ALLOWING FOR VARIABLE DEMAND IN HIGHWAY SCHEME ASSESSMENT

J J Bates
John Bates Services
D Coombe
MVA Ltd
S Porter
Department of the Environment, Transport and the Regions
D van Vliet
Institute for Transport Studies, Leeds University

1. INTRODUCTION

In their seminal report (SACTRA, 1994) the UK Standing Advisory Committee on Trunk Road Assessment concluded that "induced traffic" was likely to occur as a result of road improvements, and recommended that:
"*...variable trip matrix economic evaluations are undertaken for schemes as the cornerstone of the appraisal in every case, except where it can be shown that the trip matrix will not vary as a result of the scheme being appraised.*" (para 13.49).

While the main thrust of the Committee's recommendations was that, wherever feasible and cost-effective, the components of suppressed and induced traffic should be modelled individually, using some form of multi-modal transport model, it recognised that there are circumstances when simpler techniques would be appropriate (in small and medium-sized urban areas, for example), and that in the short term, simple elasticity models might be used (paragraphs 14.11 and 14.14).

In response to SACTRA's Report, the (then) Department of Transport has issued two versions of a Guidance Note: the second version (February 1997) is part of the Design Manual for Roads and Bridges (DMRB, Volume 12.2.2), and contains advice on the form and application of elasticity models.

Towards the end of 1997, the Department of the Environment, Transport and the Regions (DETR) commissioned a study on Improved Elasticities and Methods, in order to be sure whether, and if so how, elasticity values and methods should be further refined. It quickly became clear that it was desirable to move the emphasis of the research away from elasticity models and to tackle more general problems currently inhibiting other forms of Variable Demand modelling.

The study aims to set out the theoretical basis for the future development of Variable Demand methods. Since this has implications for the way that data is collected and how inputs are prepared, some aspects may have to be regarded as long term aspirations, rather than matters which can and should be changed immediately.

The main part of this Paper is concerned with an exploration of **assignment-based variable demand modelling**. In addition, it sets out an approach to **forecasting the**

reference case and hence establishing a realistic do-minimum case. Given the restrictions of length, detailed mathematical proofs are **not** presented here. The later sections describe the testing process with which the project is currently occupied.

2. ASSIGNMENT-BASED VARIABLE DEMAND MODELLING

The focus on 'variable demand modelling' requires an understanding of the basic principles of transport economics – in particular, the terms 'supply' and 'demand' and the related concept of an equilibrium system – at the outset.

In classical economics both supply and demand are treated as functions of cost, but the normal graph is 'inverted' by plotting cost on the vertical axis, as in Figure 1. The notion that travel demand T is a function of cost C presents no fundamental difficulties: the term 'demand model' implies a procedure for predicting what travel decisions people would wish to make, **given** the generalised cost of all alternatives. These decisions include choice of time of travel, route, mode, destination, frequency/trip suppression.

However, if these predicted travel decisions were actually realised, the generalised cost might not stay constant. This is where the 'supply' model comes in, to reflect how the transport system responds to a given level of demand: in particular, what would the generalised cost be if the estimated demand were 'loaded' on to the system. The best known 'supply' model is the conventional traffic assignment reflecting *inter alia* the deterioration in highway speeds as traffic volumes rise, but there are other important effects, such as the effects of congestion on bus operation, overcrowding on rail modes, and increased parking problems as demand approaches capacity.

Since both demand and supply curves relate volume of travel with generalised cost, the actual volume of travel must be where the two curves cross, as in Figure 1 - this is known as the 'equilibrium' point. An 'equilibrium model' ensures that the demand for travel is consistent with the network performance and other supply effects in servicing that demand. If the final estimate of demand is loaded on to the supply, the resulting costs should exactly generate the estimate of demand which is loaded.

Although the term demand is often used as if it related to a quantity which is known in its own right, it must be emphasised that the notion of travel demand **always** requires an implicit or explicit assumption about costs. The actual forecast resulting from a strategy or plan (sometimes misleadingly called **expected demand**) is the outcome of the equilibrium process referred to above.

Of course the level of demand and hence the forecasts will reflect the demographic composition of the population, together with other external changes (eg effects due to land-use, income, car ownership etc.). However, in most cases, when assessing the impact of a policy, which means essentially changing the **supply** curve, the demand curve is held constant. Hence, the appraisal of proposed changes to the transport system can usually be viewed as a comparison of two (or more) equilibrium points, using a common demand curve, but with each equilibrium point associated with a different supply curve, as shown in Figure 2.

When we consider the problem in a general transport context, we are not interested in single elements of demand, but rather a matrix of elements. Correspondingly, there is no single cost associated with the supply side, but an array of individual elements.

Additionally, the transport problem is complicated by the supply domain being that of a network of **links**, while the demand for travel relates to the inherent value of being at j, given a current location at i, and not to the particular paths used to reach j. Thus the domain of the demand model is essentially the **i-j pair**. The result is that an interface procedure is required. We return to this below.

We can generalise the simple one-dimensional example presented above, replacing the quantities T and C by multi-dimensional **matrices**, subject to some qualifications which we will ignore in this paper. Hence, the conditions for equilibrium are that: **given** $\mathbf{T}_d = f_D(\mathbf{C})$, $\mathbf{C}_s = f_S(\mathbf{T})$, we require

$$\mathbf{T}^* = f_D(\mathbf{C}^*) \quad : \quad \mathbf{C}^* = f_S(\mathbf{T}^*) \qquad (1)$$

The importance of the need to find the points of equilibrium with some accuracy cannot be ignored. The demand/supply diagrams shown in Figures 1 and 2 have been drawn with false origins for the sake of clarity. In reality, however, the benefits associated with a change in the supply curve are a very small quantity derived as the difference between two large quantities which have a certain degree of error associated with them. In order to derive the benefits accurately, it is essential that the equilibrium points are found accurately for both the do-minimum and do-something cases. Failure to do so (or, more fundamentally, to adopt modelling procedures which enable equilibrium to be found with accuracy) could easily result in erroneous decisions being taken. This point has been largely overlooked in the past.

Except for the very simplest models, there are no direct ways of calculating the equilibrium solution, and it is necessary to set up iterative procedures. Although a well-conceived iterative system should converge to a unique solution, most methods will produce only approximate equilibria, both because of inherent computational inaccuracy (eg rounding) and the desire to limit computing time.

Many complicated systems of equations can be solved by formulating an optimisation problem which by design is guaranteed to have the same solution. This is an extremely useful result, because considerable effort has been put into constrained optimisation techniques which have quite general application. Over 40 years ago an appropriate 'objective function' (to be minimised) was demonstrated by Beckmann *et al* (1956) to yield a solution consistent with Wardrop equilibrium, though it was many years before this was routinely implemented.

We may note that the equilibrium point (T*,C*) in Figure 1 can be obtained by maximising the area between the demand and the supply curves. Expressing it for convenience as a minimisation problem in T by taking the negative, we have

$$\min z(T) = \int_0^T [f_S(W) - f_D^{-1}(W)]\, dW \qquad (2)$$

where $f_D^{-1}(T)$ is the **inverse** demand curve, representing the price C at which the demand would reach level T.

Subject to certain mathematical conditions being satisfied, this one-dimensional approach can be extended to the case where the elements T, f and W are **matrices.** Our attention therefore focuses on how the demand and supply integrals can be evaluated. While the emphasis of the project is more involved with the demand side, there are some complications relating to the supply function which need to be discussed.

3. THE SUPPLY CURVE FOR HIGHWAY TRAVEL

As drawn, and implied in the matrix-based extension of the equilibrium condition, the supply function delivers the matrix of O-D costs associated with a demand matrix **T**. However, as noted, the supply mechanism operates on a link rather than O-D basis. For each element of demand there will be a set of chosen routes between i and j. In the simple "all or nothing" case, we can define an "origin-destination-link incidence matrix"

ε_{ija} = 1 if the journey from i to j uses link a, 0 otherwise.

When dealing with congested assignment, where normally we have multiple routes for a single O-D pair, we can treat ε as giving the **proportion** of i-j movements using link a.

We now sketch out the nature of the interface between Demand and Supply. The process conventionally referred to as "Assignment" subsumes a number of "modules":

- definition of the set of paths for each origin-destination movement
- estimating the choice between those paths, thereby yielding ε.
- calculating loads on links (defined by the matrix product **V**=**T**.ε).
- capacity restraint - adjusting link speed in response to changes in loads. This is the **true** supply process, giving link costs $c_a = f(\mathbf{V})$
- skim - extract O-D costs (defined by the matrix product **C** = ε. **c**).

According to this description, the interface between O-D and link is achieved by the matrix ε. The "supply curve" is, in effect, the outcome of a series of assignments of different matrices **T**, each yielding a corresponding cost matrix **C**. Except in the simplest cases, there will be no "closed form" expression for the supply function (ie, it cannot be written in a way which can be directly evaluated)..

In spite of this, we can evaluate the supply integral by proceeding along the following lines. To avoid definitional problems, we confine ourselves to the "standard" case where link costs are independent (ie, where the costs on a link are unaffected by the flows on other links). By changing the variable of integration from **T** to **V**, where. **V**=**T**.ε, the supply integral $\int_0^T f_S(W)\,dW$ becomes $\int_0^V c_S(X)\,dX = \sum_a \int_0^{V_a} c_a(x)\;dx$ where the elements **V**, **c** and **X** are now link-based vectors. Where the link costs are not independent, it is still possible to proceed by the so-called "diagonalisation" method, by forming a local approximation to the link cost functions (Sheffi, 1985).

It will be seen that the supply integral is exactly what is required for the reasonably widespread application of Wardrop equilibrium assignment, using fixed demand. Hence, in spite of the intractability of writing out the supply functions in matrix terms, it is possible to calculate the supply integral on a link basis, in a way which corresponds exactly to procedures currently used in fixed matrix assignment.

4. THE DEMAND INTEGRAL

In the fixed demand problem, where the aim is to estimate the link flows, the demand integral is constant and can therefore be dropped from the objective function. In contrast, the variable demand problem requires both the (highway) flows and the level of demand to be determined.

The specification of the objective function in terms of the inverse demand function has naturally led to the question of how these inverses might be defined, and a number of developments have flowed from this. The number of software applications for solving Variable Demand Assignment (VDA) problems is very limited, and most software is limited to demand functions having the specific property known as **separability** - that the highway demand for a particular ij movement does not depend on the highway costs for any other ij pair. This is compatible with the "simple elasticity functions" suggested by SACTRA and recommended in the current Guidance. An exception is the SATAST program (Arezki *et al*, 1996) developed under contract to the Highways Agency.

The separability restriction on the demand functions, which is mathematically identical to that on the link cost functions, allows us to follow the original specification by Beckmann *et al.*

$$z(\mathbf{V},\mathbf{T}) = \sum_a \int_0^{V_a} c_a(u)\, du - \sum_{ij} \int_0^{T_{ij}} f_{ij}^{-1}(v)\, dv \qquad (3)$$

With the reasonable proviso that demand should be a non-increasing function of cost, separability guarantees that the inverse demand functions exist, are unique and can be straightforwardly calculated, either directly by means of analytical formulae, or by simple numerical approaches.

In itself, separability does not prevent the demand being dependent on other costs (eg for other modes, or other time periods) **provided** these costs are not adjusted within the course of the program run. However, while it is possible to take into account a number of different demand responses while observing the restriction, the restriction is in principle severe, and more or less rules out any allowance for **redistribution**. For this reason, it was considered essential to understand the computational implications of attempting to remove it.

It is well known (Evans, 1976) that an appropriate objective function can be specified for one type of non-separable demand function, that associated with constrained distribution models. Further, Chapter 10 in Ortúzar & Willumsen (1994) gives a useful discussion of the way in which researchers have developed VDA approaches with more complicated demand functions, within the general sphere of random utility models, going beyond the Beckmann *et al* formulation, and a substantial generalisation is given in Oppenheim (1995). These extensions retain the essential convexity properties, so that the Frank-Wolfe algorithm is a feasible approach.

By far the most common model form used in transport planning is the logit. It turns out that the demand integral can be represented in closed form for **any** hierarchical demand model based on a logit formulation where the "utility" is linear in generalised cost. In order to sketch this out, we introduce the concept of consumer surplus (CS), defined as the difference between the area under the demand curve and the total consumption, which we write as **T.C**. In other words,

$$CS = \int_0^T f_D^{-1}(W)\,dW - T.C \qquad (4)$$

Hence, the required demand integral $\oint_{\mathbf{T}} f_D^{-1}(\mathbf{W})\, d\mathbf{W} = \mathbf{T.C} + CS$ (5)

The reason for proceeding in this way is that there is a closed form expression for CS. It can be shown that the so-called composite cost C_*, derived using the well-known "logsum" formula, is equal to the **negative** of the consumer surplus per traveller, up to an arbitrary constant. While this is easy to show for a single-level logit model, it in fact applies to the hierarchical logit model as well: the overall CS is given by the negative of the composite cost at the top of the tree.

If **I** is a unit vector with the same dimension as **C**, then $\mathbf{T.I} = T_{*.}$, the total demand over all travel possibilities. Hence, for any logit model, the demand integral can be represented, up to an arbitrary constant, as $\mathbf{T}.[\mathbf{C} - C_*.\mathbf{I}]$ where C_* is a scalar quantity. However, to use this result in the objective function, we have to develop the formula as a function of **T**. The single level logit formula can be written in vector form as: $\ln \mathbf{p} = -\beta\,(\mathbf{C} - C_*.\mathbf{I})$, whence $\mathbf{C} - C_*.\mathbf{I} = -1/\beta \ln \mathbf{p}$. We can therefore substitute in the demand integral equation, and obtain, up to an arbitrary constant

$$\oint_{\mathbf{T}} f_D^{-1}(\mathbf{W})\, d\mathbf{W} = \mathbf{T}.[\mathbf{C} - C_*.\mathbf{I}] = -\mathbf{T}.(1/\beta \ln \mathbf{p}) \qquad (6)$$

In this way, we have expressed the demand integral entirely in terms of **T**, the demand matrix, or **p**, which is equivalent to $1/T_*.\mathbf{T}$. With a hierarchical logit model, the formula is only marginally more complicated., so that for a 4-level logit model with time of day conditional on mode, conditional on destination, conditional on frequency (origin), we would obtain

$$C_{ijmt} - C_{****} = -1/\beta^T \ln (p_{t/ijm}) - 1/\beta^M \ln (p_{m/ij}) - 1/\beta^D \ln (p_{j/i}) - 1/\beta^F \ln (p_i) \quad (7)$$

where, with obvious notation, β^T is the scaling parameter for generalised cost at the time of day choice level, etc. Hence, $\oint_{\mathbf{T}} f_D^{-1}(\mathbf{W})\, d\mathbf{W} =$

$$-\sum_{ijmt} T_{ijmt} \cdot \left[\frac{1}{\beta^T}\ln(p_{t/ijm}) + \frac{1}{\beta^M}\ln(p_{m/ij}) + \frac{1}{\beta^D}\ln(p_{j/i}) + \frac{1}{\beta^F}\ln(p_i) \right] \qquad (8)$$

In applying these formulae, it is important to recognise the role of the logit model as a model of **shares**, given an overall constant total. To write out the formula for any chosen model, the analyst will need to identify at what level the constancy applies. In the illustration just given, it is assumed that the **total** travel T^{****} is constant.

Compatibility between Supply and Demand Costs

The derivation of the demand integral as the sum of CS and **T.C** needs to ensure that the costs being used for the calculation of **T.C** are compatible with those being used within the supply integral, if the objective function is to be properly defined. The closed form using the logit function relies implicitly on the fact that the costs C of individual alternatives ijmt contain **all** the required components – they are in fact pure logit 'utilities' scaled by $-\beta$ (at the lowest level). As such, the C_{ijmt} contain **all** components of utility (ie including, according to the scope of the model, modal constants, destination utilities etc.). Hence, some further adjustment is required within the objective function.

If we define the demand costs **C** as **C** = **N** + **R** (where **N** is network costs and **R** is 'residual' costs), then the implication is that the demand 'expenditure' term **T.C** should be replaced by **T.N**, to make it compatible with the supply costs. Thus we correct the demand integral formula by subtracting $\mathbf{T.R} = \Sigma_{ijmt} T_{ijmt}. R_{ijmt}$. This can be viewed **either** as bringing down the demand curve **or** shifting up the supply curve.

In cases where we **know** all the items in **R** this correction is straightforward. In other cases, it may be more tedious to calculate it. For example, if we had an incremental logit model with the same illustrative hierarchy, there is a base matrix $\mathbf{N}^0$ which, when combined with the unknown matrix **R**, generates the base demand $\mathbf{T}^0$. In this case it can be shown that, writing **q** for the base shares:

$$R_{ijmt} = C^0_{****} - 1/\beta^T \ln(q_{t/ijm}) - 1/\beta^M \ln(q_{m/ij}) - 1/\beta^D \ln(q_{j/i}) - 1/\beta^F \ln(q_i) - N^0_{ijmt} \quad (9)$$

where C^0_{****} is the base overall composite cost.

Alternatively, consider a singly constrained distribution model, with the choice of destination given by

$$p_{j\backslash i} = \frac{W_j.\exp(-\beta N_j)}{\sum_k W_k.\exp(-\beta N_{jk})} \quad (10)$$

where W_j is the "size" variable associated with destination zone j. In this case it is clear that R is given by the (implied) destination utilities, so that $R_{ij} = -1/\beta \ln W_j$.

Hence we have a theoretical basis for calculating the demand integral. In all cases, it can be seen as an expression of consumer surplus, **plus** a term relating to total consumption (in demand cost terms) **minus** a correction for "residual consumption", defined as the difference between consumption in demand and supply terms. In practical terms, this theoretical analysis means that it should be possible to extend conventional equilibrium assignment packages to enable them to handle some or all of the stages of conventional demand modelling.

5. FORECASTING FUTURE DEMAND

At any point in time, the demand curve $f_D(\mathbf{C})$ is **constant** for the scenario under consideration. However, over time, the population and land-use will vary, and this will lead to different demand curves, each related to a particular point in time. Where there are different views on how the future population and land-use will develop, different planning assumptions (termed 'scenarios') may be required for the same year. The demand model therefore needs an interface with external 'planning' or 'land-use' data (in particular, forecasts of car ownership) to reflect how the scenario assumptions affect total travel demand.

We can represent the exogenous influences at any point in time by the "**planning data**". The term "planning data" is here being used as a shorthand to include the effects of income, car ownership and licence-holding, and thus a somewhat wider definition than in normal usage.

If we assume that the **form** of the demand curve (which governs the response to changes in cost) is unchanged, it follows that the shifting of the curve over time is entirely due to changes in planning data and, to a reasonable approximation, we may assume that for that scenario the planning data changes are independent of transport costs.

Hence, as an input to a Variable Demand Model, we can predict the important effects of changes in planning data for each scenario in the absence of changes in (transport) cost. Typically, we already know a point ($\mathbf{T}^P$,$\mathbf{C}^P$) on the **base year** demand curve $f_D{}^P(\mathbf{C})$, which we may take as an equilibrium point (so that it also lies on the base year supply curve). We can then develop a methodology to obtain growth factors γ, on the assumption that costs remain at their base value of $\mathbf{C}^P$. Note that it is not only the **highway** costs which must be held constant, but the costs for all transport alternatives.

The 'product' of these factors γ and the base travel $\mathbf{T}^P$ gives an estimate of the demand $\mathbf{T}^0$ which would occur under the future scenario if there was no change in transport costs. Hence what we refer to as the '**reference**' point, ($\mathbf{T}^0$,$\mathbf{C}^P$), lies on the future demand curve, and we can use a formulation with this point as a 'pivot'.

Since there **will** almost always be a change in transport costs (not least because of the very 'exogenous' growth being modelled), the demand level associated with the Reference point is a hypothetical concept – it does not represent a realistic forecast of what will happen. That realistic forecast will have to be found by solving the equilibrium between supply and demand, using the principles set out above.

Note that the supply curve will reflect not only any changes in the network, but also (exogenous) background changes in generalised cost, including, on the one hand, fuel prices (and, in a wider context, public transport fares), as well as, on the other, changes in values of time. Hence the supply curve is time dependent as well. **In principle**, however, it is independent of the specification of the demand curve.

The result is that the costs **C** in the demand function are expressed in terms of future conditions for year y. **Nonetheless**, when using a 'pivot' formulation, they need to be compared with the base year costs $\mathbf{C}^P$. This has some potentially important implications for the definition of **units** for the parameters in the demand model. Essentially, since we may view the demand model as a random utility model, we need to ensure that we have a constant definition of **utility**.

The standard assumption is to define **C** in minutes, and to allow money costs to be deflated by the value of time. It is then assumed that the demand function parameters remain constant (in time terms): this would appear to be the most **neutral** assumption, though empirical validation would be very difficult to achieve.

Our preferred approach is thus: a) calculate growth rates γ to obtain the future year reference matrix from a base year pattern $\mathbf{T}^P$ on the basis of zero cost changes for all transport alternatives; b) hold the form of the chosen demand curve f(**C**) constant in time units; c) set the actual demand curve for the future year with reference to the initial condition ($\gamma.\mathbf{T}^P,\mathbf{C}^P$), where the costs $\mathbf{C}^P$ are in equilibrium with $\mathbf{T}^P$ in the base year; and d) ensure that the costs **C** used for the future year demand curve incorporate changes in the value of time, as well as any changes associated with fuel prices etc.

6. DEVELOPMENT OF ALGORITHMS

The standard way to solve VDA problems, given an objective function, is to use the modified Frank-Wolfe procedure suggested by Evans (1976). This is incorporated within the existing SATURN suite of programs for a limited range of separable demand functions. What follows is a very brief description of the main working of the algorithm, and does not go into the detail of all possible options. In particular no account is taken of the wider context of SATURN which allows for simulated junction cost functions etc. We also ignore questions of initial values, though this is a fruitful area of investigation.

At iteration n we have an estimate of the current matrix $\mathbf{T}^{(n)}$ and flows $\mathbf{V}^{(n)}$. The flows determine the link costs, and from these we can find the minimum paths. We require the auxiliary solutions $\mathbf{S}^{(n)}$ (matrix) and $\mathbf{W}^{(n)}$ (flows). We begin by calculating the auxiliary matrix $\mathbf{S}^{(n)}$, using the current minimum cost matrix $\mathbf{C}^{(n)}$ (dependent on $\mathbf{V}^{(n)}$) and making a **direct** estimate of the implied demand, using the selected demand function. The auxiliary flow vector $\mathbf{W}^{(n)}$ is then obtained by assigning the auxiliary matrix $\mathbf{S}^{(n)}$ on an all-or-nothing basis to the minimum cost paths.

Given the auxiliary solutions **S** and **W**, we carry out the line search (step size) with the aim of minimising the objective function, based on $(\lambda^{(n)}.\mathbf{W}^{(n)} + [1-\lambda^{(n)}].\ \mathbf{V}^{(n)}, \lambda^{(n)}.\mathbf{S}^{(n)} + [1-\lambda^{(n)}].\ \mathbf{T}^{(n)})$, with respect to $\lambda^{(n)}$. The optimum value for λ is then used to provide the current estimates for the next iteration.

The implications are that provided there is no major complexity or computational cost in evaluating the objective function (or its differential, for the step length calculation), the existing SATURN algorithm can be directly extended to more complex demand problems, by substituting the new demand function for the current SATURN formulations in the calculation of the auxiliary $\mathbf{S}^{(n)}$, and modifying the objective function accordingly. This was originally tested for the "distribution and assignment" problem, and has since been tested for hierarchical logit specifications including distribution, mode choice and time of day choice.

Note that the step length λ is doing two rather distinct things: it is "averaging" between successive estimates of the demand matrix, **and** it is shifting traffic between existing paths and new paths for the equilibrium assignment. Although performance has been good on the relatively simple demand functions tested to date, there are theoretical concerns that this "dual rôle" may not be appropriate in all cases.

For equilibrium assignment in congested conditions, it will be rare for any single new path to take a substantial proportion of the traffic: from this point of view, λ will tend to fall quite quickly as the number of iterations increases. If the demand matrix does not stabilise quickly, this tendency to low values could cause problems in convergence. For this reason, a bi-level approach is also being considered in which separate line searches are carried out, first for the equilibrium assignment for a current fixed demand, using the supply integral only, and then using the full objective function to adjust the demand matrix. Effectively, this can be seen as an alternative definition of the auxiliary **W** in the Evans algorithm.

7. TESTING PROGRAMME

Based on the theoretical review summarised in this paper, an ambitious programme of tests has been designed, with a number of aims in mind. A critical emphasis has been on the development of **practical** approaches: for this reason, the main test bed is a sizeable network based on an existing model for Glasgow, with both highway and public transport information, for three separate time periods.

It remains a central task to provide some guidance as to the level of complexity appropriate in different modelling situations. The development of variable demand methods does not imply a commitment to using them in all circumstances. There are also questions about the appropriate level of **segmentation** (trip purposes etc.).

As an example of a highly specified demand model, we are applying the model developed in our work for the Department on the Manchester Motorway Box (Skinner *et al*, 1997). This involves a detailed zoning system, with travel demand segmented into: home to work and employers' business, work and employers' business to home, home to education, education to home, home to other, other to home, and non-home based, and further segmentation between car-owning and non-car-owning households; three modelled time periods; car, public transport and walk/cycle modes (and freight); (macro) time of day, mode, destination and frequency choices modelled using a hierarchy of logit models, with the order of choices varying by trip purposes.

A number of key issues are being addressed. By **Convergence**, we intend both the **methods** (or algorithms) by which equilibrium is attained, and the **accuracy** with which it is attained. A related issue is the definition of convergence criteria themselves. We have a number of possibilities, including damped 'cobweb' approaches, and the single and bi-level fully integrated approaches. These reflect two factors – the necessary level of complexity, and the current availability of software.

In terms of the **Modelling of Demand**, there are a large number of possibilities: the aim is to cover the range between the elasticity-based approaches as envisaged in the current Guidance (ie, incorporating the "separability" restriction) and the "full" multimodal demand model incorporating a number of different transport choices, as described above. A crucial element here is the level of complexity necessary to obtain results adequate for appraisal. If comparable results can be obtained by simpler demand functions, then such methods would be preferred on practical grounds.

One of the main comparisons is of the **efficiency** of the computational process - that is, for a given demand model, how do the different algorithms compare in terms of the level of convergence achieved and in the time taken to reach that level? Other comparisons relate to **economic evaluation.** In this case, there are questions of sensitivity **both** to the level of convergence **and** to the choice of demand function

Preliminary indications are that after allowing some "tuning" with reasonable heuristic ("cobweb") methods, objective function methods could achieve similar levels of convergence in about one-fifth of the computing time. At the same time, the uncertainty about convergence with heuristic methods is removed.

These techniques potentially offer a significant improvement in the accessibility of multi-modal transport modelling techniques and the DETR is exploring how the techniques can be most effectively promulgated when testing is completed. However, an extensive programme of testing must be completed before they can be released for general use.

ACKNOWLEDGMENT

The work reported in this paper is being carried out for HETA Division of the UK Department of Environment Transport and the Regions (DETR). The views ex4resped are the authors' and do not represent those of the DETR.

REFERENCES

Arezki Y, Hall M D, Hyman G & Mackenzie I (1996). Simultaneous Assignment Distribution and Mode Choice in NAOMI (also indexed as Arezki, Hall & Hyman: A Simultaneous Incremental Equilibrium Assignment and Dual Trip Distribution for Multiple User Classes), PTRC European Transport Forum, Seminar D Proceedings.

Beckmann M J, McGuire C B & Winston C B (1956), Studies in the Economics of Transportation, Yale University Press, New Haven, Conn

Design Manual for Roads and Bridges, Induced Traffic Appraisal, Vol 12, Section 2, Part 1

Evans S P (1976), Derivation and Analysis of some Models for Combining Trip Distribution and Assignment, Transportation Research, **10**(1), 37-57

Oppenheim N (1995), Urban Travel Demand Modelling: From Individual Choices to General Equilibrium, John Wiley & Sons.

Ortúzar J de D & Willumsen L G (1994), Modelling Transport, John Wiley & Sons

SACTRA (1994), Trunk Roads and the Generation of Traffic, HMSO

Sheffi Y (1985), Urban Transportation Networks, Englewood Cliffs, NJ: Prentice Hall

Skinner A, Coombe D, Bates JJ, Fowkes A and Hyman G (1997), Feasibility of Measuring the Induced Traffic Effects of the Manchester Motorway Box, PTRC European Transport Forum, Transportation Planning Methods Seminar (Volume I)

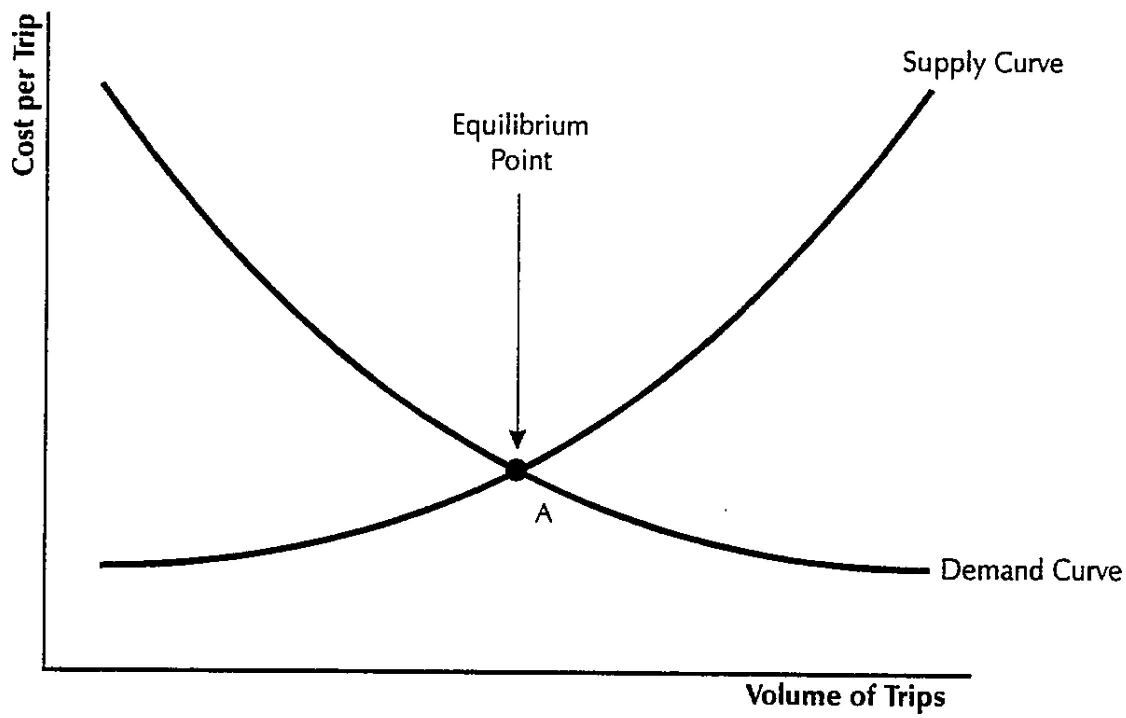

Figure 1: Demand/Supply Equilibrium

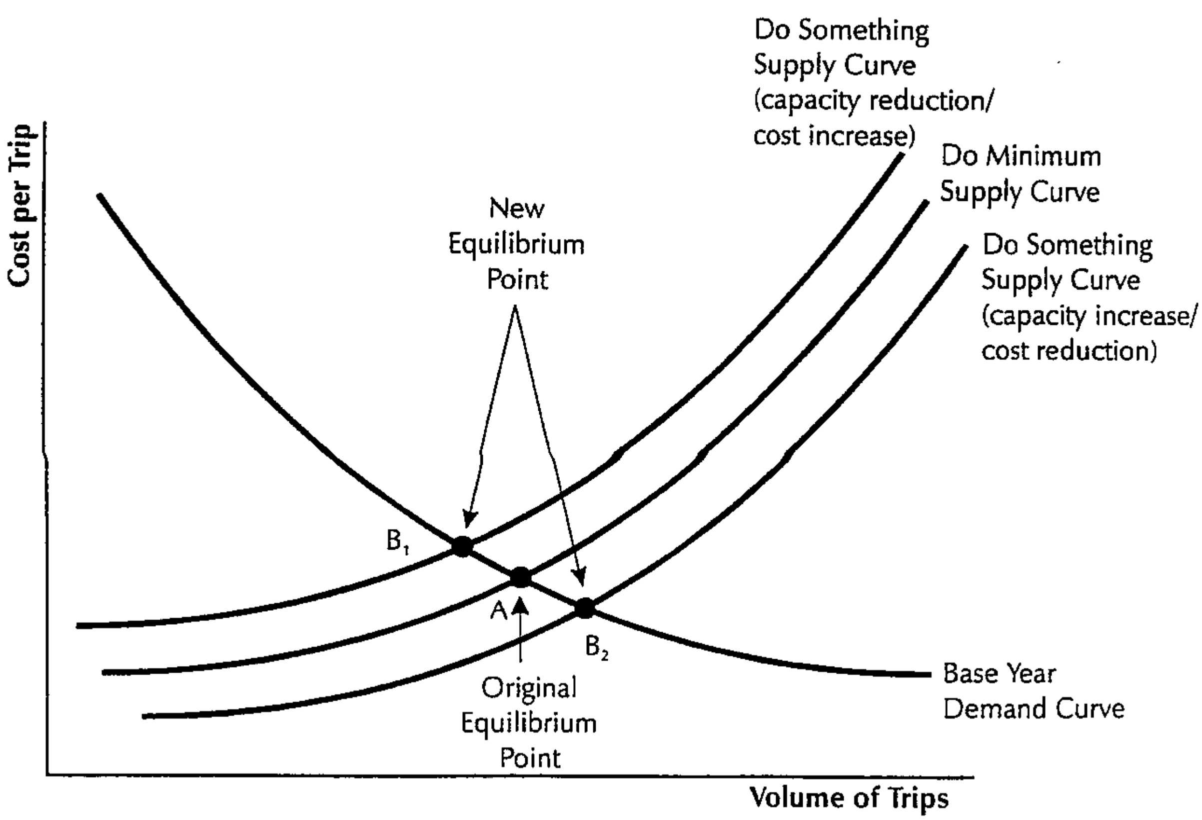

Figure 2: Appraisal in the Base Year