PIVOT-POINT PROCEDURES IN PRACTICAL

TRAVEL DEMAND FORECASTING

Andrew Daly [*], James Fox, Jan Gerrit Tuinenga

RAND Europe

1.   INTRODUCTION

For many cities, regions and countries, large-scale model systems have been developed to support the development of transport policy. These models are intended to predict the traffic flows that are likely to result from assumed exogenous developments and transport policies affecting people and businesses in the relevant area. The accuracy of the model over a wide range of policy measures and exogenous developments is crucial to determining the quality of the information that can be extracted as input to the planning and policy analysis process.

A frequent approach to modelling, which can substantially enhance the accuracy of the model, is to formulate the model as predicting *changes* relative to a base-year situation. Often, base-year traffic flows can be observed rather accurately and the restriction of the model to predicting differences reduces the scope for errors in the modelling – whether they be caused by errors in the model itself or in the inputs to the model – to influence the outputs. Such approaches are called 'pivot point' methods, the name given to them by Manheim (1979), or sometimes incremental models. The approaches have proved themselves beneficial in practical planning situations and now form part of the recommended 'VaDMA' procedures in the UK.

While the error-reducing principle of the pivot point is clear, the implementation of the principle in practical model systems can be done in a number of ways and the choice between these can have substantial influence on the model forecasts. In particular modellers need to consider:

[*]   Corresponding author, e-mail: daly@rand.org .

- whether the change predicted by the model should be expressed as an absolute difference or a proportional ratio, or whether a mixed approach is necessary;

- how to deal with growth in 'green-field' situations, when the future is likely to be very different from the present situation, when applying these approaches;

- at what level in the model should the pivoting apply, i.e. at the level of mode choice, destination choice, overall travel frequency or combinations of these;

- whether the pivoting is best undertaken as an operation conducted on an explicit 'base matrix' or the model is constructed so that it automatically reproduces the base year situation with base year inputs.

This paper reviews the alternative approaches to each of these issues, discussing current practice and attempting to establish the basis on which alternative approaches might be established; in particular, whether pivoting is treated as a correction to a model which is in principle correctly specified but incorporates some error, perhaps from faulty data, or as a partial replacement for a model that handles at best part of the situation. These views of the pivoting lead to different procedures.

In this paper 'pivoting' is used as a broad term, describing the use of a model to predict changes relative to a fixed or more reliable base point. The contrast is with the use of a model alone to produce purely 'synthetic' (sometimes called 'absolute') forecasts. One may distinguish two broadly separate procedures for pivoting:

- the use of two clearly identifiable synthetic model results, applicable to base and forecast cases, which are then used, either as a ratio or as a difference, to adjust the base point – we call these factor pivoting and difference pivoting;

- the specification of a model which can only predict changes relative to a base, – we use the common name of incremental modelling.

It should be clear that either of these pivoting methods will reproduce the base when the forecast case has inputs equal to the base inputs. Moreover, as we shall see, there is little difference in practice between an absolute forecast in which large numbers of correction terms (sometimes called 'K factors') have been included, factor pivoting and incremental modelling. Details are explained below.

The following section of the paper reviews current practice in using pivot-point and related procedures in the United Kingdom, chosen because UK practice is reasonably

2

well-developed and gives a good basis for illustrating the issues. We then discuss theoretical approaches to the problem, followed in section 4 by a presentation of our own experience. The final section presents conclusions and recommendations for future practice.

2. CURRENT UK PRACTICE

2.1 Multi-Modal Studies

The programme of Multi-Modal Studies was commissioned by the UK Government in 1999 and 2000, and all of the studies are now complete and have reported their findings. The Multi-Modal Studies provide a good picture of current UK practice, as most of the main UK transport consultancies were involved in the studies.

At the outset of the Multi-Modal Studies, the then Department of the Environment, Transport and the Regions issued the Guidance on the Methodology for the Multi Modal Studies (GOMMMS). However the GOMMMS advice contained no discussion of pivoting procedures and therefore no guidance was given to study teams on this issue.

The Department for Transport's WebTAG site provides links to a number of the Multi-Modal Studies[†]. Where available, information has been assembled on the modelling approach used in these studies. Table 1 summarises the application of pivoting approaches in those studies for which information was available.

[†]   http://www.webtag.org.uk/links/lmmstudies.htm

Table 1:   Incremental Modelling in the UK Multi-Modal Studies

Study	Modelling Team	Pivoting?	Notes
SWYMMS	MVA	Incremental adjustment factors at matrix level	Essentially factor pivoting
SEMMMS	SDG, Atkins	Growth factors by mode	Report lacks detail
Norwich to Peterborough	Atkins	Trip end growth factors from TEMPRO	Separately for car & PT
M1J19	Halcrow	Trip end growth factors from TEMPRO	Segmented by area type and time period
CHUMMS	Mouchel/WSP/ Oscar Faber	None, synthetic forecasts applied directly	Generation & distribution predicted by land use model

Following the completion of the studies, Bates *et al* (2003) were commissioned to evaluate the modelling and appraisal aspects of the studies. Although this review considers a number of modelling aspects in depth, it does not discuss the decision as to whether to employ incremental approaches.

2.2   VaDMA Advice

In the UK, the Department for Transport has issued the Variable Demand Modelling Advice (VaDMA), which has been designed as a reference document for use during the model design stage, with sections outlining the issues to be considered at each stage of model development. VaDMA discusses the decision as to whether to apply the 'synthetic' model predictions directly, or to employ an 'incremental' approach where changes are predicted relative to an observed base.

VaDMA states that 'the Department's preferred approach is to use an incremental rather than a synthetic model, unless there are strong reasons for not doing so'. The advice qualifies this statement by noting that 'a purely incremental approach may not be sensible where there are large changes in land use between the base and forecast years, which will significantly change the distributions of origins and destinations'.

To deal with the problem of empty cells in the observed base matrices, VaDMA suggests that a weighted average of observed and synthetic matrices are used in the base. This approach gives greater weight to cells where there are more observed trips than expected from the synthetic model.

While VaDMA is clear in its recommendation that incremental approaches should usually be used, it does not provide information on the particular approaches that

should be employed. It is understood that the forthcoming WebTAG guidance[‡] on variable demand modelling will provide more detailed guidance on this particular issue.

2.3 Other Sources

Ortúzar and Willumsen (2001) discuss what they term pivot-point modelling in a chapter on simplified demand models in their widely used *Modelling Transport* text. Pivot-point modelling in their terminology is the application of a logit model using differences in utility between base and future situations:

$$p_k^{'} = \frac{p_k^0 \exp(V_k - V_k^0)}{\sum_j p_j^0 \exp(V_j - V_k^0)} \tag{2.1}$$

where: p_k' is the proportion of trips using alternative k;

p_k^0 is the original proportion of trips by alternative k; and

$(V_k - V_k^0)$ is the change in utility of using mode k.

Ortúzar and Willumsen note that such incremental forms are not difficult to develop or implement, and have the advantage of preserving the base matrices in application. This approach is more often termed 'incremental logit' modelling.

Bates *et al* (1987) discuss the theory and application of a nested incremental logit model for access choice to London airports. The basic form of the incremental models they use is as given in Equation (2.1), but they develop the model further to include nested structures. They achieve this by representing the nested structure as a series of linked binary choices. The authors present a procedure for inferring the mode constant of new modes by rating the mode on a series of attributes which can be compared against existing modes and their associated attributes.

Abraham et al. (1992) describe the development on an incremental four stage model for London to evaluate major rail schemes. The model was developed from the existing LTS model, a conventional four stage model applied in an absolute fashion. It was important that the model produced results consistent with the LTS model, while

---

[‡]   Web-based transport analysis guidance (WebTAG) is issued by the UK Department for Transport and is made freely available via the internet .

at the same time producing results more rapidly than LTS, which had significant run times. The incremental formulation allows the development task to be broken into individual stages, although in this case there is no incremental trip generation model, rather forecasts are taken directly from LTS. The focus on rail schemes, and the need to achieve more rapid run times than the LTS, meant that only home-work trips were modelled, for all other purposes fixed trip matrices from the LTS were used.

The basic formulation for the incremental models used is as given in Equation (2.1) above. Separate mode choice and distribution models were estimated for blue collar and white collar workers, who form separate segments in the LTS. Trips are also segmented into 0, 1, 2+ car owning households in the mode split model, and public zero car, public some car and private in the distribution model. The distribution model was estimated and applied in a doubly-constrained fashion to ensure full consistency with both trip generations and trip attractions.

3.   PIVOT THEORY

In most large-scale model systems, forecasts are developed of changes in a series of traveller decisions (frequency, destination, mode and departure time choice, for example) which are expressed in the form of forecast matrices which are then assigned to highway and/or public transport networks. The objective of the pivoting approach in this context is therefore to predict changes in the base matrix.

The reasoning behind the choice of a pivoting approach is chiefly that it reduces error. When the base matrix can be estimated with greater accuracy than can be achieved by a model, it makes sense to use the base matrix to predict 'most' of the trips, using the model to predict the changes relative to that, obtaining a smaller error from the model because a smaller number of trips is being predicted.

In this section of the paper, three theoretical issues are discussed. First, the equivalence of several factoring approaches is established. Then a comparison is made between factoring and additive approaches, and this is followed by a discussion of the issues arising in multi-stage modelling. Finally, a review is given of the competing advantages of the different approaches.

3.1   Equivalence of factoring approaches

To define the notation, suppose there is a synthetic (logit) model predicting the probability of choosing each of a number of alternatives

$$S_i = \exp V_i / \Sigma_j \exp V_j$$

where V gives the utility of each alternative.

The incremental approach (Bates, Ashley and Hyman (1987) quote a 1980 paper by Kumar) predicts P_i choices for alternative i (our notation)

$$P_i = B_i \exp \Delta V_i / \Sigma_j B_j \exp \Delta V_j$$

where B is the observed (base) probability of choosing each alternative and

ΔV gives the change in utility.

The synthetic model can be used to make synthetic predictions S*

$$S^*_i = \exp (V_i + \Delta V_i) / \Sigma_j (\exp V_j + \Delta V_j)$$

So we can calculate

$$S^*_i/S_i = \{ \exp (V_i + \Delta V_i) / \exp V_i \} . \{ \Sigma_j (\exp V_j) / \Sigma_j (\exp V_j + \Delta V_j) \}$$

$$= \exp \Delta V_i . \text{constant}$$

since the term in the second { } is a global constant. Since the global constant cancels out the incremental calculation can then be expressed

$$P_i = (B_i S^*_i/S_i) / (\Sigma_j B_j S^*_j/S_j)$$

The incremental calculation is therefore exactly the same as factoring the observed choices B by the ratio of the synthetic forecasts, ***providing the total number of choices is normalised to remain constant***.

Similarly, consider a model in which a sufficient number of correction terms ('K factors') have been included so that the synthetic model in the base case gives exactly the observed results. This will happen if we use a revised utility function

$$W_i = V_i + K_i = V_i + \log B_i / S_i$$

This then gives a revised base model result

$$R_i = \exp W_i / \Sigma \exp W_j = (B_i/S_i) \exp V_i / \Sigma (B_j/S_j) \exp V_j = B_i$$

since once again a global constant cancels out and then $\Sigma\, B_j = 1$. We can then make the forecast

$$R^*_i = \exp W^*_i / \Sigma \exp W^*_j = (B_i / S_i) \exp V^*_i / \Sigma\, (B_j / S_j) \exp V^*_j$$

$$= B_i \exp \Delta V_i / \Sigma\, B_j \exp \Delta V_j = P_i$$

since once again the global constant cancels out. Thus a model with sufficient K factors is simply the same as using an incremental model.

It would be possible to argue that the approach using K factors has the advantage that factors can be considered for their plausibility and decisions made to include or exclude specific factors on the basis of empirical information. In the incremental or factoring approaches it is not possible to make decisions of this type, although the factoring approach can be adapted to use mixtures of synthetic and factored (and even difference) approaches, as will be seen in Section 4.

Any of these methods may require adaptation because of the existence of zero or near-zero values in one or other of the inputs, for example in 'greenfield' cases. Examples of such adaptations are also given in Section 4.

3.2   Additive approach

In the previous section, the basic equivalence of incremental, factoring and K-factor methods was noted. However, a substantially different approach is the use of synthetic matrices to adjust base matrices by adding the difference between forecast and base-year synthetic, i.e. in the notation of the previous section

$$P_i = B_i + S^*_i - S_i$$

Zero inputs do not cause a specific problem with this approach. However, a problem that may be encountered is that a negative result may be obtained, at which point the output value zero would normally be substituted. However, if all the results remain positive without correction the calculation will maintain the number of forecast trips as expected, i.e.

$$\Sigma_i\, P_i = \Sigma_i\, S^*_i$$

To set against this advantage, the chief disadvantage of the difference approach are that it relates poorly to the specification of most of the models that are used. Specifically, the models predict proportions of trips making each choice and there is

no way to correct the models so that a given difference in the number of trips will be predicted. The models are fundamentally models of choice and the appropriate correction for such models is by factoring, not by adding a correction.

In another important example, the numbers of trips predicted by our models are proportional to the 'sizes' of the generating and attracting zones. An error in measuring these sizes cannot easily be corrected in a way that would yield a fixed number of additional car trips in the morning peak.

For this reason, the use of differences must be considered as an unusual approach, for application only in special circumstances.

In 'greenfield' cases, where the base observed and base synthetic are both zero or near zero, an additive approach is more-or-less equivalent to using the synthetic forecast. In fact there is little choice in this context – the synthetic forecast is the only information available.

3.3   Pivoting in multi-stage models

Many model systems are composed of multiple stages, with a number of travel choice decisions treated in sequence: travel frequency, destination, mode, departure time, etc., each stage further splitting the demand forecast by the previous stage. In these cases, pivoting can be considered at each stage in sequence or as a final step.

When the model is formulated incrementally (see Bates, Ashley and Hyman, 1987) or with K factors, pivoting is effectively conducted for each step in sequence. However, when pivoting is performed by factoring with synthetic model output, the choice of whether or not a normalisation of the total trips at each stage is a real issue. This is because there is no guarantee that the sum of factored trips is equal to the forecast total, as has already been noted in section 3.1.

An important special case is when the base matrix is defined with a different level of geographical aggregation than some components of the model. For example, the base may be defined with an aggregated zoning system or with fewer purpose categories than the model to enhance the reliability of the base data. Alternatively the matrix may be defined over public transport stations, which occurs particularly often when ticket sales information is an important data source. These approaches can be useful to maintain proper respect for the base matrix data.

3.4   Considerations in the choice of approach

In choosing between pivoting approaches, we have argued that the additive approach should be used only in special circumstances.   Some of these circumsta nces are described in our practical experience set out in section 4.

In choosing between the various factoring approaches, it should be noted that the factoring procedure with explicit base, base synthetic and forecast synthetic matrices is clear, simple to program (because the same program can be used for the two synthetic matrices) and the base matrix is prominent as a component of the modelling process.

On the other hand, the K factor and incremental approaches give automatic normalisation at each stage, while the K factor approach has the additional feature of making it easier to test whether corrections are needed in each cell of the model. But if normalisation is carried out at each stage in multi-stage models there is effectively no difference betwee n any of these methods in terms of the results obtained.

In practice, a decisive consideration may be the need to adapt the pivoting method to accommodate zeros and other difficult cases. This is rather easier to undertake if the factoring method is used – a description of some of the adaptations is given in Section 4.

4.   RAND EUROPE EXPERIENCE

4.1   Introduction

In this section we describe some of our own experience in a number of large -scale modelling studies. In these studies, we have consistently applied pivoting using the factoring of synthetic matrices, for the reasons given in section 3.4.

Pivoting is carried out at matrix cell level. That is, for a specific origin, destination, mode, time of day and purpose, adjustments are made relative to the corresponding cell in a base matrix.

As has been noted, issues may arise that affect marginal totals in the matrices and corrections may be necessary for those. However, for the sake of clarity of this section of the paper, this issue is neglected here. The specific procedures proposed in this section can also be adapted for row or matrix pivoting.

The procedures set out below are based on RAND Europe's experience with a number of pivot-point models used in their transport demand forecasting systems. Some of these models have special procedures to adjust the calculation when the growth in a specific cell is considered to be 'extreme' (some of our models identify 'extreme' cells automatically, while some others rely on manual selection). Extreme growth situations can often be identified with greenfield sites. Our preferred approach involves automatic selection of the 'extreme' cases and is therefore reflected in the procedures set out below.

4.2   Calculation procedures

For the reasons explained above, the preferred approach to pivot-point forecasting is to apply the ratio of model outputs for base and forecast situations as a growth factor to the base matrix, i.e. in a given cell the predicted number of trips P is given by

$$P = B \cdot \frac{S_f}{S_b} \qquad\qquad (4.1)$$

where: S_f is an OD-matrix containing synthetic trips for a future year
S_b is an OD-matrix containing synthetic trips for the base year
B is the observed (base) matrix

However, two considerations mean that it is not always possible to apply this calculation as simply as stated.

First, any combination of the three components on the right hand side of this equation may be zero (or very small) making the calculation impossible or meaningless. Eight possible cases arise (combinations of zero values) and these are dealt with separately in the recommendations below.

Second, particularly when there is a land-use change affecting a currently undeveloped zone, the change may be quite extreme and strict application of the formula above can lead to an 'explosion' in the number of trips. In these cases it is better to pivot by applying the difference method, i.e. $(S_f - S_b)$, to the base matrix, rather than a factor as shown above. In the recommendations below, difference pivoting is applied to all cases when S_b is zero and to other cases when S_f / S_b exceeds a specified factor (including 'infinity' when $S_b=0$ and S_f is non-zero).

The eight possible cases and the recommended treatments are set out in the table below.

Table 2: **Eight possible cases**

Base (B)	Synthetic Base (S_b)	Synthetic Future (S_f)	Predicted (P)		Cell Type
0	0	0	0		1
0	0	>0	S_f		2
0	>0	0	0		3
0	>0	>0	Normal growth	0	4
			Extreme growth	$S_f - X_1$	
>0	0	0	B		5
>0	0	>0	$B + S_f$		6
>0	>0	0	0		7
>0	>0	>0	Normal growth	$B \cdot S_f / S_b$	8
			Extreme growth	$B \cdot X_2 / S_b + (S_f - X_2)$	

To complete the specification of the calculation it is necessary to specify the X variables, to define when a cell is considered to be zero (our experience has led us to use a test value of 10^{-3}) and when and how extreme growth (for cases 4 and 8) is to be applied.

With respect to the last point, in the extreme growth cases (4) and (8) it can be seen that the standard factor function is used initially, up to the limit when S_f is X_1 (case 4) or X_2 (case 8), and from that point an absolute growth is applied. In case (4) the starting point for absolute growth is 0, in case (8) it is $B \cdot X_2 / S_b$.

The definitions for X_1 and X_2, using a switching factor G and parameters k_1 and k_2, are given by:

$$X_1 = B \cdot G \qquad\qquad (4.2)$$

$$X_2 = S_b \cdot G \qquad\qquad (4.3)$$

$$G = k_1 + k_2 \cdot \max (S_b/B, \ k_1/k_2)^{\S} \qquad\qquad (4.4)$$

---

§ Note that G can be written $G = 1 + \dim(k_2 \cdot S_b/B, \ k_1)$, i.e. G is never less than 1 and exceeds 1 when B $< k_1 \cdot k_2 \cdot S_b$.

where: $k_1,\ k_2 > 0$ [**]

G is the switching factor

Note: that when $B \to 0$ $X_1 \to k_2.S_b$ $X_2 \to \infty$;

The formula indicated thus determines the switch point to absolute growth by the relation between B and S_b.   The formula has proved to be satisfactory in practice.

Recent RAND Europe experience in the application of pivoting theory to PRISM, a disaggregate model system for the West Midlands region of the UK, demonstrated that careful consideration needs to be given to the choice of the switching factor, and the importance of testing the application of the pivoting procedure to ensure the plausibility of the model forecasts.

In the PRISM model, a detailed zoning system containing around 900 zones was used. The base matrices were based on expansion of survey data, and consequently were significantly sparser in coverage than the synthetic model matrices[††]. This meant that when *both* B and S_b are non-zero, B was larger on average than S_b because a significant volume of S_b occurs in cells where B is exactly zero. The impact of this characteristic was that the trigger point for extreme growth was not being reached even in cases where Sf >> Sb, because both values were smaller than B. Modification to the switch point calculation to use the formula above solved this problem and following this change plausible model forecasts were obtained.

4.3   Continuity

Continuity of the results is an important issue in the use of the procedures as set out in the previous section. Clearly, transitions between any of the eight cases as a result of small changes in either of the matrix cells should not lead to substantial changes in predicted values. Therefore we have tested the procedures for continuity in a number of directions.

[**] The values k_1=0.5 and k_2=5 are commonly used.

[††] Probabilistic models predict trips to all available destinations, but the number of trips predicted to unattractive destinations is very small. By contrast an expanded matrix will contain no trips to unattractive destinations unless they happen to be sampled.

First of all, continuity in S_f should be considered to be the most important by far, since discontinuities could produce counter-intuitive differences between forecasts made from the same base. The procedures specified in the table are fully continuous in S_f, so there are no problems in this most important group of transitions.

Secondly, transitions in B, S_b and between X's are much less important as they do not occur in standard applications of pivot-point based models but only when changes in base values are considered. In some of these transitions we have identified a discontinuity resulting from the use of these procedures, they would however not occur in standard applications, but are cases for potential further investigation at a later point in time.

4.4   Results

The method described here has been used in the past years in a number of transport models implemented in various countries and has proven to yield acceptable results. Illustration of the results of the method is not very straightforward, however. Therefore we show the results of the most common case from the table above: case 8, where all OD-matrices have positive values and we show the damping of the effects when extreme growth arises.

First of all, it is important to note that we are mostly interested in the effect that a change in the synthetic future year matrix (S_f) would have on the pivoted outcome (P). For that reason, we can rewrite the normal growth part of case 8 from

$$P = B \cdot \frac{S_f}{S_b}$$

as

$$P = \frac{S_f}{S_b/B},\qquad\qquad(4.5)$$

which – in the normal growth situation – shows the linear relationship between P and S_f, apart from the factor as given in the denominator of (4.5), which can be seen as an indicator of the match between the synthetic base matrix S_b and the base year value (B). As, in general, the denominator of (4.5) remains equal for a given cell between different model applications, it is interesting to analyse the effects of different switch

points G (which are linearly dependent on S_b/B as can be seen in (4.4)) on the predicted outcome.

The results of applying the formulae for case 8 (using the commonly used values for the parameters: $k_1=0.5$ and $k_2=5$), are demonstrated in the figure below. In this figure we have used several different values for the various matrices involved and consequently seve ral sets of predictions are derived, clustered by the match between the two base year matrices: S_b/B (using lines for each value of this ratio as given in the legend).

Pivot results by base ratios (S_b/B)
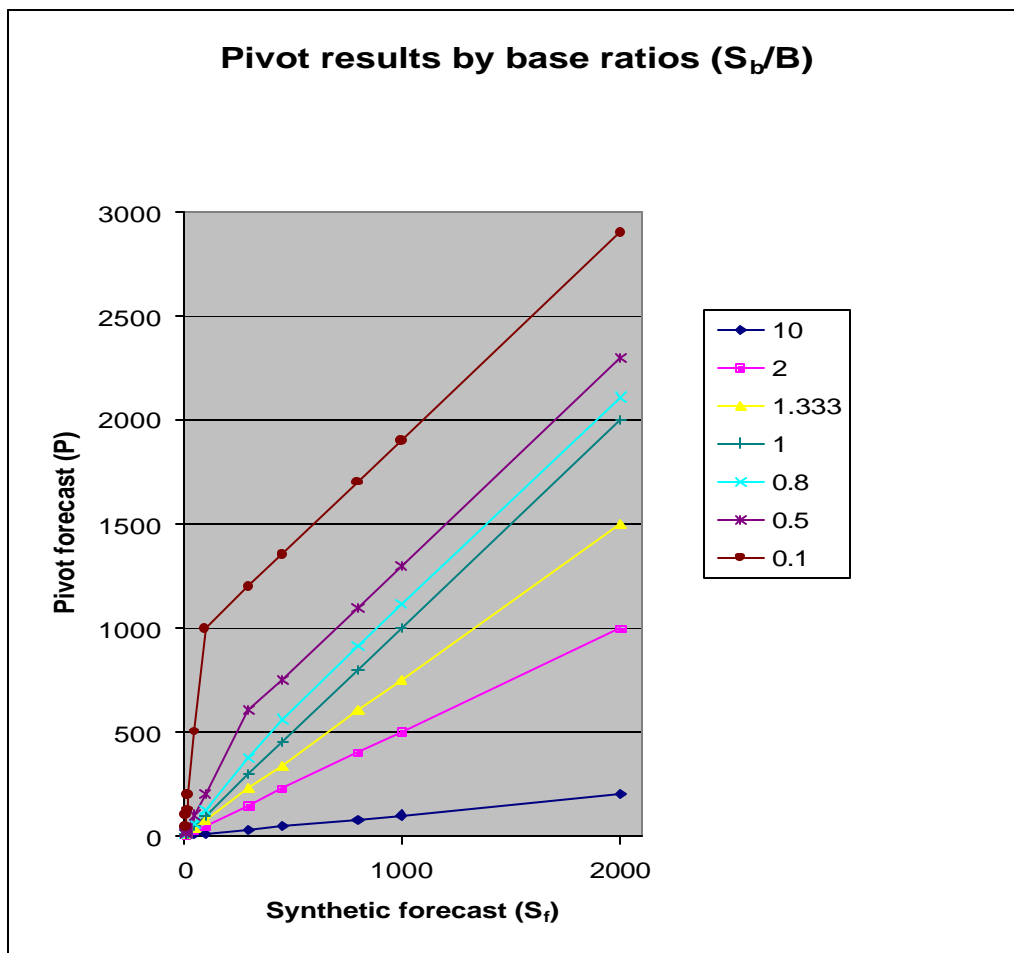
As is clear from this figure the results demonstrate that the switch from relative to absolute pivoting is applied when the ratio S_b/B is smaller than 1.0 and also that the switch point is applied earlier (with lower values for S_f) when the ratio S_b/B gets much smaller than 1.0. This reflects the requirement to damp the predicted outcomes and to use more absolute pivoting when cells of these two matrices are less comparable to one another.

A somewhat more intuitive picture is given next, demonstrating the switch point more clearly. In this picture, the axes of the picture above are normalised to their respective base values to compare the synthetic growth factors with the obtained growth factors after applying the pivot procedures as set out above. The picture shows clearly that when base ratios (i.e. S_b/b) are lower the switch point to absolute pivoting is applied earlier (i.e. with lower synthetic growth ratio S_f/S_b). For base ratios over 1, the predicted growth factor is equal to the synthetic growth factor (as a consequence in the figure there is an overlap for all lines with base ratios equal and over 1).

Growth factors by base ratios (S_b/B)
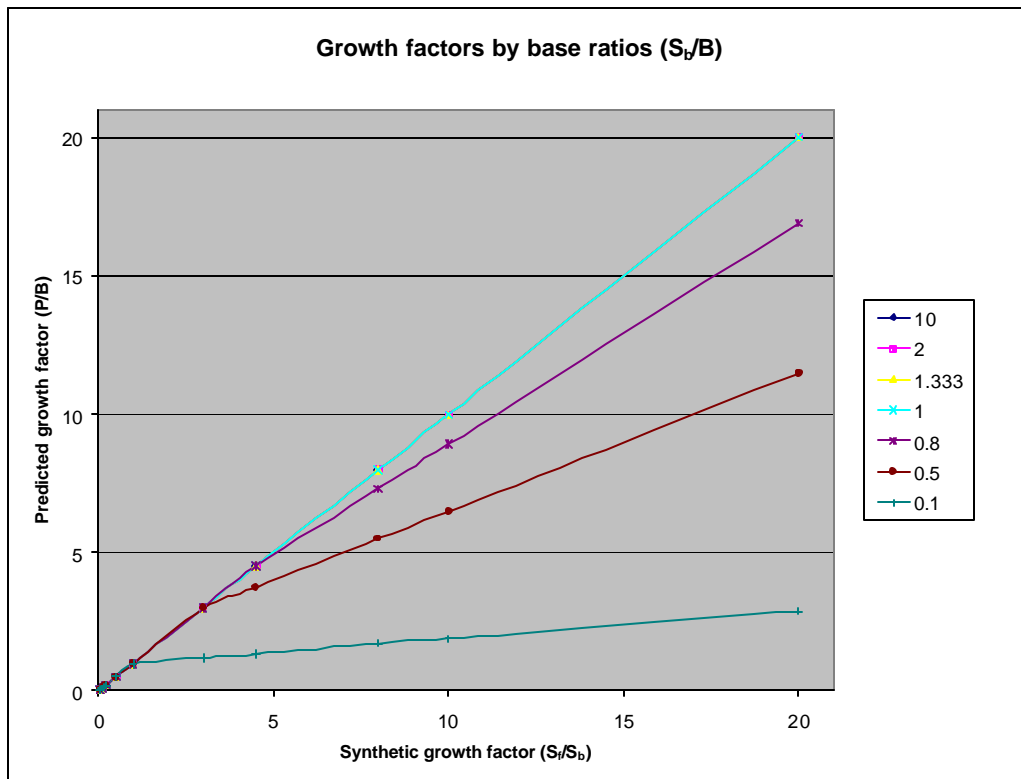
5. CONCLUSIONS

The use of pivoting is widespread in practical transport planning and has substantial advantages for reducing error in forecasting. However, it has not generally been made clear what the alternative procedures are and how a selection between these should be made to obtain the best results.

In British practice both factoring and difference approaches to pivoting are used. The Department for Transport strongly advises the use of pivoting but does not advocate

one specific method. Other literature that we have found primarily focuses on the use of incremental modelling.

However, incremental modelling, factor pivoting and the use of K factors are closely equivalent, except when factor pivoting is used without 'normalising' to correct the total number of choices made. In general, forecasts should be normalised, because the function of each model is to predict proportions choosing each alternative, while the total number of choices to be made is usually the function of another model. The use of difference pivoting should be restricted to special cases.

In practical experience, a system of automatic selection of the appropriate pivot procedure for a range of cases can be set up. We have explained the procedures that we have developed in our experience over a number of large-scale modelling studies. These procedures have a number of attractive properties, such as that they are largely continuous, i.e. small changes in the inputs are unlikely to lead to large changes in the outputs.

In summary, we find that pivoting is a useful procedure which can improve the accuracy of model forecasts, but that care is required in practice. We hope our experience presented here will be of assistance to others.

REFERENCES

Abraham, H, N .Shaw and L Willumsen (1992) *A Micro-Based Incremental Four Stage Transportation Model for London*, *Proceedings 20th PTRC Summer Annual Meeting*, UMIST, September 1992, England.

Bates, J., D. Ashley and G. Hyman (1987) *The Nested Incremental Logit Model: Theory and Application to Modal Choice*, 15[th] PTRC Summer Annual Meeting, University of Bath, England.

Bates J., P. Mackie and D. Forster (2003) *Evaluation of the Multi-Modal Study Process*, prepared for the UK Department for Transport.

Manheim, M.L. (1979) *Fundamentals of Transportation System Analysis*. The MIT Press, Cambridge, Mass.

Ortúzar J. de D., L. G. Willumsen (2001) *Modelling Transport*, *Third Edition*, John Wiley & Sons Ltd, Chichester, UK.