cost function, and thus its own routeing pattern through the road network. A further logit model spreads the tonnage of freight between HGV types within each category, according to cost. Each detailed HGV type has its own average load factor (including an allowance for empty running and typical patterns of use).

Other categories of traffic

With the highway assignment model giving an estimate of growth in car traffic, and the freight model giving growth in HGV traffic, it remains to estimate traffic growth by light goods vehicles (LGVs) and public service vehicles (PSVs - buses and coaches).

There is little data on the use of LGVs. These vehicles are used for a mixture of trips, including:

- some trip purposes covered by the National Travel Survey, which are assumed in the modelling to be car trips
- some journeys which are essentially for the movement of goods, and thus might be approached as an extension to the freight model

- some trips (such as a series of calls made by a plumber in the course of his work) which do not fall into either category.

and thus any consistent behavioural approach would be relatively complex. This is an area flagged for further development of the model in the longer term. Currently, LGV traffic is projected by a simple time series model relating LGV km in a given year to the levels of GDP and fuel price.

Like LGV traffic, PSV traffic is modelled for the purpose of assessing the impact of these vehicles on congestion. The number of bus-kms in peak periods is assumed to grow in proportion to passenger numbers, but off-peak bus kms remain unchanged (on the assumption that off-peak any change in bus passenger numbers alters load factors but does not increase or decrease the service level).

For more detailed information, see

Research report #6 - freight model

Research report #9 - LGV growth

Chapter 6 - FORGE

The FORGE component is a development of the approach used in the NRTF97 forecast to constrain traffic growth to the capacity of the road network.

It is built around a database of base year traffic (the current FORGE base year is 2000). Each cell in the database contains the total traffic for a given combination of

- Road type (Motorway, trunk A road, principal A road, B road, minor road)
- Area type (10 area types, similar to those in the demand model)
- Sub-region (GovtOffice regions divided into busy & less-busy county groups)
- Time period (weekdays, Saturdays, Sundays, each split appropriately)
- Vehicle type (with cars further split by trip purpose)
- Direction (busier and less-busy, to allow for peak period tidality)
- Busyness (measured as volume:capacity ratio).

For each forecast run, the inputs to the model are:

- Traffic growth (by road type, area type, sub-region, for each vehicle type) between base and future year
- Any changes to the road network resulting from policy action, either in terms of additional infrastructure (represented as an increase in capacity) or tolling (as a cost per km, which may vary with banded congestion levels).
- Vehicle data giving future year fuel cost per kilometre and emissions per kilometre as a function of speed.

The program applies the input traffic growth to the base year traffic database, to get a future year "demand" traffic volume in each cell. Speed-flow relationships are used to calculate the revised speed in each time period. From the speed, the fuel cost is calculated, added to any toll charges, and converted to minutes using the appropriate future year value of time (separate for each vehicle type). This future year generalised travel time is compared with the base year value for each cell, and a set of rule-based responses applied. These shift traffic between cells, using elasticity values to determine the proportion of traffic which shifts, and the relative change in generalised time in the possible destination cells to determine where it shifts to.

The allowed responses are

- Reassignment to another, less costly (in time), road of the same type in the same area type and sub-region.
- Reassignment to another road of a lower road type in the same area type and sub-region (so motorway traffic can trickle down onto A-roads, and A-road traffic onto minor roads, but not vice versa)
- Retiming of traffic from the weekday peak hours to the adjacent time periods
- Reallocation of traffic across time periods, to reflect the extent to which responses modelled elsewhere have differential impact across different times of day.

Having shifted traffic, the program recalculates speeds and generalised times, and repeats the process until it converges to an equilibrium.

The outputs, which are all disaggregated by road type, area type and time of day include:

- traffic (further disaggregated by vehicle type or car purpose)
- total delay relative to free-flow speed (used for calculating congestion)
- Traffic times, costs and any toll charges,
- total tail-pipe emissions of the three major pollutants - CO_2, NO_x, PM_{10} - using emission equations as a function of speed at the detailed level

Car **link** speeds and any tolls are multiplied back by the mileage profiles (see chapter 4) to give a change in car driver **journey** speeds and money costs for each distance band and O-D area types. These revised generalised costs of car travel are fed back into the demand model for the next iteration.

Once the model has converged, the final traffic, journey time, cost and emissions results are used for detailed analysis of policy impacts on the highway network.

For more detailed information, see

Research report #8 - FORGE

Chapter 7 - the national rail model

In parallel with the highway modelling, the rail trips from the Demand model are sent to the National Rail Model (NRM). This assigns a CAPRI-based estimate of rail trips to a detailed geographic network of rail services, which can be varied to represent different rail investment scenarios. Resulting journey time, fare and crowding costs data are sent back to the Demand model, and assignment results are used for detailed analysis of policy impacts on rail.

Rail network

Whereas costs of travel by bus and slow modes can be described fairly naturally in terms of generic area types, the same is not true of rail. Similar towns can have significantly different levels of rail service, and a policy of rail investment tends to benefit particular services rather than all rail travel across the board. For this reason, the NRM uses a rail network based on real geography.

The network represents all passenger train services in GB, which connect a system of 1318 zones, each of which corresponds to one of the area types in the demand model. The capacities of the trains comprising the services are an important input into the assignment procedures. Each service is coded into the model with a particular rolling stock type so that the individual service capacities can be modelled accurately when calculating overcrowding costs.

The model includes a representation of all National Rail services, plus London Underground and Docklands Light Railway. The modelled periods are 0700 - 0959, representing the morning peak, and 1000 - 1600, representing the inter-peak.

The assignment procedure reflects the effect of service frequency, journey time, interchange, access/egress time and overcrowding on passengers' choice of route. In the peak periods, particularly in south-eastern England, capacity limitations have an important impact on travel behaviour. The NRM uses a standard approach for producing converged public transport assignments, undertaking incremental assignments of rail demands and revising rail costs based on the demand/supply ratio. Three trip purposes - commuting, business and leisure - are distinguished.

The NRM outputs components of the cost of travel which are used in further iterations of the demand model, which in turn enable the distributional and modal transfer effects of rail schemes to be determined. The assignment results are also used to provide detailed analysis of impacts of policies on rail demand, overcrowding and emissions.

Rail matrix

Unlike the highway side, for rail we have an estimate of the base year trip pattern. The base year rail zone-to-zone trip matrix is originally derived from CAPRI ticket sales data.

For each future year, this is factored up using rail industry-standard evidence on how rail demand grows over time, incorporating both GDP and time trend effects, to generate a reference trip matrix. These factors vary by 3 trip purposes and 4 groups of rail services. This is an interim measure, pending investigation of the different estimates of rail patronage growth over time obtained from econometric and multi-modal analysis.

Each run of the NTM demand model produces a matrix of rail trips at the level of area types and distance bands, responsive to cost changes on all modes. In each iteration, the current demand matrix is compared with a reference case demand matrix, and the ratio of the two applied to the reference case zone-to-zone trip matrix.

However, the 15x15 x13 (area type and distance band) cells in the demand model are at a much coarser level than the 1318x1318 cells in the detailed rail matrix. Some of the rail travel in each coarse cell will have been attracted by improvements to specific rail services in only some of the corresponding detailed matrix cells. To avoid the impact of any improvement being dispersed over all other rail services which happen to share the same area type / distance band combination but might be in a different rail corridor, a two-step process is used:

An initial estimate of the demand response at the detailed zone-to-zone level is obtained by applying standard elasticities to the changes in the NRM zone-to-zone rail costs. Then these initial estimates are controlled to the results from the demand model at the level of area type and distance band. This means that increases in rail patronage due to service improvements are allocated to the right services, while increases due to switching from other modes apply equally to improved and unimproved services.

Rail policies

The assignment to a rail network allows the representation of rail investment policies in detail, either as generic changes in service frequency, capacity or speed, or as packages of changes to individual services, representing particular enhancement schemes. Changes to fares can also be represented at the level of broad route corridors.

For more detailed information, see

Research report #7 - National Rail Model

Chapter 8 - Local Policy Interface - EPD and SID

The EPD (Expenditure Policy Database) is simply a database of Local Transport Plan (LTP) expenditures, as agreed in the most recent LTP settlement, categorised into different policy types. This is then projected forward to cover the full ten year period, so as to estimate what total Ten Year Plan expenditure is likely to be in terms of spending on each type of policy within each area type.

The more challenging component is SID (Spend Impact Database) which translates this expenditure into changes in the generalised cost of travel for each area type / mode combination, for input into the Pass 1 mode split/distribution model.

Where there is experience of policy impact from local modelling work or case studies, the modelled effects are based on these. Where there is to date little quantified evidence of the impact of a particular policy, the modelled effect is based largely on judgement.

The impact of policies can be represented in a number of forms, either singly or in combination:

- A factor applied to individual cost components (eg. Bus fares reduced by 10%)
- Mode constants adjusted by an amount that is calibrated to give a given modal switch in the base year
- A proportion of travellers affected by the policy (e.g. x% of trips to a given area type assumed to pay an additional cordon charge)

SID has changed since its use in the original 10 Year Plan modelling work. Then it produced traffic change factors (TCF) which were applied to the background growth forecasts input to FORGE. Within the NTM framework it does not need to go this far - the cost changes are input into the Demand model, leaving that model to calculate the impact on traffic.

For infrastructure improvement policies which have a capital cost, and fare support policies which have a revenue cost, it is relatively straightforward to convert expenditure into some physical measure of the amount of policy applied, and then estimate the impact of this on travel costs.

For charging policies which in general will not only cover their operating costs but actually raise revenue, it is assumed that the money raised becomes revenue expenditure and is re-invested in public transport. These policies are described in terms of the number of schemes rather than the money cost.

Less-quantifiable policies (such as green travel plans), are assumed to be capable of application at three levels - high, medium and low (or not at all). Rather than affecting the money or time cost elements within the demand model, expert opinion has been sought as to the likely scale of impact, and mode constant terms within the model are adjusted by an amount that is calibrated to give this effect.

SID is a convenient way of representing local policies in an aggregate, largely non-geographic, national model. But there is inevitably some debate as to how representative are the policy impacts reported by case studies, and how they should be scaled up or down when applied to different sizes of urban area. This is therefore an area of the modelling work which we will wish to review regularly to ensure that latest information is incorporated.

Different towns within different local authorities, but with the same area type, will in practice apply different policies. But SID and the demand model can only deal in averages. Arguably, the model is representing some fraction of a light rail line being applied to every urban area of a given area type. We have undertaken tests to check that linearity holds - e.g. that an average of one-tenth of a light rail scheme in each large urban area gives roughly one tenth of the impact of one light rail scheme in each area - and this proved to be an acceptably good modelling approximation.

For more detailed information, see

Research report #3 - Spend-Impact Database

Chapter 9 - Disaggregation

Different components of the model operate at different levels of geographical and market detail, so as to better focus on the key drivers of that particular process. This chapter sets out the degree of segmentation of the major datasets involved.

Trip ends

Within the trip end model, trip end figures are calculated at a very high level of disaggregation:

15 purposes x 88 person types/household types x approx. 10,000 wards.

These trip ends are aggregated differently at different stages in the modelling.

Demand model

Pass1 reduces the 15 x 88 = 1320 market segments to 105, using additional data to proportion trips by socio-economic group for some trip purposes. The table below shows the definition of these 105 segments, which were derived from analysis of NTS data to see for which groups there is statistical evidence of significantly different behaviour.

The geography is compressed to 15 area type / region combinations.

The demand model therefore reads in trip production totals for each of 15 x 105 = 1575 combinations. The trip attractions are split only by 15 area types and 8 purposes, and act as a constraint on the Demand model results, ensuring that each area type is allocated its appropriate share of trip destinations.

The structure of the model allows the costs that are read into the demand model to differ for each combination of

- 15 production area types
- 15 attraction/destination area types
- 13 distance bands
- 105 market segments

but in practice there is much less variation involved - most of the costs are represented with considerably fewer than the theoretical maximum of over 300,000 levels.

Bus fares are the same for all market segments except that children and pensioners (where identified separately) get concessionary fares. Bus times and fares vary by distance band, according to one of seven different tariffs, depending on the destination area type.

Purpose	**Person type**	**SEG / Income**	**1 adult / 0 car**	**1 adult /1+ car**	**2+ ad / 0 car**	**2+ ad / 1 car**	**2+ad / 2+ car**	**All**
HB Work	Full time empl	High Medium Low	1 6 11	2 7 12	3 8 13	4 9 14	5 10 15	
	Rest of pop'n	All	16	17	18	19	20	
HB EB	Full time empl	High Medium Low	21 26 31	22 27 32	23 28 33	24 29 34	25 30 35	
	Rest of pop'n	All	36	37	38	39	40	

HB Educ	Child (0-15) Full time empl Other 16-64 Pensioner		41 46 51 56	42 47 52 57	43 48 53 58	44 49 54 59	45 50 55 60	
HB PB / Shopping	Child (0-15) Full time empl Other 16-64 Pensioner		61 66 71 76	62 67 72 77	63 68 73 78	64 69 74 79	65 60 75 80	
HB Rec / Visiting friends	Child (0-15) Full time emp Other 16-64 Pensioner		81 86 91 96	82 87 92 97	83 88 93 98	84 89 94 99	85 90 95 100	
HB Hols / Day trips	All persons							101
NHB EB	All persons	High Medium Low						102 103 104
NHBO	All persons							105

The purpose split is based on the NTS categorisation of trip purposes:

- Home-based (HB) trips are all those starting or ending at home.
- HB Work trips are those between home and the person's usual place of work - trips in the course of work count as Employers Business (EB).
- Personal Business (visiting bank, doctor etc) trips are combined with shopping trips, as these have similar characteristics.

Non-home based trips are divided only into Employers Business and Other.

There are three different rail tariffs, representing season tickets (for commuting and education trips), full-fare tickets (for business trips) and saver fares (for all other trip purposes). Children and pensioners (where identified separately) get concessionary fares. Commuting, education and home-based employers business trips get peak period rail journey times and service levels; for other purposes off-peak characteristics are used. In each case, the National Rail Model calculates figures for every valid combination of area type and distance band.

Walk and cycle speeds take only two different values, being slightly slower in the area types representing the biggest urban areas.

There are two sets of car journey times, with different purposes taking peak and off-peak values as for rail. Times vary for each area type/distance band combination, and vehicle operating costs are calculated from average speeds.

Highway distribution & assignment

The detailed highway model used to convert trips into traffic growth for input to FORGE has nearly 10,000 wards, but combines all purposes and market segments.

FORGE

FORGE itself reads in for each of the six car purposes and four other vehicle types a "demand" traffic growth for each combination of 7 road types, 10 area types, and 20 sub-regions. Not every sub-region contains every area type, but this still gives over 500 levels of traffic growth for each purpose or vehicle type.

National Rail Model

The NRM receives rail trips in the format used by the Demand model, but needs to aggregate to three market segments, and disaggregate to 1318 geographic zones in order to assign to its detailed geographic rail network. After assignment, the journey time components and fares data are aggregated, and market segments are disaggregated, back into Demand model format.

Chapter 10 - Model calibration & runs

This chapter looks briefly at how the different components are calibrated, and then how they work together to run the model in a typical forecasting exercise. In the current version, the base year for the NRM and demand model is 1998, while the base year for the highway assignment and FORGE is 2000.

Calibration

The calibration of the demand model aims to reproduce:

- mode shares and trip length distributions for each market segment, as recorded in the 1989-1996 NTS data
- target response elasticities of car traffic to fuel cost and rail passenger-km to rail fares
- mode shares for journey to work to conurbations and London, from the Labour Force Survey (because NTS does not have data on destination area types).

The parameters available to achieve this are

- sensitivity parameters at each level of the choice hierarchy, some of which vary by trip purpose
- mode-specific constants for each mode and distance band, varying by market segment
- further mode constants that vary by destination area type.

The calibration is not straightforward, and adopts a heuristic approach. Partly this is due to the fact that the model is doubly-constrained at the middle level of the choice hierarchy, so as to match attraction totals by trip purpose and area type. Partly it is because the aggregate elasticities that are to be matched apply across all trip purposes and market segments combined, so that it is not possible to look at these individually and tailor the model to individual markets.

The highway assignment model is calibrated, varying the value of time used for assignment so as to reproduce car traffic statistics on the proportion of total traffic using each road type.

The rail assignment model takes an overcrowding function, and values of service headway and access/egress time relative to journey time, from other studies. The assignment has been validated, comparing the results to SRA data for passenger counts on cordons and estimates of passenger-miles by train operator.

FORGE is not calibrated as such. The elasticities have been derived from the same input elasticity of car traffic to fuel costs as is used for Pass1 calibration; the relative strengths of the different responses in FORGE are taken from stated preference work.

The freight model uses generalised cost parameters (one for value of time and one for value of reliability) from stated preference work. There are two dimensions where calibration is used: cross channel services, where we have precise link flows, and rail traffic by commodity. The model sets impedances on the two sets of links to reproduce these. There is no calibration for road traffic.

See the various research reports for more detail on calibration.

Setting up the model

The first step is to compile the planning data inputs for all base and future years, and to run the car ownership and trip end models to give a database of numbers of trips made by each category of person. Trip attractions are used to weight each destination zone.

The highway network has been updated to 2000; a unit matrix is assigned to this to calculate the "size terms" - an index of how many destinations of each area type lie within each distance of each other area type.

The NRM is run, assigning a base year rail matrix derived from ticket sales data, to give base year rail costs. Base year costs for other modes are derived from the National Travel Survey.

The demand model is then run once to forecast trips for 2000, using NRM results for rail costs and external estimates of change in highway costs between 1998 and 2000. The resulting car driver trips are fed through the distribution model and highway assignment model, to give an estimate of 2000 car traffic. This is factored to match published 2000 traffic totals by road type, to allow for car traffic other than GB residents in households.

Running the FORGE module with zero growth gives base year traffic, emissions and congestion results.

Reference case (do-Minimum) forecasts

For each future year of interest, the model is first run to equilibrium with no policy inputs, or only those policies that were in place before the announcement of the Ten-Year Plan.

This starts with the demand model. Trend assumptions are used for rail and bus fares, and initially, 2000 highway costs and rail service levels.

On this first iteration, the demand model is followed by the highway distribution and assignment models. This picks up any change in car routeing patterns caused by land use change, trip lengthening over time, or fuel price change between 2000 and the forecast year. These routeing patterns are now fixed, and used in subsequent runs for the same forecast year.

The freight model and LGV growth models are run, followed by FORGE, and the increased highway costs caused by congestion are calculated and fed back to the demand model.

In parallel with FORGE, the NRM is run. This has its own procedure for projecting rail growth over time, depending on GDP growth and time trends. Assigning this higher rail demand to the 2000 rail network gives a first estimate of future year do-Minimum rail costs.

On the second and subsequent iterations,

- the demand model reads in higher rail overcrowding and highway congestion numbers, and estimates a revised mode split.
- the car driver trips are converted directly into traffic using the same routeing patterns as previously
- the same LGV and HGV demand growth is fed into FORGE
- FORGE re-estimates the future year congestion costs
- In parallel, NRM estimates demand by taking the difference between first and current runs of the demand model, & applying that to its future year base demand.

The process usually converges acceptably within 4 or 5 iterations. Convergence is monitored using both aggregate statistics for modal share, and disaggregate statistics which calculate a root-mean-square average change over all cells in the demand model results between successive iterations.

Policy tests

Having established a reference case scenario, the model is ready to examine how this future year picture of total travel might be altered by the application of different transport policies. A typical list of policies to examine might be as follows:

Policy	**NTM Representation**
National road policy	
Trunk road improvements - widenings	Additional lanes on busiest or road links in FORGE, reducing journey times for any given flow
Trunk road improvements - bypasses	Change to average capacity of trunk roads in relevant area types in FORGE
Local road improvements Trunk road improvements - junctions	Changes to the average speed-flow relationships for the relevant road types in FORGE
Inter-urban road charging	Additional money cost per km on relevant road types in FORGE; extra costs feed back to demand model
Fuel price, measures to increase fuel efficiency	Changes in money cost for car driver and car passenger trips in the demand model and emissions calculations in FORGE.
Local transport policy	
Revenue support for bus and local rail	Reduced bus fares and rail fares for short journeys, reducing money costs in the demand model
Capital investment for bus and light rail	Reduced bus and rail access times and interchange penalties in the demand model
Raising parking charges Road user (cordon) charges	Higher average money cost on some or all car trips destinating in the relevant area types in the demand model
Changes to parking supply	Changes to parking search time - an element of gen.cost for car modes
Bus quality partnerships	Reduce the mode constant (disutility of bus travel) for trips destinating in the relevant area type, in the demand model
Initiatives to encourage walking and cycling	Vary the mode constants to reproduce expected mode shifts
Workplace parking levy	Equivalent to parking cost, but for commuting purposes only
School/workplace transport plans	Increase the mode constant (disutility of car driver trips) for education/commuting purposes
Land-use planning policies	
Concentrating growth in urban areas	Allocation of growth between urban and rural parts of local authority districts in the trip-end model
Mixed land use, higher densities, avoiding edge-of-town development	Reduction of average car trip lengths for <10 mile journeys to the same area type as the

	origin zone. Reduced PT access/egress times.
National passenger rail policy	
Fare changes by ticket type	Money cost changes by journey purpose in the demand model
Investment in infrastructure	Higher rail speeds and frequencies on affected services in the NRM
Longer trains, rolling stock replacement	Increased capacity leading to less overcrowding for a given load, in the NRM
Freight policy	
Rail freight improvements	Money or time costs of rail mode within the freight model
"Sustainable distribution"	Increased average load factors for lorries in the freight model, & improved fuel efficiency
Lorry VED, permitted lorry weights	Money costs of different vehicle types in the freight model

For more detailed information, see

Working paper 2 - model inputs

Welfare Module

Having obtained results, it remains to estimate the associated economic welfare and environmental impacts.

The economic welfare calculation uses the Department's standard TUBA package. This is run using the demand model outputs in terms of person-trips and their associated costs, so as to estimate the transport user benefits. This has to be supplemented using outputs from FORGE, to pick up both the welfare impacts on freight traffic, and the externalities (emissions etc).

For more detailed information, see

Working paper 4 - welfare module