



EUROPE

CHILDREN AND FAMILIES
EDUCATION AND THE ARTS
ENERGY AND ENVIRONMENT
HEALTH AND HEALTH CARE
INFRASTRUCTURE AND
TRANSPORTATION
INTERNATIONAL AFFAIRS
LAW AND BUSINESS
NATIONAL SECURITY
POPULATION AND AGING
PUBLIC SAFETY
SCIENCE AND TECHNOLOGY
TERRORISM AND
HOMELAND SECURITY

The RAND Corporation is a nonprofit institution that helps improve policy and decisionmaking through research and analysis.

This electronic document was made available from www.rand.org as a public service of the RAND Corporation.

Skip all front matter: [Jump to Page 1](#) ▼

Support RAND

[Browse Reports & Bookstore](#)

[Make a charitable contribution](#)

For More Information

Visit RAND at www.rand.org

Explore [RAND Europe](#)

View [document details](#)

Limited Electronic Distribution Rights

This document and trademark(s) contained herein are protected by law as indicated in a notice appearing later in this work. This electronic representation of RAND intellectual property is provided for non-commercial use only. Unauthorized posting of RAND electronic documents to a non-RAND Web site is prohibited. RAND electronic documents are protected under copyright law. Permission is required from RAND to reproduce, or reuse in another form, any of our research documents for commercial use. For information on reprint and linking permissions, please see [RAND Permissions](#).

This product is part of the RAND Corporation technical report series. Reports may include research findings on a specific topic that is limited in scope; present discussions of the methodology employed in research; provide literature reviews, survey instruments, modeling exercises, guidelines for practitioners and research professionals, and supporting documentation; or deliver preliminary findings. All RAND reports undergo rigorous peer review to ensure that they meet high standards for research quality and objectivity.

TECHNICAL REPORT

Enhancement of the pivot point process used in the Sydney Strategic Model

James Fox, Andrew Daly, Bhanu Patruni

Prepared for the Bureau of Transport Statistics, Transport for NSW

The research described in this report was prepared for the Bureau of Transport Statistics, Transport for NSW.

RAND Europe is an independent, not-for-profit research organisation whose mission is to improve policy and decision making for the public good. RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

RAND® is a registered trademark.

© Copyright 2012 Transport for NSW

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the copyright holders.

Published 2012 by the RAND Corporation
1776 Main Street, P.O. Box 2138, Santa Monica, CA 90407-2138
1200 South Hayes Street, Arlington, VA 22202-5050
4570 Fifth Avenue, Suite 600, Pittsburgh, PA 15213-2665
Westbrook Centre, Milton Road, Cambridge CB4 1YG, United Kingdom
RAND URL: <http://www.rand.org>
RAND Europe URL: <http://www.rand.org/randeurope>
To order RAND documents or to obtain additional information, contact
Distribution Services: Telephone: (310) 451-7002;
Fax: (310) 451-6915; Email: order@rand.org

Preface

RAND Europe was commissioned by the Bureau of Transport Statistics (BTS) of Transport for NSW to enhance the pivoting procedure used in the application of the Sydney Strategic Transport Model (STM). The pivoting procedure combines the forecasts from the travel demand models with base matrix information describing base travel patterns to provide best-estimate forecasts of the future travel matrices for car and public transport modes.

The STM was designed by Hague Consulting Group (1997). In Stage 1 of model development (1999–2000), Hague Consulting Group developed mode-destination and frequency models for commuting travel, as well as models of licence ownership and car ownership. In addition a forecasting system was developed incorporating these components. In Stage 2 of model development (2001–02), RAND Europe, incorporating Hague Consulting Group, developed mode and destination and frequency models for the remaining home-based purposes, as well as for non-home-based business travel. Then, during 2003 and 2004, RAND Europe undertook a detailed validation of the performance of the Stage 1 and 2 models. Finally, Halcrow undertook Stage 3 of model development (2007), re-estimating the home-work mode-destination models, and at the same time developing models of access mode choice to train for home-work travel.

By 2009, some model parameters dated back to 1999, raising concerns that the model may no longer reflect with sufficient accuracy the current behaviour of residents of Sydney. Furthermore, changes to the zone structure of the model occurred with the number of zones approximately trebling in number and the area of coverage increased to include Newcastle and Wollongong. Therefore, the BTS commissioned RAND Europe to re-estimate the STM models using more recent information on the travel behaviour of Sydney residents.

Following the completion of the re-estimation project, RAND Europe was commissioned to undertake three parallel projects to implement the new models, and improve the performance of the pivoting process.

The first project was to implement the new frequency, mode and destination components in the STM. For each journey purpose represented in the STM, the frequency, mode and destination models are implemented within a single structure referred to as a travel demand model.

The second project was to update the Population Synthesiser used to generate future forecasts of the population in the STM study area by segment and zone. In application, the

outputs from the Population Synthesiser are used as inputs into the Travel Demand models.

The third project was to assess and enhance the pivoting procedure used in the STM. The work to update the pivoting procedure is described in this report.

This document is intended for a technical audience familiar with transport modelling terminology.

RAND Europe is an independent not-for-profit policy research organisation that aims to improve policy and decision making in the public interest, through research and analysis. RAND Europe's clients include European governments, institutions, NGOs and firms with a need for rigorous, independent, multidisciplinary analysis. This report has been peer-reviewed in accordance with RAND's quality assurance standards.

For more information about RAND Europe or this document, please contact James Fox:

RAND Europe
Westbrook Centre
Milton Road
Cambridge CB4 1YG
United Kingdom
Tel. +44 (1223) 353 329
jfox@rand.org

Contents

Preface.....	iii
Table of tables	vii
Acknowledgements.....	ix
CHAPTER 1 Introduction	1
1.1 Objective of work.....	1
1.2 Original pivoting process.....	1
1.3 Structure of remainder of report	3
CHAPTER 2 Performance criteria	5
2.1 Base matrix sparsity	5
2.2 Comparison of synthetic and predicted growth.....	5
2.3 Distribution of demand across the 8 cases.....	6
CHAPTER 3 Performance of original pivoting process.....	9
3.1 Test matrices	9
3.2 Zonal matrices.....	10
3.2.1 Definition of zero trips.....	10
3.2.2 Base matrix sparsity.....	11
3.2.3 Comparison of synthetic and predicted growth.....	12
3.2.4 Distribution of demand across the 8 cases	13
3.3 Station matrices.....	15
3.3.1 Base matrix sparsity.....	15
3.3.2 Comparison of synthetic and predicted growth.....	15
3.3.3 Distribution of demand across the 8 cases	15
3.4 Summary of issues with the original pivoting process.....	16
3.4.1 Zonal matrices	16
3.4.2 Station matrices	17
CHAPTER 4 Enhancements to the pivoting process	19
4.1 Zonal aggregation.....	19
4.1.1 Initial tests	19
4.1.2 Selecting the appropriate level of aggregation	20
4.1.3 Disaggregating demand back to detailed zone level	22

4.2	Modification to the extreme growth switch point	23
4.2.1	Why the case 8 extreme growth rule was modified	23
4.2.2	Selection of switch point parameter values.....	24
4.3	Normalisation.....	26
4.3.1	Mode-origin normalisation	26
4.3.2	Overall mode normalisation	27
CHAPTER 5 Performance of enhanced pivoting process		29
5.1	Zone to zone matrices.....	29
5.1.1	Base matrix sparsity	29
5.1.2	Comparison of synthetic and predicted growth	30
5.1.3	Distribution of demand across the 8 cases	32
5.2	Station to station matrices	34
5.2.1	Base matrix sparsity	34
5.2.2	Comparison of synthetic and predicted growth	35
5.2.3	Distribution of demand across the 8 cases	35
CHAPTER 6 Summary and recommendations		37
6.1	Zone to zone matrices.....	37
6.2	Station to station matrices	38
6.3	Recommendations	38
REFERENCES.....		39
	Reference list.....	41
APPENDICES		43

Table of tables

Table 1:	Eight pivoting cases	2
Table 2:	Sign change example.....	6
Table 3:	Example distribution across the eight cases	7
Table 4:	Test matrix sets.....	9
Table 5:	Zero trips definition tests, rescaled JT _W test matrix set, car driver.....	11
Table 6:	Zero trips definition tests, rescaled JT _W test matrix set, car passenger.....	11
Table 7:	Sparsity indices, rescaled JT _W and expanded HTS base matrices.....	11
Table 8:	Synthetic and predicted growth, rescaled JT _W and expanded HTS base matrix sets	12
Table 9:	Synthetic and predicted growth, case 8 cells, rescaled JT _W matrix set.....	13
Table 10:	Distribution of car driver demand, rescaled JT _W matrix set.....	14
Table 11:	Distribution of car driver demand, expanded HTS matrix set	14
Table 12:	Sparsity indices, station matrices.....	15
Table 13:	Comparison of synthetic and pivoted growth, station to station matrices	15
Table 14:	Distribution of rail demand, RailCorp matrix set.....	16
Table 15:	Impact of zonal aggregation on car driver predictions	20
Table 16:	Impact of zonal aggregation on predicted demand: synthetic growth ratios, all modes.....	20
Table 17:	SLA zone characteristics.....	21
Table 18:	Detailed 2690 zone characteristics	21
Table 19:	Sparsity indices after aggregation to 80 SLA zone level.....	23
Table 20:	k_2 tests, rescaled JT _W matrices.....	25
Table 21:	k_2 tests, expanded HTS matrices	25
Table 22:	Sparsity indices, rescaled JT _W base matrices.....	29
Table 23:	Sparsity indices, expanded HTS base matrices	30
Table 24:	Synthetic and predicted growth comparisons, rescaled JT _W base matrix set	30
Table 25:	Synthetic and predicted growth comparisons, expanded HTS base matrix set	32
Table 26:	Distribution of bus demand, expanded HTS matrix set, original pivoting process	33
Table 27:	Distribution of bus demand, expanded HTS matrix set, impact of aggregation	33

Table 28: Distribution of bus demand, expanded HTS matrix set, impact of revised case 8 switch point for extreme growth	34
Table 29: Sparsity indices, station to station rail matrices.....	34
Table 30: Synthetic and predicted growth comparisons, station to station matrices	35
Table 31: Distribution of rail demand, RailCorp matrix set, original pivoting process	36
Table 32: Distribution of rail demand, RailCorp matrix set, impact of revised case 8 switch point for extreme growth.....	36

Acknowledgements

The project team acknowledges the contribution of Frank Milthorpe, who has made a substantial contribution to the development of the STM over time. In this project, Frank has extracted matrix data for analysis, and run a number of pieces of supporting analysis.

The project team would also like to acknowledge the contribution of the two quality assurance reviewers, Charlene Rohr and Dimitris Potoglou, whose comments have improved the clarity and accessibility of this report.

1.1 **Objective of work**

The overall objectives of this work were to assess the performance of the pivoting procedure originally used in the Sydney Strategic Model (STM), and then to go on to specify enhancements to improve the operation of the procedure. In the brief for this work, the Bureau of Transport Statistics (BTS) stated that it were seeking an enhanced pivoting procedure that had the following features:

- is transparent, i.e. the working of the procedure should be clear to users of the STM
- does not require user intervention for each application
- retains the implied growth forecasts (as much as possible) from the STM at an appropriate geographic level, while accepting that there will be differences at the (origin-destination) cell basis
- can be applied to zone to zone forecasts
- can be applied to station to station forecasts
- can be applied to sparse (e.g. Household Travel Survey) matrices
- can be applied to relatively full (e.g. journey to work and rail operator station to station) matrices

1.2 **Original pivoting process**

This section describes the pivoting process used in the STM before the work to enhance the process was begun.

In forecasting travel demand it is fairly common to base future-year forecasts on an accurately known pattern of base-year observed travel flows defined in the base matrices. By focusing the modelling effort on predicting *changes* it is possible to make significant reductions in the expected forecasting error. Following Manheim (1979), the process of taking a fixed base point and making forecasts relative to that is called pivoting.

The pivoting approach used in the STM is the ‘eight-case method’, and was taken from the definition of the eight-case method given in Daly, Fox and Tuinenga (2005). The eight-case method implements factor pivoting, where the ratio of future and base demand model predictions is applied to the base matrices, as the default approach. Factor pivoting is applied for what is termed ‘normal growth’. However, the eight-case method also uses additive pivoting, where the difference between future and base demand model predictions

is applied to the base matrix for ‘extreme growth’ cases, for example in cases of brownfield sites where there is no travel demand in the base year. The pivoting method is carried out at matrix cell level and so can be termed cell factor pivoting. That is, for a specific origin, destination, mode and purpose, adjustments are made relative to the corresponding cell in a base matrix.

Table 1 sets out how the predicted matrix P is obtained from the cell factor pivoting process as a function of the:

- base matrix B
- ‘synthetic’ base Sb – base year output from the demand models before pivoting
- ‘synthetic’ future Sf – future year output from the demand models before pivoting

Eight specific cases are defined, allowing for situations where one or more of the items are 0.

The pivoting procedures have been chosen to optimise, as far as possible, the *continuity* of the process, i.e. to ensure that small changes in the inputs do not lead to large changes in the outputs. The ‘switching points’ X_1 and X_2 between normal and extreme growth, defined in Equations (1.1) and (1.2) below, are chosen on the basis of experience to fit with this requirement.

Table 1: Eight pivoting cases

Case	Base (B)	Synthetic base (Sb)	Synthetic future (Sf)	Predicted (P)	
1	0	0	0	0	
2	0	0	>0	Sf	
3	0	>0	0	0	
4	0	>0	>0	Normal growth, ($Sf < X_1$)	0
				Extreme growth ($Sf > X_1$)	$Sf - X_1$
5	>0	0	0	B	
6	>0	0	>0	$B + Sf$	
7	>0	>0	0	0	
8	>0	>0	>0	Normal growth ($Sf < X_2$)	$B.Sf/Sb$
				Extreme growth ($Sf > X_2$)	$B.X_2/Sb + (Sf - X_2)$

where:

$$X_1 = k_2 \cdot Sb \tag{1.1}$$

$$X_2 = k_1.Sb + k_2.Sb.\max\left[\frac{Sb}{B}, \frac{k_1}{k_2}\right] \quad (1.2)$$

The values used for the cell pivoting process originally used in the STM were $k_1 = 0.5$ and $k_2 = 5$, which are the example values given in Daly, Fox and Tuinenga (2005). These values for k_1 and k_2 emerged from practical experience in applying the Dutch National Model. Note that with these values for k_1 and k_2 , when $S_b/B < 0.1$ Equation (1.2) reduces to $X_2 = S_b$.

An issue that arises when working with synthetic trip matrices is that very low numbers of trips can be predicted for cells where the predicted probability of choosing that destination is very low. Rather than use these small numbers to predict very small numbers of trips, any cell where the number of trips is less than 0.001 is taken to be 0.

The ‘standard case’ is case 8 where none of the items is 0 and growth is ‘normal’ (termed case 8n), where the formula is:

$$P = B \frac{Sf}{Sb} \quad (1.3)$$

The analysis described in this report was undertaken during 2011, and at that time the pivoting process worked with the 2006 zoning system, which has a total of 2690 travel zones. Therefore, all analysis undertaken in this project used the 2006 zoning system. A new 2011 zoning system will to be developed during 2012 for use with 2011 Census data, resulting in minor changes to the travel zones.

1.3 Structure of remainder of report

The structure of the remainder of the report is set out as follows.

In Chapter 2, a set of performance criteria is defined to enable the performance of the pivoting process to be assessed. Then in Chapter 3, these criteria are used to assess the performance of the original pivoting process, and a number of issues with the performance of the original process are highlighted. In order to address these issues, Chapter 4 sets out three enhancements designed to improve the performance of the pivoting process. Next, Chapter 5 assesses the performance of the enhanced pivoting process, capturing the performance improvement that results from each of the three enhancements. Finally, Chapter 6 summarises the improvements that have been made to the process.

Three key performance criteria have been used in this work to assess the performance of the pivoting process:

- measures of base matrix sparsity
- comparisons of synthetic and predicted growth
- distributions of demand across the 8 pivoting cases defined in Section 1.2.

These three measures are defined and discussed in the following three sections.

2.1 **Base matrix sparsity**

To assess the sparsity of the base matrices for each mode, a new sparsity index has been defined for use in this project:

$$\text{Sparsity Index} = \frac{\text{number of cells}[Sb > 0]}{\text{number of cells}[B > 0]} \quad (2.1)$$

where: *number of cells[Sb>0]* is the number of cells in *Sb* (the synthetic matrix) that are greater than 0

number of cells[B>0] is the number of cells in *B* (the base matrix) that are greater than 0

We would usually expect a sparsity index greater than 1, as cells would only be 0 in *Sb* when it is not possible to travel between the origin and destination by the relevant mode. A high sparsity index indicates high sparsity in the base matrix, i.e. a lower proportion of cells with observed trips.

2.2 **Comparison of synthetic and predicted growth**

Synthetic growth expresses the growth predicted by the demand model. This is calculated as a percentage:

$$\frac{(Sf - Sb)}{Sb} * 100 \quad (2.2)$$

where: *Sb* is the synthetic base – base year output from the demand models before pivoting

S_f is the synthetic future – future year output from the demand models before pivoting

Predicted growth is the growth that is predicted after pivoting. This is calculated as a percentage:

$$\frac{(P - B)}{B} * 100 \quad (2.3)$$

where: B is the base matrix;

P is the predicted demand after pivoting.

It can be seen that when Equation (2.3) is applied at the individual cell level, synthetic and pivoted growth are the same in percentage terms. Thus the ‘standard case’ (case 8n) formula ensures that synthetic and predicted growth are consistent at the cell level. However, there is no guarantee of consistency at higher levels of aggregation than individual cells, because the distribution of base demand is usually different between the base matrix B and the synthetic base S_b . The simple example given in Table 2 illustrates how synthetic and predicted growth can match at the cell level but differ at the aggregate level.

An important issue when comparing synthetic and predicted growth is *sign changes*. It is possible that for a given mode there can be a difference between the sign of the synthetic and pivoted changes. The simple example shown in Table 2 for a single mode with two destination zones D1 and D2 illustrates how sign change can occur with the cell factor process currently used in the STM.

Table 2: Sign change example

	D1	D2	Total
S_b	10	10	20
S_f	9	12	21
$(S_f - S_b) / S_b$	-10%	+20%	+5%
B	15	5	20
P	13.5	6	19.5
$(P - B) / B$	-10%	+20%	-2.5%

This simple example shows how at the cell level, synthetic and pivoted growth match exactly, but that at the total mode level there is a sign change, which follows from the different distributions of demand in B and S_b across destination zones.

2.3 Distribution of demand across the eight cases

A third approach that has been used to investigate the performance of pivoting in other studies¹ is to analyse the percentage of cells that fall in each of the different cases. Table 3 illustrates this approach with an example based closely on real data from another project (for a full definition of each case, refer back to Table 1).

¹ In particular in analysis of the PRISM model system for the West Midlands region of the UK. Notes summarising this analysis were circulated to the PRISM project team but are not publicly available.

Table 3: Example distribution across the eight cases

Case	Formula	% Cells	% of B	% of S_b	% of S_f	% of P
1	0	5 %	0 %	0 %	0 %	0 %
2	S_f	0 %	0 %	0 %	0 %	0 %
3	0	0 %	0 %	0 %	0 %	0 %
4n	0	80.5 %	0 %	46.0 %	45.0 %	0 %
4e	$\max(S_f - 5.S_b, 0)$	0.5 %	0 %	0.1 %	0.6 %	0.3 %
5	B	0 %	2.0 %	0 %	0 %	2.0 %
6	$B + S_f$	0 %	0 %	0 %	0 %	0 %
7	0	0 %	0 %	0 %	0 %	0 %
8n	$B \times S_f / S_b$	14.0 %	98.0 %	53.9 %	54.4 %	97.7 %
8e	$B \times S_f / S_b + (S_f - S_b)$	0 %	0 %	0 %	0 %	0 %

Two key patterns emerge from this table. First, 97.7% of demand in matrix P is predicted by case 8n, which is the standard pivoting formula. Second, 81% of cells have zero values for the base matrix but non-zero values for S_b and S_f .

This example clearly illustrates the issue of *sparsity* in the base matrix. Nearly half of the synthetic demand (46%) occurs in the row corresponding to 4n ($B=0, S_b>0, S_f>0$) where the base matrix is 0. For these cells, changes in the synthetic demand have no impact after pivoting, because as per Table 3 the predicted matrix $P=0$ for case 4n.

In the simple example given in Table 3, there are no observations in the case 8 extreme growth row, but tabulations of this type are also useful in assessing the impact of extreme growth on the percentage of total demand.

If geographical aggregation is used as an approach to overcome sparsity in the base matrices, we would expect to see a shift in the distribution of cells away from cases 1–4 to corresponding cases 5–8.

In summary, the key performance measures when looking at the distribution across the eight cases are:

- the percentage of synthetic demand that occurs in cell types 1–4, where the base matrix is 0, and in particular the percentage of demand that occurs in case 4n, which is typically the most populated of cell types 1–4 and where predicted demand is zero
- of the predicted demand P that occurs within case 8, the proportions that are classified as normal (8n) and extreme (8e)

Other measures may be extracted for particular analyses, for example when analysing growth in new development areas where there are no trips in the base year the percentage of demand in case 2 is relevant.

This chapter assesses the performance of the original process using the performance criteria defined in Chapter 2. Section 3.1 describes the four sets of test matrices that BTS has supplied to allow assessment of the pivoting process. Both zone to zone matrices and station to station matrices have been supplied. The two sets of matrices are distinct in characteristics, and so have been assessed separately. Section 3.2 documents analysis of the zone to zone matrices, whereas Section 3.3 describes analysis of the station to station matrices. The chapter concludes in Section 3.4 with a summary of the issues identified with the original pivoting process.

3.1 Test matrices

BTS supplied four sets of base matrices to allow the operation of the original pivoting process to be analysed. All four sets of matrices have been obtained by running the existing version of the STM for a 2006 base year and a 2036 forecast year. It is emphasised that the existing version of the STM is not the recently re-estimated version of the model. The recently re-estimated version of the model is in the process of being implemented. The test matrices have been supplied using the 2006 travel zones system, which has a total of 2690 zones.

Table 4: Test matrix sets

Set	Purpose(s)	Base matrices	Modes	Zone system
1	Home-work	Rescaled JTW	Car driver, car passenger, train, bus and ferry	2690 zones
2	Home-work	Expanded HTS	Car driver, car passenger, train, bus and ferry	2690 zones
3	All home-based	RailCorp	Train only	343 stations
4	All home-based	HTS rail	Train only	343 stations

Matrix set 1 uses 2006 Census journey to work (JTW) data to define the base matrix, scaled to match Household Travel Survey (HTS) control totals. Scaling is applied to correct for two factors. First, the JTW data are collected on Census day when fewer people are on holiday, whereas the HTS control totals are for an average workday. Second, the JTW data are believed to include a mix of home-based work and some home-based business travel. For both of these reasons, the JTW contains more tours than would be implied by compatibility with HTS, on which the STM is based. For each mode of travel,

a single scaling factor is defined, which is applied to all cells to ensure compatibility with HTS. Matrix set 1 is referred to as simply 'rescaled JTW' in this report.

Matrix set 2 uses the same set of synthetic matrices as matrix set 1, but replaces the base matrices with expanded HTS data. The HTS base matrices are significantly sparser than the JTW matrices, and so comparison of results between matrix sets 1 and 2 provides a test of the impact of sparsity on the pivoting results. Matrix set 2 is also important in demonstrating the performance of the original pivoting process, because for purposes other than home-based commuting expanded HTS data are used to define the base matrices. Matrix set 2 is referred to as simply 'expanded HTS' in this report; it should be compatible with STM.

Matrix set 3 uses the rail operator's (RailCorp) 3.5 hour AM peak matrix as the base matrix. This matrix set allows the operation of the pivoting process to be tested at the station to station level. The synthetic matrices are generated by summing together rail demand for all the home-based purposes, and then applying a factoring approach to adjust from an average weekday (the STM definition) to a 'rail busy day', which is defined as a Tuesday, Wednesday or Thursday in school term time. Thus matrix set 3 contains matrices of all-day demand for a rail-busy day. Matrix set 3 provides another test of the impact of sparsity on the pivoting process, because the number of stations is significantly lower than the number of model zones. This means that the number of cells in the matrix is much lower than in a zone to zone matrix, and so the proportion of cells with zero trips is expected to be much lower. Matrix set 3 is referred to as simply 'RailCorp' in this report.

Matrix set 4 uses expanded HTS data to generate a rail station to station matrix. The matrix generated is substantially sparser than the RailCorp matrix, and so the matrix would not be used in actual applications of the STM. However, tests with this matrix have been undertaken to test the performance of the pivoting process with different input matrices. Matrix set 4 is referred to as simply 'HTS rail' in this report.

3.2 Zonal matrices

This section documents analysis of the performance of the original pivoting process using two sets of zone to zone matrices, namely the rescaled JTW and expanded HTS matrix sets.

3.2.1 Definition of zero trips

As noted in Chapter 1, in the current pivoting process if the number of synthetic trips calculated is less than 0.001 the cell is taken to be 0. Analysis was undertaken using the JTW test matrix set to investigate how much demand is lost from the process by using this definition of zero trips. Table 5 summarises analysis for car driver that investigates the numbers of cells by different ranges of value, and the percentage of demand that lies in each range.

Table 5: Zero trips definition tests, rescaled JTW test matrix set, car driver

Value range	Base matrix B		Synthetic base S_b		Synthetic future S_f		Predicted P	
	% cells	% dem.	% cells	% dem.	% cells	% dem.	% cells	% dem.
0	96.58%	0.00%	1.39%	0.00%	1.03%	0.00%	86.60%	0.00%
0–0.001	0.00%	0.00%	33.24%	0.06%	29.19%	0.04%	0.17%	0.00%
0.001–0.01	0.00%	0.00%	24.21%	0.69%	23.70%	0.49%	5.79%	0.09%
0.01–0.1	0.00%	0.00%	25.27%	6.70%	26.79%	5.14%	2.49%	0.51%
0.1–1	0.28%	1.76%	13.35%	29.35%	15.72%	25.27%	1.41%	3.04%
1 and over	3.13%	98.24%	2.54%	63.20%	3.57%	69.06%	3.54%	96.35%
Total	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%

One-third of synthetic base and almost 30% of synthetic future trips lie in the range 0–0.001, the range where cells are set to 0. However, these cells account for just 0.06% of synthetic base and 0.04% of synthetic future demand. Therefore only tiny fractions of synthetic demand are lost though the current zero test assumption. Larger fractions of synthetic demand are observed in the 0.001–0.01 range and therefore there is no evidence from this analysis to suggest the zero trips cut-off should be increased.

This analysis was repeated for the car passenger mode, where demand is lower and so we'd expect more demand in low value cells. The results are presented in Table 6.

Table 6: Zero trips definition tests, rescaled JTW test matrix set, car passenger

Value range	Base matrix B		Synthetic base S_b		Synthetic future S_f		Predicted P	
	% cells	% dem.	% cells	% dem.	% cells	% dem.	% cells	% dem.
0	99.55%	0.00%	21.75%	0.00%	21.31%	0.00%	95.51%	0.00%
0–0.001	0.00%	0.00%	43.75%	0.42%	43.44%	0.37%	0.06%	0.00%
0.001–0.01	0.00%	0.00%	18.83%	3.89%	19.04%	3.49%	2.91%	0.57%
0.01–0.1	0.00%	0.00%	12.30%	22.08%	12.54%	20.19%	0.79%	1.86%
0.1–1	0.05%	3.11%	3.12%	44.02%	3.39%	43.34%	0.30%	7.34%
1 and over	0.40%	96.89%	0.24%	29.58%	0.28%	32.61%	0.42%	90.24%
Total	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%

Higher fractions of cells and demand lie in the 0–0.001 range for car passenger, however the cells account for less than 0.5% of synthetic demand. More than 3% of synthetic demand lies in the 0.001–0.01 range and therefore it is clear from this analysis that the zero test definition should not be increased.

Similar analysis was undertaken for the rail, ferry and bus modes. The highest fraction of synthetic demand lying in the 0–0.001 range was 0.59% in the synthetic base for bus.

Overall it was judged that the current zero test assumption of 0.001 was reasonable and therefore it has been retained in the enhanced pivoting process.

3.2.2 Base matrix sparsity

Table 7 gives the sparsity indices for each mode for the rescaled JTW and expanded HTS base matrices.

Table 7: Sparsity indices, rescaled JTW and expanded HTS base matrices

Mode	Rescaled JTW	Expanded HTS	Ratio
Car driver	19.1	949.3	49.6
Car passenger	77.4	3770.6	48.7
Rail	42.5	2373.5	55.8
Ferry	9.0	371.4	41.4
Bus	42.2	1848.9	43.8

The rescaled JTW base matrices are much sparser than the synthetic base. For car driver, the percentage of non-zero cells in B is about one-twentieth of the percentage of non-zero

cells in the synthetic base. For car passenger, rail and bus, which have lower shares, the rescaled JTW matrices are sparser still. It is not entirely clear why ferry is the least sparse mode; probably this relates to a relatively low incidence of ferry in the paths for Sb in the wet-rail network, as journeys which use both rail and ferry are considered as rail tours.

The expanded HTS base matrices are even sparser, with the ratio of the indices indicating them to be 50 times sparser than the rescaled JTW. Thus base matrix demand is concentrated over a much lower proportion of cells than the synthetic matrices.

A factor that contributes to these high sparsity indices is the large number of zones in the 2006 zoning system, which means that there is a total of 7,236,100 (O-D) cells in each matrix. Once travel demand has been split by purpose and mode a substantial proportion of these cells would be expected to be zero, even in a fully observed base matrix data like the JTW. For example, for car driver, the mode with the largest share of travel, just 3.4% of cells in the JTW base matrix are non-zero. The synthetic matrices are probabilistic and so demand is distributed over all cells for which that mode is available.

3.2.3 Comparison of synthetic and predicted growth

Table 8 presents a comparison of synthetic and predicted growth for the rescaled JTW and expanded HTS matrix sets.

Table 8: Synthetic and predicted growth, rescaled JTW and expanded HTS base matrix sets

Mode	Synthetic	Predicted: rescaled JTW	Predicted or synthetic	Predicted: expanded HTS	Predicted or synthetic
Car driver	40.5 %	28.4 %	0.70	10.4 %	0.26
Car pass.	12.6 %	6.5 %	0.52	2.7 %	0.22
Rail	86.0 %	40.9 %	0.48	29.1 %	0.34
Ferry	9.4 %	2.4 %	0.26	4.6 %	0.49
Bus	74.3 %	22.9 %	0.31	14.6 %	0.20

The predicted growth when pivoting from the rescaled JTW matrices is consistently lower than the synthetic growth values. For car driver, predicted growth is 70% of the synthetic; for the other modes the predicted growth is 30–50% of the synthetic value. It is noteworthy that the ratio of the difference between predicted and synthetic growth varies considerably between modes.

The predicted growth levels when pivoting from the expanded HTS matrices are even lower than those observed when pivoting from the rescaled JTW matrices for four of the five modes. This results in predicted growth values that are much lower than the synthetic growth values.

Thus a significant issue with the original pivoting process is that the predicted values damp² the growth forecast by the demand models. The effect of this damping is related to sparsity in the base matrices, with the worst correspondence between synthetic and predicted growth observed for the sparsest matrix set (expanded HTS). However, the relationship is less clear at the level of individual modes. For example, car passenger has the highest sparsity index for the rescaled JTW and expanded HTS matrix sets, but a closer

² By damp we mean that we observed lower growth after pivoting relative the synthetic growth.

correspondence between synthetic and predicted growth than the three public transport modes.

Further analysis has been undertaken for the rescaled JTW matrix set to understand better how synthetic and predicted growth differ. The synthetic and pivoted growth figures were compared for the case 8 cells, where Sb , Sf and B are all greater than 0, and in particular for case 8 normal (8n) where the standard pivot formula applies:

$$P = B \frac{Sf}{Sb} \tag{3.1}$$

This comparison allows investigation of how well synthetic and predicted growth match for the case 8n cells where the standard formula applies.

Table 9: Synthetic and predicted growth, case 8 cells, rescaled JTW matrix set

		Synthetic growth	Predicted growth	Ratio
Car driver	Case 8n cells	22.7 %	15.3 %	0.673
	Case 8e cells	174.3 %	25.9 %	0.148
	All cells	40.5 %	28.4 %	0.701
Car passenger	Case 8n cells	4.6 %	-7.7 %	-1.674
	Case 8e cells	86.9 %	8.7 %	0.101
	All cells	12.6 %	6.5 %	0.519
Rail	Case 8n cells	8.0 %	5.3 %	0.656
	Case 8e cells	114.7 %	14.5 %	0.127
	All cells	86.0 %	40.9 %	0.476
Ferry	Case 8n cells	0.5 %	0.2 %	0.343
	Case 8e cells	45.2 %	8.7 %	0.193
	All cells	9.4 %	2.4 %	0.259
Bus	Case 8n cells	10.3 %	-2.6 %	-0.253
	Case 8e cells	85.9 %	9.1 %	0.106
	All cells	74.3 %	22.9 %	0.309

These results (Table 9) demonstrate that synthetic growth is lower than pivoted growth for case 8n cells, as well as for all cells. So even when the standard formula is applied, the pivoting procedure is damping the synthetic growth. **This means that differences between synthetic and predicted growth are occurring because of differences in the distribution of B and Sb over cells.**

The results for case 8n for car passenger and bus also illustrate the sign change problem discussed in Section 3.2, where synthetic growth is positive but predicted growth is negative. Again, the sign change problem occurs because of differences in the distribution of B and Sb over cells.

3.2.4 Distribution of demand across the eight cases

Analysis of the distribution of demand across the eight cases used in the pivoting procedure has been undertaken for each of the five travel modes in the rescaled JTW and expanded HTS matrix sets. Below, the distributions for the car driver mode are presented and discussed. Table 10 presents the comparison for car driver for the rescaled JTW matrix set.

Table 10: Distribution of car driver demand, rescaled JTW matrix set

Case	B	Sb	Sf	B > 0	Sb > 0	Sf > 0	P > 0
1	0	0	0	0.0 %	0.0 %	0.0 %	0.0 %
2	0	0	> 0	0.0 %	0.0 %	0.3 %	0.3 %
3	0	> 0	0	0.0 %	0.0 %	0.0 %	0.0 %
4n	0	> 0	> 0	0.0 %	43.9 %	38.8 %	0.0 %
4e	0	> 0	> 0	0.0 %	0.9 %	11.0 %	8.6 %
5	> 0	0	0	0.2 %	0.0 %	0.0 %	0.1 %
6	> 0	0	> 0	0.0 %	0.0 %	0.0 %	0.0 %
7	> 0	> 0	0	0.0 %	0.0 %	0.0 %	0.0 %
8n	> 0	> 0	> 0	82.9 %	53.6 %	46.8 %	74.5 %
8e	> 0	> 0	> 0	16.8 %	1.6 %	3.1 %	16.5 %
Total				100.0 %	100.0 %	100.0 %	100.0 %

Almost 83% of the distribution of the base matrix B occurs under case 8n, the standard pivot formula. However, 44% of Sb demand, and 39% of Sf demand, occur in case 4n where no demand is predicted ($P=0$). This result further illustrates the impact of base matrix sparsity, with close to half of synthetic demand occurring in cells where the base matrix is 0.

Table 11 presents the same comparison for car driver for the expanded HTS matrix set.

Table 11: Distribution of car driver demand, expanded HTS matrix set

Case	B	Sb	Sf	B > 0	Sb > 0	Sf > 0	P > 0
1	0	0	0	0.0 %	0.0 %	0.0 %	0.0 %
2	0	0	> 0	0.0 %	0.0 %	0.3 %	0.4 %
3	0	> 0	0	0.0 %	0.0 %	0.0 %	0.0 %
4n	0	> 0	> 0	0.0 %	95.8 %	84.5 %	0.0 %
4e	0	> 0	> 0	0.0 %	1.2 %	12.4 %	10.8 %
5	> 0	0	0	0.1 %	0.0 %	0.0 %	0.1 %
6	> 0	0	> 0	0.1 %	0.0 %	0.0 %	0.1 %
7	> 0	> 0	0	0.0 %	0.0 %	0.0 %	0.0 %
8n	> 0	> 0	> 0	28.4 %	1.6 %	1.3 %	23.1 %
8e	> 0	> 0	> 0	71.4 %	1.5 %	1.5 %	65.7 %
Total				100.0 %	100.0 %	100.0 %	100.0 %

With the expanded HTS matrix set, where the base matrices are much sparser, over 95% of synthetic demand occurs in cells 1–4 where the base matrix is 0. Most of this demand is in case 4n, where no demand is predicted after pivoting.

A further issue is that for those cells that lie in case 8 (where B , Sb and Sf are all greater than 0) the larger contribution to demand comes from the case 8e extreme growth formula. In this case, the pivoting formula deliberately predicts lower growth than the standard formula to prevent the predicted number of trips exploding when Sf is much greater than Sb . However, with these sparse matrices, the trigger point for using the extreme growth rule, $Sf > X_2$, is being triggered too frequently:

$$X_2 = Sb.k_1 + Sb.k_2 \cdot \max \left[\frac{Sb}{B}, \frac{k_1}{k_2} \right] \tag{3.2}$$

The reason for this result is that with very sparse matrices Sb/B is frequently smaller than the ratio $k_1/k_2=0.1$, so that $X_2=Sb$, and in most cases $Sf>Sb$. **This result demonstrates that an issue with the original pivoting procedure is that the switch point for extreme growth X_2 does not work well when the base matrices are sparse relative to the synthetic matrices.**

3.3 Station matrices

This section describes analysis of the performance of the existing pivoting process using the two sets of station to station matrices, namely the RailCorp and HTS rail matrices.

3.3.1 Base matrix sparsity

The sparsity indices for the two station to station base matrices are presented in Table 12.

Table 12: Sparsity indices, station matrices

Matrix set	Sparsity index
RailCorp	3.75
HTS rail	65.55

The sparsity indices for the RailCorp matrix set, the station to station matrix set used for pivoting when the STM is applied, are relatively low. The HTS rail matrices are considerably sparser.

3.3.2 Comparison of synthetic and predicted growth

Table 13 compares synthetic and predicted growth for the two station to station matrix sets.

Table 13: Comparison of synthetic and pivoted growth, station to station matrices

Matrix set	Synthetic	Predicted	Predicted and synthetic
RailCorp	78.2 %	67.7 %	0.87
HTS rail	78.2 %	46.9 %	0.60

For the less sparse RailCorp matrix set, predicted and synthetic growth match reasonably well. For the sparser HTS rail matrix set, the predicted growth significantly under-predicts the synthetic.

3.3.3 Distribution of demand across the eight cases

Table 14 presents the distribution of demand across the eight cases for the RailCorp matrix set.

Table 14: Distribution of rail demand, RailCorp matrix set

Case	B	Sb	Sf	B > 0	Sb > 0	Sf > 0	P > 0
1	0	0	0	0.0 %	0.0 %	0.0 %	0.0 %
2	0	0	> 0	0.0 %	0.0 %	25.0 %	24.7 %
3	0	> 0	0	0.0 %	0.2 %	0.0 %	0.0 %
4n	0	> 0	> 0	0.0 %	6.3 %	5.0 %	0.0 %
4e	0	> 0	> 0	0.0 %	0.5 %	2.9 %	1.5 %
5	> 0	0	0	1.2 %	0.0 %	0.0 %	0.7 %
6	> 0	0	> 0	2.3 %	0.0 %	2.9 %	4.3 %
7	> 0	> 0	0	9.3 %	8.4 %	0.0 %	0.0 %
8n	> 0	> 0	> 0	70.9 %	81.2 %	54.8 %	48.6 %
8e	> 0	> 0	> 0	16.2 %	3.5 %	9.5 %	20.2 %
Total				100.0 %	100.0 %	100.0 %	100.0 %

The percentage of synthetic base demand that occurs in cases 1–4 where the base matrix is 0 is relatively low at 7.0%. A quarter of synthetic future demand occurs in case 2; however, for this case demand is predicted after pivoting ($P=Sf$) so this growth is not lost.

Focusing on those cells in case 8, it can be seen that most demand is predicted by case 8n, where the standard pivoting formula is applied.

In summary, for the RailCorp station to station matrix the distribution of demand over the eight cases is good, with low percentages of synthetic base matrix demand occurring in cases 1–4 where the base matrix is 0, and most case 8 growth being predicted by the normal growth regime.

3.4 Summary of issues with the original pivoting process

3.4.1 Zonal matrices

The analysis of the operation of the original pivoting process for the zone to zone matrices has identified two closely inter-related issues. First, the base matrices are sparse at the 2690 zone level:

- Even for the fully observed JTW matrices, the high number of cells in the matrices leads to sparsity.
- When expanded HTS data are used, the approach used for pivoting all purposes apart from commute, the base matrices are very sparse indeed.
- As a result of this sparsity, substantial fractions of synthetic demand occur in cases where the base matrix is 0, and as a consequence a lot of synthetic demand growth has no impact on the predicted demands.
- Even for cells where the base matrix is defined, the mismatch between the mean size of non-zero base and synthetic cells means that too much growth is classified as ‘extreme’ when the expanded HTS data are used.

Second, synthetic and predicted growth do not match well:

- Predicted growth is consistently lower than synthetic.

- Even for the cells where the standard pivoting formula applies, differences between synthetic and pivoted growth exist.
- These differences get worse as base matrix sparsity increases.
- These differences vary between modes.

To address these issues, three enhancements to the pivoting process were proposed which are set out in detail in Chapter 4:

- pivoting at an aggregate level, which will reduce the sparsity of the base matrices, and therefore lead to a better correspondence between synthetic and predicted growth
- revisiting the definition of the switch-point for extreme growth
- ‘normalising’ at some level, with the normalisation ensuring that synthetic and predicted growth match exactly at the level at which the normalisation is applied (e.g. total growth for a mode, growth for a mode and origin combination)

3.4.2 Station matrices

The original pivoting process performs better for the RailCorp station to station matrix set, where the base matrices are significantly less sparse than the zonal matrices. Synthetic and predicted growth match fairly well, and the distribution of demand over the eight cases is good. The performance is worse for the sparser HTS Rail matrix set, but this matrix set is not used when the pivoting process is applied in the STM.

Given the issues with the performance of the original pivoting process discussed in Chapter 3, three enhancements to the pivoting process were suggested: pivoting at a more aggregate level, revising the formula that controls the switch point between normal and extreme growth, and applying some form of normalisation so that the growth in trips predicted by the demand models – the synthetic growth – is matched by the growth in trips predicted by the pivoting process. Sections 4.1 to 4.3 set out each of these three performance improvements in more detail.

4.1 **Zonal aggregation**

The detailed 2690 zone system used in the new STM results in sparse base matrices, even for the fully observed Census JTW data. This sparsity results in the original pivoting process performing poorly; specifically synthetic growth is systematically under-predicted.

Therefore, it was proposed to undertake the pivot at a more aggregate zonal level. This approach was expected to improve the performance of the pivoting process for two reasons. First, after zonal aggregation the percentage of zero cells in the base matrix will be significantly lower. This shifts synthetic demand from case 4 normal, where the predicted demand is zero, to case 8 normal, where positive demand is always predicted. Second, after zonal aggregation the mean size of non-zero values of Sb and B would be expected to be closer. This is expected to improve the correspondence between synthetic and pivoted growth for case 8 normal cells.

4.1.1 **Initial tests**

To assess the performance improvement zonal aggregation offered, BTS made initial tests where the test matrices were aggregated from the detailed 2690 zones to the 80 statistical local area (SLA) zones used in the 2006 zoning system, and then the pivoting process was applied at the 80 SLA zone level. Table 15 summarises the performance improvement observed for the car driver mode for the rescaled JTW and expanded HTS test matrix sets.

In addition to the sparsity index, Table 15 summarises the percentage of synthetic base matrix demand that occurs for pivoting cases 1–4 where the base matrix demand is zero, and compares the ratio between predicted and synthetic growth, for all cells and for case 8 normal cells where the standard factor pivoting formula applies ($P=B.Sf/Sb$).

Table 15: Impact of zonal aggregation on car driver predictions

		Sparsity index	% Sb cases 1–4	Growth ratio, predicted / synthetic	
				All cells	Case 8 normal cells
JTW test matrix set	2690 zones	19.1	44.8 %	0.70	0.67
	80 SLA zones	1.25	0.12%	0.98	0.97
HTS test matrix set	2690 zones	949.3	97.0 %	0.26	-0.61
	80 SLA zones	4.06	15.0 %	0.95	0.94

These initial tests demonstrate significant performance improvements from zonal aggregation. At the more aggregate SLA level, base matrix sparsity is much lower, there are significant reductions in the percentage of *Sb* that occur in cases 1 to 4, and the correspondence between predicted and synthetic growth is much improved. Furthermore, the sign change issue observed for the case 8 normal cells when pivoting at the 2690 zone level is overcome.

Table 15 demonstrates that zonal aggregation works very well for car driver. To assess the performance improvement aggregation offers across all the modes, Table 16 summarises the ratios of predicted and synthetic growth for each of the five modes.

Table 16: Impact of zonal aggregation on predicted demand: synthetic growth ratios, all modes

		Car driver	Car pass.	Rail	Ferry	Bus
JTW test matrix set	2690 zones	0.70	0.52	0.48	0.26	0.31
	80 SLA zones	0.98	0.87	0.85	-0.38	0.60
HTS test matrix set	2690 zones	0.26	0.22	0.34	0.50	0.20
	80 SLA zones	0.95	0.79	0.58	-1.21	0.37

With the exception of ferry, the correspondence between synthetic and predicted growth is significantly improved by zonal aggregation, although there is still a tendency for predicted growth to fall below synthetic, particularly for bus. For ferry, the results are worse with zonal aggregation, which may follow from the fact that ferry is only used for a few limited origin-destination (OD) movements within an SLA-SLA matrix.

Overall it was concluded that zonal aggregation significantly improves the performance of the pivoting process, and so should form part of the enhanced pivoting approach.

4.1.2 Selecting the appropriate level of aggregation

The 80 SLA zones defined in the 2006 zoning system performed well in these initial tests, and was therefore a candidate level of aggregation for the enhanced pivoting approach. Before deciding for certain that SLA level was the appropriate level of zonal aggregation it was important to check that the maximum population and employment per SLA to ensure that there are not areas where too high a fraction of demand is concentrated. Table 17 summarises key headline statistics at the SLA zone level, which can be compared to the equivalent figures at the detailed 2690 zone level, which are presented in Table 18.

Table 17: SLA zone characteristics

	Population	Employment	Zonal area (km ²)
Mean	65,178	30,838	306
Standard deviation	27,045	32,008	727
Coefficient of variation	0.415	1.038	2.381
Minimum	13,746	4527	4
Maximum	139,163	274,473	4322
Total	5,214,203	2,467,000	24,444

For population, there is reasonable variation in total population between SLA, with a coefficient of variation of 0.415. The maximum population is 139,163 which represents 2.7% of the total population of the study area. For employment, the coefficient of variation shows that there is more variation between SLA. This is expected as employment tends to be concentrated in centres. The SLA with the highest employment is Inner Sydney, which accounts for 11.1% of total employment. This is also the area with the smallest area, just 4 square km.

Table 18: Detailed 2690 zone characteristics

	Population	Employment	Zonal area (km ²)
Mean	1938	917	9.1
Standard deviation	1412	1400	58.8
Coefficient of variation	0.729	1.527	6.474
Minimum	0	0	0.003
Maximum	7563	16,723	1202
Total	5,214,203	2,467,000	24,444

At the detailed zone level there is greater variation in population, employment and zonal area between zones relative to the 80 SLA zone level, shown by the higher coefficients of variation relative to those presented in Table 17. In application of the STM, zones are set to have a minimum of ten households and minimum employment of ten persons. This approach means that all cells can have $S_b > 0$ and $S_f > 0$, though for public transport modes it is also necessary to have a path through the network for the cell in question.

Overall, the maximum population and employment numbers per SLA zone were judged to be acceptable; in particular, the area defined by the Inner Sydney SLA is small and so does not account for too high a fraction of total employment. Therefore, the SLA zone system was adopted as the zone system for the aggregate pivot procedure.

During 2012, a new 2011 travel zone system will be adopted. The number of travel zones will only change slightly from the current total of 2690. However, the SLA level of geography will be replaced with a hierarchy of areas, with Statistical Area 1 (SA1), Statistical Area 2 (SA2), etc. defined. There will be around 350 SA2 areas, which implies matrices with 112,500 cells compared with 6400 in the SLA zone system. On the basis of the analysis in this report, we believe that using the more detailed SA2 zone system would significantly reduce the benefits of zonal aggregation as the aggregate matrices would be sparser. Further analysis will be required when matrices are available defined with the new 2011 zoning system, but ideally the aggregate pivot procedure would be applied using aggregate zones of comparable size to the 80 SLAs defined with the 2006 zoning system.

4.1.3 Disaggregating demand back to detailed zone level

In order to implement pivoting at the aggregate level, a procedure was required to distribute the predicted demand at the aggregate level back down to the detailed 2690 zone level. The best available information for this disaggregation process taking into account future travel demand is the synthetic future matrix Sf , except for a special case where the base matrix B has to be used instead. The following formulae set out how this process operates.

Let B^A , Sb^A , Sf^A be the base matrix, synthetic base matrix and synthetic future matrix at an aggregate level A , and B_a , Sb_a , Sf_a be the corresponding matrices at a disaggregate level a , where:

$$B^A = \sum_a B_a \quad (4.1)$$

$$Sb^A = \sum_a Sb_a \quad (4.2)$$

$$Sf^A = \sum_a Sf_a \quad (4.3)$$

P^A , which is the aggregate predicted matrix calculated, is a function of the aggregate base matrix B^A , aggregate synthetic base matrix Sb^A and aggregate synthetic future matrix Sf^A . P^A is obtained by applying the eight-case method taking into account the revised definition of the switch point between normal and extreme growth discussed in Section 4.2.

To determine the predicted matrix at the disaggregate level, two disaggregation formulae are required. The default is to distribute demand to the detailed zonal level using the distribution of Sf :

$$P_a = P^A * \frac{Sf_a}{Sf^A} = Sf_a * \frac{B^A}{Sb^A} \quad (4.4)$$

where: P_a is the predicted matrix at disaggregate level, which is subsequently normalised as is laid out in Section 4.3 to obtain the final pivoted matrix.

It is clear from Equation (4.4) above that adding up P_a over the aggregated cells will give P^A . Moreover, the adjustment that is made relative to the standard disaggregate pivoting

process is simply to replace the factor $\frac{B_a}{Sb_a}$, which is unreliable because B_a is often 0 and for the low fraction of cells where B_a is not zero it is typically much larger than Sb , by $\frac{B^A}{Sb^A}$, which is a more reliable value.³

For aggregate pivot case 5 ($B^A > 0$, $Sb^A = 0$, $Sf^A = 0$) the default formulation given in Equation (4.4) cannot be used, as $Sf^A = 0$ but pivoted demand at the aggregate level is non-zero, and

³ At the disaggregate level, the standard pivoting formula for case 8 normal can be written as $P_a = Sf_a \cdot B_a / Sb_a$. By moving to aggregate pivoting this formula is modified to $P_a = Sf_a \cdot B^A / Sb^A$.

therefore an alternative formulation is required. In this instance, the approach is to simply take the distribution of demand in the base matrix:

$$P_a = B_a \tag{4.5}$$

Note that the disaggregation procedure takes place at the very end of the pivoting process, after the two normalisation steps set out in Section 4.3 have been applied.

ALOGIT code to implement the disaggregation procedure has been written and the procedure has been tested by comparing the matrix totals by mode before and after disaggregation.

4.2 Modification to the extreme growth switch point

4.2.1 Why the case 8 extreme growth rule was modified

The formulation that determines whether growth is normal or extreme was classifying too much case 8 growth as ‘extreme’, particularly where the base matrices are sparse, contributing to the under-prediction of growth. As demonstrated in Table 15, zonal aggregation reduces this problem. Nonetheless, an improved formulation for extreme growth was required that was able to work with a range of different base matrices with different expansion factors.

For case 8, the key ratio is Sb/B , which is compared with k_1/k_2 to determine how the switch point for ‘extreme growth’ pivoting is calculated. This relates closely to the sparsity ratio. Suppose we assume that $E(B)=Sb$, which is basically saying that the model is good, and that B is sampled at a rate R . Then the probability of observing a trip in a cell is Sb/R , which is a small number in most cases, so that observing two trips is rare.⁴ If a trip is sampled, B would then have R trips in that cell. Then we can calculate, for cells with $Sb>0$:

$$\text{Probability of positive cell in } B = Sb/R \approx Sb/B$$

Assuming that if $Sb=0$ there is no probability of $B>0$, the sparsity ratio for the matrix would then also be equal to this.

The sparsity indices are significantly reduced by zonal aggregation. However, indices greater than 10 still occur for car passenger and public transport modes with the HTS test matrix set, as shown in Table 19.

Table 19: Sparsity indices after aggregation to 80 SLA zone level

	Car driver	Car pass.	Rail	Ferry	Bus
JTW test matrix set	1.25	1.74	2.0	1.7	2.8
HTS test matrix set	4.1	13.3	12.5	15.5	17.0

For cases where the sparsity index is greater than 10, it is reasonable to conclude that in most cells the ratio $Sb/B < 0.1$.

⁴ In fact, in the JTW matrices small values in the range 0 to 1 are rounded to either 0 or 3, so this discussion is not exactly applicable.

The original pivoting process defined the switch point between normal and extreme growth for case 8 as follows, with the extreme growth rule applied if $Sf > X_2$:

$$X_2 = Sb.k_1 + Sb.k_2 \cdot \max \left[\frac{Sb}{B}, \frac{k_1}{k_2} \right] \quad (4.6)$$

with $k_1=0.5$ and $k_2=5$.

This gives a ratio for $k_1/k_2=0.1$. If $Sb/B < 0.1$, then the second term in the max is operative and this formula reduces to $X_2=Sb$, which is inappropriate because it means that the extreme growth rule is applied where $Sf > Sb$, which as a result of growth in population and employment would be expected for nearly all cells. When a high fraction of case 8 cells are classified as extreme, predicted growth is lower than synthetic growth.

Table 19 illustrates that the STM pivoting process has to operate with base matrices with a wide range of different sparsity values, and the discussion above illustrates that the relationship of k_2 to the ratio Sb/B is problematic. Therefore the formula for the switch point between normal and extreme growth X_2 was simplified to:

$$X_2 = Sb.k_2 \quad (4.7)$$

This gives a switch point between normal and extreme growth that is able to work with matrices at a range of different sparsity values, consistent with BTS's objective for this work of developing an enhanced pivoting process that is able to operate with a range of different base matrices.

4.2.2 Selection of switch point parameter values

A number of tests were made to determine the appropriate value for k_2 in the simplified formulation given in Equation (4.7) for case 8 pivoting. A relevant consideration is that for case 4 the switch point for extreme growth is formulated as:

$$X_1 = Sb.k_2 \quad (4.8)$$

with $k_2 = 5$.

Thus consistency with the case 4 definition of extreme growth is an advantage of moving to the simpler formulation for case 8 given in Equation (4.7).

Three sets of tests were run to investigate different switch point parameter values:

- simple form 1: $k_2 = 5$ for cases 4 and 8, ensuring consistency in the k_2 values between cases 4 and 8, and maintaining the factor of $5.Sb$ as the switch point between normal and extreme growth
- simple form 2: $k_2 = 3$ for cases 4 and 8, ensuring consistency between cases 4 and 8, but defining extreme growth at a lower level
- simple form 3: $k_2 = 3$ for case 4, but $k_2 = 5$ for case 8

The third approach aims to maximise the correspondence between synthetic and pivoted growth because for case 4 normal growth the pivoting rule is $P=0$, so defining more cases as 'extreme' increases the predicted demand, whereas for case 8 moving cells from normal to extreme cases reduces predicted demand. This approach introduces an inconsistency in the definition of extreme growth between cases 4 and 8.

The results were compared to the rules used in the original process, where for case 8 the complex switch point formulation given in Equations (1.1) and (1.2) is used. All tests were undertaken pivoting at the aggregate 80 SLA zone level, as pivoting takes place at this level in the enhanced pivoting process.

The sets of tests were run for the rescaled JTW and expanded HTS base matrices. Table 20 summarises the correspondence between synthetic and predicted growth obtained for the rescaled JTW matrices in the results.

Table 20: k_2 tests, rescaled JTW matrices

	Growth ratio, predicted : synthetic			
	Complex formulation $k_1=0.5$ $k_2=5$	Simple form 1: $k_2=5$ case 4 $k_2=5$ case 8	Simple form 2: $k_2=3$ case 4 $k_2=3$ case 8	Simple form 3: $k_2=3$ case 4 $k_2=5$ case 8
Car driver	0.976	0.997	0.993	0.997
Car passenger	0.871	0.935	0.968	0.936
Train	0.854	0.940	0.925	0.953
Bus	0.603	0.677	0.686	0.691
Ferry	-0.379	-0.187	-0.189	-0.178

All three combinations of the simplified formulation give clear performance improvements (growth ratios closer to 1) relative to the complex formulation. Comparing simple forms 1 and 2, simple form 1 performs better for three of the five modes. Simple form 3 offers a slight improvement in performance relative to simple form 1.

Table 21 presents the same set of comparisons for the sparser expanded HTS matrices.

Table 21: k_2 tests, expanded HTS matrices

	Growth ratio, predicted / synthetic			
	Complex formulation	Simple form 1: $k_2=5$ case 4 $k_2=5$ case 8	Simple form 2: $k_2=3$ case 4 $k_2=3$ case 8	Simple form 3: $k_2=3$ case 4 $k_2=5$ case 8
Car driver	0.948	1.038	1.019	1.046
Car passenger	0.792	1.013	1.073	1.077
Train	0.584	0.892	0.899	0.959
Bus	0.373	0.600	0.635	0.660
Ferry	-1.214	-0.716	-0.677	-0.677

More substantial improvements in performance are observed in moving from complex to simple forms with the sparser HTS matrices, in particular for train and bus modes. Simple form 2 performs better than simple form 1 in four of the five modes, though the difference in performance is not large. Simple form 3 performs best for train and bus, but increases the over-prediction of growth for car driver and car passenger.

Following these tests, it was decided to adopt simple form 1 in the enhanced pivoting process, with $k_2=5$ for both cases 4 and 8. Moving to a lower value of $k_2=3$ in simple form 2 gave mixed results relative to simple form 1, and while simple form 3 with mixed k_2 values gives the best results overall, the improvements relative to simple form 1 are relatively small and this approach introduces a discontinuity between the case 4 and 8 switch points.

4.3 Normalisation

Both zonal aggregation and modifying the switch point for normal and extreme growth are expected to improve the performance of the pivoting procedure. However, they will not guarantee that synthetic and pivoted growth match exactly. Furthermore, sign changes may still occur in particular cases where (for example) small positive synthetic growth results in small negative growth after pivoting, and such results are difficult to explain to users of the STM.

Normalisation means applying a procedure so that synthetic and pivoted growth match exactly at the level at which normalisation is applied, which might be growth in total trips, or growth in trips by a given mode.

Daly, Fox and Patrui (2011) describe a mathematical procedure that allows the pivoted and synthetic growth to match exactly at each level represented in the choice structure, rather than at only one level in the choice structure. The detailed thinking is complex, and is presented in Appendix A, but in summary there are two approaches which could be followed:

- implementing a utility-based correction, which requires the travel demand models to be modified, i.e. to calculate a set of K-factors for each mode and destination, so that base and synthetic base trips match exactly
- implementing a ‘top-down normalisation’ so that synthetic and predicted growth are consistent for each level in the choice structure

Both these approaches would take account of the complex and varied tree structures used in the STM, which include toll road choice and park-and-ride nests for a number of the purposes. The two approaches would give identical answers, and in both cases would ensure that base and synthetic base trips match exactly at each level in the model structure for a given origin zone.

However, implementation of either of these procedures requires significant additional programming, which was not possible within the resources available for this project. Therefore, a simpler normalisation consistent with the thinking in the original proposal was specified.

4.3.1 Mode-origin normalisation

The first normalisation approach that has been specified is to apply a mode normalisation for each aggregate SLA origin zone. This step is termed ‘mode-origin normalisation’, and the normalisation factors are calculated as follows:

$$MOF = \left(\frac{\sum_d B}{\sum_d P} \right) * \left(\frac{\sum_d Sf}{\sum_d Sb} \right) \quad (4.8)$$

where: MOF is the mode-origin normalisation factor for the origin zone

$\sum_d B$ is the sum over destinations of B for the origin zone

$\sum_d P$ is the sum over destinations of pivoted demand P before normalisation for the origin zone

$\sum_d Sf$ is the sum over destinations of Sf for the origin zone

$\sum_d Sb$ is the sum over destinations of Sb for the origin zone

If one or more of these summations is 0 for a given origin zone, then the mode normalisation factor is set to a value of 1.

Then the mode normalisation factor is applied to each cell for the origin zone in question:

$$P' = MOF . P \tag{4.9}$$

where: P' is the predicted demand after mode-origin normalisation

The mode normalisation process has been applied at the SLA origin zone level, i.e. at the aggregate zonal level at which pivoting has been implemented in the enhanced process.

4.3.2 Overall mode normalisation

This mode-origin approach ensures that synthetic and predicted growth in tours by each mode are matched exactly for each origin zone. However, as will be seen in Chapter 5, this does not ensure that synthetic and pivoted growth in tours match when demand is summed over the whole matrix. Therefore, a second normalisation was applied after the mode-origin normalisation, termed ‘overall mode normalisation’. The normalisation factors are calculated as follows:

$$MF = \left(\frac{\sum_o \sum_d B}{\sum_o \sum_d P'} \right) * \left(\frac{\sum_o \sum_d Sf}{\sum_o \sum_d Sb} \right) \tag{4.10}$$

where: MF is the overall mode normalisation factor

$\sum_o \sum_d B$ is the total base matrix demand

$\sum_o \sum_d P'$ is the total pivoted demand after mode-origin normalisation

$\sum_o \sum_d Sf$ is the total synthetic future demand

$\sum_o \sum_d Sb$ is the total synthetic base demand

Then the mode normalisation factor is applied to each cell in the matrix as follows:

$$P^* = MF . P' \tag{4.11}$$

where: P^* is the predicted demand after overall mode normalisation

It is emphasised that when a multi-stage normalisation is applied, only the final normalisation is guaranteed to hold. So in the enhanced procedure, the overall mode normalisation is guaranteed by the final step, but mode normalisation at the origin level is not. The further the overall mode normalisation factor MF is from 1, the further the overall mode normalisation takes the results from ensuring the results are normalised at the mode-origin level.

CHAPTER 5 **Performance of enhanced pivoting process**

To assess the performance of the enhanced pivoting process, the analysis of the original pivoting process presented in Chapter 3 has been repeated. For the enhanced pivoting process, the performance measures have been calculated after each of the three performance improvements is applied so that the relative impact of each improvement can be assessed. Section 5.1 presents the analysis of the performance of the enhanced process applied to the zone to zone matrices, and Section 5.2 presents analysis of the enhanced process applied to the station to station matrices.

5.1 **Zone to zone matrices**

5.1.1 **Base matrix sparsity**

Table 22 summarises the impact of zonal aggregation on the sparsity indices for the matrix set that uses the rescaled JTW base matrices. It should be noted that with 2690 zones there are 7.2 million matrix cells, whereas with 80 zones there are only 6400 matrix cells, and therefore zonal aggregation reduces the number of matrix cells by a factor of just over 1100.

Table 22: Sparsity indices, rescaled JTW base matrices

Mode	Original process: pivoting at 2690 zone level	Enhanced process: pivoting at 80 SLA zone level	Ratio
Car driver	19.1	1.25	15.3
Car passenger	77.4	1.74	44.4
Rail	42.5	2.03	20.9
Ferry	9.0	1.69	5.3
Bus	42.2	2.75	15.4

As would be expected, zonal aggregation has significantly reduced the sparsity of the matrices, with a more than ten-fold reduction in sparsity for all modes except ferry. Following zonal aggregation the levels of sparsity are comparable between the base and synthetic base matrices, i.e. the sparsity indices are no higher than 3.

Table 23 summarises the impact of zonal aggregation on the sparsity indices for the matrix set that uses the expanded HTS base matrices.

Table 23: Sparsity indices, expanded HTS base matrices

Mode	Original process: pivoting at 2690 zone level	Enhanced process: pivoting at 80 SLA zone level	Ratio
Car driver	949.3	4.06	233.6
Car passenger	3770.6	13.3	284.4
Rail	2373.5	12.5	189.9
Ferry	371.4	15.5	24.0
Bus	1848.9	17.0	108.8

The expanded HTS base matrices are considerably sparser than the JTW matrices at the 2690 zone level. As a result, the reduction in sparsity that follows from zonal aggregation is more dramatic than that observed for the JTW base matrix set, with a more than 100-fold reduction in sparsity observed for all modes except ferry. Nonetheless, for all modes except car driver the base matrices remain more than ten times as sparse as the synthetic matrices after zonal aggregation.

5.1.2 Comparison of synthetic and predicted growth

Table 24 presents a comparison of the ratio between predicted and synthetic growth for each mode for the rescaled JTW base matrix set. The table shows the impact of each of the three performance improvements on the predicted to synthetic growth ratio. The three performance improvements (zonal aggregation, revised switch point, normalisation) are applied cumulatively, so for example the revised switch point is applied in addition to zonal aggregation, and so the table allows the relative impact of each performance improvement to be assessed.

Table 24: Synthetic and predicted growth comparisons, rescaled JTW base matrix set

Mode	Synthetic growth	Original pivoting process	Predicted / synthetic growth ratios			
			Zonal aggregation	Revised switch point	Normalisation, step 1: origin mode	Normalisation, step 2: total mode
Car driver	40.5 %	0.70	0.98	1.00	1.02	1.00
Car pass.	12.6 %	0.52	0.87	0.93	1.00	1.00
Rail	86.0 %	0.48	0.85	0.94	0.96	1.00
Ferry	9.4 %	0.26	-0.37	-0.19	0.09	1.00
Bus	74.3 %	0.31	0.60	0.68	0.72	1.00

Of the various measures, zonal aggregation results in the largest improvement in the correspondence between synthetic and predicted growth (with the exception of ferry, where positive synthetic growth is translated into negative predicted growth).

The revised switch point leads to a further performance improvement, i.e. predicted/synthetic growth ratios move closer to 1 (predicted growth for ferry still has the wrong sign but the growth ratio does move closer to 1). This performance improvement comes about because more demand is classified as case 8 normal where the standard factor pivoting formula is applied. At this stage, synthetic and pivoted growth match fairly well for car driver, car passenger and rail; ferry remains problematic, however, with a sign change issue; and for bus synthetic growth remains significantly under-predicted (this should be improved when the new STM is implemented).

The first normalisation step, mode normalisation for each origin, results in a small performance improvement for car passenger, rail and bus. Furthermore, at the SLA level,

the application of origin-mode normalisation has ensured that the percentage growth in demand by mode as predicted by the travel demand models is exactly matched after pivoting.⁵ However, as illustrated in Table 24, this condition does not ensure that synthetic and pivoted growth rates at the whole matrix level match exactly after pivoting, and in particular for bus synthetic growth remains significantly under-predicted.

In order to understand better the results that have been obtained, the results for bus were analysed further. This analysis revealed that at the 80 SLA origin level at which the mode normalisation factors are defined, there was a strong negative correlation of -0.39 between base matrix demand ($\Sigma_a B$) and synthetic growth ($\Sigma_a Sfl \Sigma_a Sb$). This means that higher synthetic growth occurs for SLA origins associated with lower base matrix demand, and as a result the percentage growth in trips after pivoting is lower than predicted by the synthetic matrices. This result follows from significant differences in the distributions of demand over the 80 SLA origin zones between B and Sb .

Given the significant difference between synthetic and predicted growth for bus, it was decided to implement a second normalisation step, an overall mode normalisation, so that synthetic and predicted growth match exactly for each mode. The overall mode normalisation is applied by calculating a single normalisation factor for each mode, applied across all origins and destinations for that mode, i.e. at the whole matrix level. For car driver, car passenger and rail the impact of this step is small – i.e. normalisation factors close to 1 are applied – as synthetic and pivoted growth matched well after origin-mode normalisation. However, for bus a factor of 1.39 is required, which means that predicted growth at the origin level will be significantly higher than predicted by the travel demand models. When the new version of the STM is operational it is recommended that analysis is undertaken to compare the distributions of B and Sb for bus to better understand any differences, and it is hoped that with a better correspondence between the two the overall mode normalisation factor required for bus will be substantially closer to 1.

Pivoting is problematic for ferry. More than three times as many trips exist in the ferry base matrices as in the synthetic base, so the correspondence between B and Sb is poor, and zonal aggregation does not work well for this mode. BTS is aware of issues with the current base matrices for ferry, and given these issues and the low mode share for ferry it was agreed with BTS that it was not worthwhile developing a separate pivoting procedure just for ferry. Given these results BTS may choose not to use the pivoted results for ferry when it applies the model.

Table 25 presents the same set of comparisons for the expanded HTS base matrix set.

⁵ For a given SLA origin, the normalisation can only be calculated if there is demand in the base matrix and the synthetic base matrix to at least one destination zone. If there is no base matrix demand and/or no synthetic base matrix demand in any destination zone then the normalisation factor is set to 1.

Table 25: Synthetic and predicted growth comparisons, expanded HTS base matrix set

Mode	Synthetic growth	Predicted / synthetic growth ratios				
		Original pivoting process	Zonal aggregation	Revised switch point	Normalisation, step 1: origin mode	Normalisation, step 2: total mode
Car driver	40.5 %	0.26	0.95	1.04	1.03	1.00
Car pass.	12.6 %	0.22	0.79	1.01	0.87	1.00
Rail	86.0 %	0.34	0.58	0.89	0.92	1.00
Ferry	9.4 %	0.49	-1.21	-0.41	-0.12	1.00
Bus	74.3 %	0.20	0.37	0.60	0.69	1.00

Consistent with the results obtained from the rescaled JT \mathcal{W} base matrix set, the enhancement that yields the largest performance improvement is zonal aggregation, though once again ferry is problematic with a sign change occurring, i.e. after pivoting at the aggregate level a reduction in trips is predicted.

The revised switch point gives a further improvement, with predicted growth for car driver and car passenger matching synthetic within 5%, and for rail the synthetic growth ratio improves from 0.58 to 0.89. The results for bus improve from 0.37 to 0.60, but 0.60 is still a significant under-prediction of synthetic growth (again, this should be improved when the new STM is implemented).

Based on this all matrix comparison of synthetic and predicted growth, the application of origin-mode normalisation yields slight performance improvements for car driver, rail and bus, but a worsening of the performance for car passenger. For bus, pivoted growth remains significantly below that predicted from the synthetic matrices.

The second normalisation step ensures that overall synthetic and pivoted growth by mode match exactly. For bus, a factor of 1.45 is required which means that at the origin level predicted growth will be significantly higher than synthetic. When the new STM is operational the overall mode normalisation factors required for bus when working with the expanded HTS base matrices should be reviewed to check that they are not too far from 1, i.e. that the expected improvement that should follow from the greater similarity between B and Sb for bus has in fact occurred.

5.1.3 Distribution of demand across the eight cases

Tabulations of the distribution of demand across the eight cases have been examined for each of the five modes at each step of the enhancements to the process. To illustrate the patterns of changes, the distributions for the bus mode pivoted using the expanded HTS matrix set are presented here. This matrix set has been chosen as the base matrices are relatively sparse.

Table 26 presents the distribution of demand across the eight cases in the original pivoting process, where pivoting is undertaken at the 2690 zone level.

Table 26: Distribution of bus demand, expanded HTS matrix set, original pivoting process

Case	B	Sb	Sf	B > 0	Sb > 0	Sf > 0	P > 0
1	0	0	0	0.0 %	0.0 %	0.0 %	0.0 %
2	0	0	> 0	0.0 %	0.0 %	10.4 %	9.7 %
3	0	> 0	0	0.0 %	0.5 %	0.0 %	0.0 %
4n	0	> 0	> 0	0.0 %	96.8 %	74.1 %	0.0 %
4e	0	> 0	> 0	0.0 %	2.0 %	15.0 %	8.7 %
5	> 0	0	0	1.8 %	0.0 %	0.0 %	1.6 %
6	> 0	0	> 0	0.3 %	0.0 %	0.0 %	0.3 %
7	> 0	> 0	0	0.0 %	0.0 %	0.0 %	0.0 %
8n	> 0	> 0	> 0	40.0 %	0.3 %	0.1 %	29.1 %
8e	> 0	> 0	> 0	57.9 %	0.4 %	0.3 %	50.6 %
Total				100.0 %	100.0 %	100.0 %	100.0 %

This distribution illustrates two significant issues. First, 97% of *Sb* demand and 74% of *Sf* demand occur in case 4n where no demand is predicted after the pivoting. Second, over 50% of the predicted demand falls under extreme growth regime case 8e, and this demand corresponds to less than 0.5% of the model’s demand predictions *Sb* and *Sf*.

Table 27 presents the distribution of bus demand across the eight cases following zonal aggregation to the SLA zone level.

Table 27: Distribution of bus demand, expanded HTS matrix set, impact of aggregation

Case	B	Sb	Sf	B > 0	Sb > 0	Sf > 0	P > 0
1	0	0	0	0.0 %	0.0 %	0.0 %	0.0 %
2	0	0	> 0	0.0 %	0.0 %	0.0 %	0.0 %
3	0	> 0	0	0.0 %	0.0 %	0.0 %	0.0 %
4n	0	> 0	> 0	0.0 %	37.8 %	41.3 %	0.0 %
4e	0	> 0	> 0	0.0 %	2.3 %	13.3 %	5.7 %
5	> 0	0	0	0.0 %	0.0 %	0.0 %	0.0 %
6	> 0	0	> 0	0.0 %	0.0 %	0.0 %	0.0 %
7	> 0	> 0	0	0.0 %	0.0 %	0.0 %	0.0 %
8n	> 0	> 0	> 0	77.6 %	56.2 %	41.1 %	73.5 %
8e	> 0	> 0	> 0	22.4 %	3.7 %	4.2 %	20.8 %
Total				100.0 %	100.0 %	100.0 %	100.0 %

Substantial improvements are observed as a consequence of zonal aggregation. The percentage of *Sb* and *Sf* demand that occurs in case 4 is significantly reduced, though a substantial proportion of case 4 demand remains, and in case 8 the balance between normal and extreme growth has altered with close to three-quarters of demand now predicted by the standard case 8n growth regime.

Table 28: Distribution of bus demand, expanded HTS matrix set, impact of revised case 8 switch point for extreme growth

Case	B	Sb	Sf	B > 0	Sb > 0	Sf > 0	P > 0
1	0	0	0	0.0 %	0.0 %	0.0 %	0.0 %
2	0	0	> 0	0.0 %	0.0 %	0.0 %	0.0 %
3	0	> 0	0	0.0 %	0.0 %	0.0 %	0.0 %
4n	0	> 0	> 0	0.0 %	37.8 %	41.3 %	0.0 %
4e	0	> 0	> 0	0.0 %	2.3 %	13.3 %	5.1 %
5	> 0	0	0	0.0 %	0.0 %	0.0 %	0.0 %
6	> 0	0	> 0	0.0 %	0.0 %	0.0 %	0.0 %
7	> 0	> 0	0	0.0 %	0.0 %	0.0 %	0.0 %
8n	> 0	> 0	> 0	99.3 %	59.8 %	44.9 %	92.2 %
8e	> 0	> 0	> 0	0.7 %	0.1 %	0.4 %	2.7 %
Total				100.0 %	100.0 %	100.0 %	100.0 %

The modification to the case 8 switch point has a substantial impact on the distribution of demand between cases 8n and 8e. The proportion of demand that is predicted under the case 8e extreme growth rule has reduced from 20.8% to 2.7%. This shift results in the improved correspondence between synthetic and predicted growth highlighted in Table 25. On the other hand, around 40% of synthetic demand still occurs for case 4n where predicted demand is zero.

As discussed in Section 4, synthetic growth for bus remains significantly below predicted growth after the application of the revised case 8 switch point, and so two normalisation steps were applied. The application of the normalisation factors has no impact on the distribution of demand across the eight cases and so the results given in Table 28 represent the distribution at the end of the enhanced pivoting process.

5.2 Station to station matrices

5.2.1 Base matrix sparsity

An aggregation procedure over stations has not been applied for the station to station matrices used for rail. Aggregating over stations is outside the scope of this project, and furthermore would have less benefit relative to the zonal matrices because, as shown in Table 29, the sparsity indices for the RailCorp matrices are significantly lower than those observed for the JTW and HTS matrices. The lower sparsity follows from the fact that the number of stations (343) is significantly lower than the number of zones (2690). Higher sparsity indices are observed in Table 29 for the HTS rail station to station matrices. However, this matrix set has been used for testing purposes only and will not be used when the STM is used in application.

Table 29: Sparsity indices, station to station rail matrices

Matrix set	Sparsity index
RailCorp	3.75
HTS rail	65.55

5.2.2 Comparison of synthetic and predicted growth

For the station to station matrices, there is no zonal aggregation step, but nonetheless the impact of the revision to the switch point between normal and extreme growth can be assessed.

Tests of origin station normalisation revealed that application of the procedure resulted in a *worse* correspondence between synthetic and predicted growth. The demand models predict growth in demand at the origin zone level, rather than the origin station, and so normalising at the origin station does not have the same behavioural rationale as normalising at the origin zone level. Given that applying an origin station normalisation resulted in a worse correspondence between synthetic and predicted growth only, an *overall* normalisation has been applied for the station to station matrices.

Table 30 presents synthetic and predicted growth comparisons for the station to station matrices.

Table 30: Synthetic and predicted growth comparisons, station to station matrices

Matrix set	Synthetic growth	Predicted / synthetic growth ratios		
		Original pivoting process	Revised switch point	Normalisation: total mode
RailCorp	78.2 %	0.87	1.20	1.00
HTS rail	78.2 %	0.60	1.14	1.00

For the RailCorp matrix set, moving to the revised switch point results in higher predicted growth, with a change from a 13% under-prediction of growth to a 20% over-prediction, so performance is slightly worse. For the HTS rail matrix moving to a revised switch point again results in higher predicted growth, but in this case the performance is significantly improved with a change from a 40% under-prediction to a 14% over-prediction of synthetic growth.

The application of a total mode normalisation by definition ensures that synthetic and predicted growth match exactly.

5.2.3 Distribution of demand across the eight cases

Table 31 presents the distribution of demand across the eight cases for the RailCorp matrix set for the original pivoting process.

Table 31: Distribution of rail demand, RailCorp matrix set, original pivoting process

Case	B	Sb	Sf	B > 0	Sb > 0	Sf > 0	P > 0
1	0	0	0	0.0 %	0.0 %	0.0 %	0.0 %
2	0	0	> 0	0.0 %	0.0 %	25.0 %	24.7 %
3	0	> 0	0	0.0 %	0.2 %	0.0 %	0.0 %
4n	0	> 0	> 0	0.0 %	6.3 %	5.0 %	0.0 %
4e	0	> 0	> 0	0.0 %	0.5 %	2.9 %	1.5 %
5	> 0	0	0	1.2 %	0.0 %	0.0 %	0.7 %
6	> 0	0	> 0	2.3 %	0.0 %	2.9 %	4.3 %
7	> 0	> 0	0	9.3 %	8.4 %	0.0 %	0.0 %
8n	> 0	> 0	> 0	70.9 %	81.2 %	54.8 %	48.6 %
8e	> 0	> 0	> 0	16.2 %	3.5 %	9.5 %	20.2 %
Total				100.0 %	100.0 %	100.0 %	100.0 %

The percentage of synthetic base demand that occurs in cases 1–4, where the base matrix is 0 is relatively low at 7.0%. A quarter of synthetic future demand occurs in case 2; however, for this case demand is predicted after pivoting ($P=Sf$) so this growth is not lost.

Table 32 shows the impact on the distribution of rail demand over the eight cases of revising the switch point for case 8 extreme growth.

Table 32: Distribution of rail demand, RailCorp matrix set, impact of revised case 8 switch point for extreme growth

Case	B	Sb	Sf	B > 0	Sb > 0	Sf > 0	P > 0
1	0	0	0	0.0 %	0.0 %	0.0 %	0.0 %
2	0	0	> 0	0.0 %	0.0 %	25.0 %	21.3 %
3	0	> 0	0	0.0 %	0.2 %	0.0 %	0.0 %
4n	0	> 0	> 0	0.0 %	6.3 %	5.0 %	0.0 %
4e	0	> 0	> 0	0.0 %	0.5 %	2.9 %	1.3 %
5	> 0	0	0	1.2 %	0.0 %	0.0 %	0.6 %
6	> 0	0	> 0	2.3 %	0.0 %	2.9 %	3.7 %
7	> 0	> 0	0	9.3 %	8.4 %	0.0 %	0.0 %
8n	> 0	> 0	> 0	82.5 %	83.1 %	56.5 %	58.4 %
8e	> 0	> 0	> 0	4.6 %	1.6 %	7.8 %	14.7 %
Total				100.0 %	100.0 %	100.0 %	100.0 %

The revised switch point for case 8 extreme growth results in an increase in demand predicted under the normal growth regime (8n) and a corresponding slight reduction in the demand predicted under the extreme growth regime (8e).

6.1 **Zone to zone matrices**

The zone to zone matrices used for this analysis are defined using the 2006 zoning system.

The current assumption of setting cell values with less than 0.001 trips to be zero is reasonable, with only tiny fractions of synthetic demand (at most 0.6%) lost as a result of this assumption.

Pivoting at the aggregate 80 SLA level rather than the 2690 zone level results in a significant improvement to the correspondence between synthetic and predicted growth. As would be expected, matrix sparsity is significantly reduced by aggregation. Over the eight cases, demand is shifted from case 4, where the base matrix is 0, to case 8 where the base matrix is defined, and this contributes to the improved correspondence between synthetic and predicted growth.

Revising the switch point between normal and extreme growth for case 8 results in a further improvement to the correspondence between synthetic and predicted growth. This improvement is observed because the revised switch point results in a higher fraction of case 8 demand being classified as 'normal', where the standard factor pivoting formula is applied. The revised switch point formulation is more robust for working with matrices at a range of different levels of sparsity.

Incorporation of origin-mode normalisation ensures that synthetic and predicted growth by mode match exactly at the aggregate SLA origin level. However, this does not ensure that overall synthetic and predicted growth match exactly when summed across all origins. In particular, significant differences are observed for bus due to differences in the distributions of demand between the base matrices B and the synthetic base Sb .

Therefore a second normalisation step has been applied so that overall synthetic and predicted growth match exactly for each mode. The second step is termed overall mode normalisation. For car driver, car passenger and rail, this step applies factors relatively close to 1, so the changes relative to the origin-mode normalised results are modest. However, for bus normalisation factors of up to 1.45 are required, which means that at the origin SLA level synthetic growth will be over-predicted by up to 45%. A difference of this level is not desirable, and so when the new STM is operational it is recommended that this analysis be repeated to check that overall normalisation factors closer to 1 are obtained for bus.

The results for ferry are problematic. Significant differences between the base matrices B and the synthetic base Sb for ferry underlie this result, and if the results are as poor when matrices from the new STM are available BTS may decide not to use the pivoting process for ferry.

6.2 Station to station matrices

For the RailCorp matrix set, the station to station matrix which will be used in implementation, the existing pivoting process gives a fairly good correspondence between synthetic and predicted growth.

No station aggregation step is applied, as the matrices are much less sparse than the zone-zone matrices and as a result the original process works fairly well. Note that a zonal aggregation would not be appropriate because trips in the matrix are associated with stations rather than zones.

Revising the switch point between normal and extreme growth for case 8 results in synthetic growth being slightly over-predicted.

Tests demonstrated that applying an origin station normalisation did not improve the performance of the pivoting process, and therefore an origin station normalisation has not been implemented.

A total mode normalisation step has been applied so that the percentage synthetic growth in total rail trips is matched exactly at the end of the pivoting process.

6.3 Recommendations

It is recommended that BTS implements the enhanced pivoting system for use in the STM, as it has been demonstrated to give significant improvements in performance relative to the original pivoting process.

All of the tests presented in this report have been run using matrices generated from the existing version of the STM. When the new version of the STM is operational in Sydney it is recommended that the analysis presented in this report for the enhanced pivoting process is repeated using forecasts from the new version of the STM. Particular issues to examine when the analysis is repeated are the performance of the pivoting process for bus, where it is hoped an improved correspondence between synthetic base and base matrices will be achieved, and a review of the performance of the process for the problematic ferry mode.

REFERENCES

Reference list

Daly, A., J. Fox and B. Patrui (2011) “Pivoting in Travel Demand Models”, paper presented at European Transport Conference, Glasgow.

Daly, A., J. Fox and J. G. Tuinenga (2005) “Pivot-Point Procedures in Practical Travel Demand Forecasting”, European Congress of the Regional Science Association, Amsterdam. Manheim, M. L. (1979) *Fundamentals of Transportation System Analysis*, Cambridge, Mass: MIT Press.

APPENDICES

Appendix A: Mode normalisation with a utility maximisation approach

Calculations within a utility maximisation approach

Following Daly (1987) we define a tree logit model using the RUM1 (“non-normalised”) specification, which is used in the STM models. The specification depends on a tree function that gives the ancestor node t_j for each node j in the model. However, the notation used here is changed to simplify the discussion of pivoting.

A tree logit model (RUM1) is defined recursively by the following

$$\log p_{k|t_k} = V_k - \log \sum_{t_h=t_k} \exp V_h \quad (1)$$

where V gives the measured utility for each alternative.

For composite alternatives, i.e. those that are the ancestor of some other node, utility is transmitted through the ‘logsum’

$$V_j = \theta_j \log \sum_{t_h=j} \exp V_h \quad (2)$$

where $0 < \theta_j \leq 1$.

The intention is to set up a pivoted model, with the same form as the original model but with utility functions amended. Terms in this model are indicated in the equations by a superscript, i.e. p^* and V^* , so we get

$$\log p_{k|t_k}^* = V_k^* - \log \sum_{t_h=t_k} \exp V_h^* \quad (3)$$

$$V_j^* = \theta_j \log \sum_{t_h=j} \exp V_h^* \quad (4)$$

The amendment to the utility function is required to serve two functions:

- It must be consistent with the utility transmission (4).
- It must give the pivoting that is required at each level of the tree.

It turns out that the amendment that is needed is:

$$V_h^* = V_h + K_h + L_{t_h} \quad (5)$$

where $K_h = \log \left(\frac{b_h}{p_h} \right)$;

b_h is the observed fraction of total trips choosing alternative h (note that this is the *marginal* fraction, not conditional on t_h);

$$L_g = \sum_{i \in A_g} K_i (1 - \theta_i) \frac{\prod_{l \in A_{t_i}} \theta_l}{\prod_{l \in A_g} \theta_l};$$

A_g is the set of ancestor nodes of g : $\{g, t_g, t_{t_g}, \dots, i \mid t_i = r\}$; note that this includes the node itself but excludes the root r of the tree.

The effect of the fraction in L_g is to divide by the θ parameters for all nodes in the ancestor set from g to i inclusive.

Because of the way it is defined, L_{t_h} in (5) involves only t_h and higher nodes, so when we calculate the conditional probability for h that term disappears and the probabilities (using equations 3 and 5) are determined by $(V + K)$ only

$$\log p_{h|t_h}^* = (V_h + K_h) - \log \sum_{t_k=t_h} \exp(V_k + K_k) \quad (6)$$

Similarly, when we apply the recursive formula (4) to get the utility for composite alternative j , L_j is constant and can be taken outside the logsum as a constant to obtain

$$\begin{aligned} V_j^* &= \theta_j \log \sum_{t_h=j} \exp V_h^* \\ &= \theta_j \log \sum_{t_h=j} \exp(V_h + K_h + L_j) \\ &= \theta_j L_j + \theta_j \log \sum_{t_h=j} \exp(V_h + K_h) \end{aligned} \quad (7)$$

Moreover, we can calculate the first term in (7) further, cancelling out the factor θ_j

$$\begin{aligned} \theta_j L_j &= \theta_j \sum_{i \in A_j} K_i (1 - \theta_i) \frac{\prod_{l \in A_{t_i}} \theta_l}{\prod_{l \in A_j} \theta_l} \\ &= K_j (1 - \theta_j) + \sum_{i \in A_{t_j}} K_i (1 - \theta_i) \frac{\prod_{l \in A_{t_i}} \theta_l}{\prod_{l \in A_{t_j}} \theta_l} \\ &= K_j (1 - \theta_j) + L_{t_j} \end{aligned} \quad (8)$$

The second term in (7) can also be calculated:

$$\theta_j \log \sum_{t_h=j} \exp(V_h + K_h) = \theta_j \log \sum_{t_h=j} \left(\frac{b_h}{p_h} \right) \exp V_h \quad (9)$$

Now if $t_h = j$, we can use (1) and (2) to calculate

$$p_h = p_j \cdot p_{h|j} = p_j \cdot \frac{\exp V_h}{\sum_{t_k=j} \exp V_k} = p_j \cdot \frac{\exp V_h}{\exp \frac{V_j}{\theta_j}} \quad (10)$$

so we can substitute in (9) to obtain

$$\begin{aligned} \theta_j \log \sum_{t_h=j} \exp(V_h + K_h) &= \theta_j \log \sum_{t_h=j} \left(\frac{b_h \cdot \exp \frac{V_j}{\theta_j}}{p_j} \right) \\ &= \theta_j \log \left(\exp \frac{V_j}{\theta_j} \right) + \theta_j \cdot \log \sum_{t_h=j} \left(\frac{b_h}{p_j} \right) \\ &= V_j + \theta_j \log \left(\frac{b_j}{p_j} \right) = V_j + \theta_j K_j \end{aligned} \quad (11)$$

because $\sum_{t_h=j} b_h = b_j$.

Given the results (8) and (11) we are now in a position to simplify equation (7)

$$V_j^* = V_j + K_j + L_{t_j} \quad (12)$$

We note that equation (5) reappears at the next higher level, indicating that the change to the model is carried through to the next level and hence to all levels. That is, we satisfy the first requirement of consistency in the amended model.

In equation (6) we noted that the marginal probabilities at each level in the tree are determined by $(V + K)$ only. We can then obtain the ratio of probabilities for two alternatives in the same nest:

$$\frac{p_k^*}{p_h^*} = \frac{\exp V_k}{\exp V_h} \cdot \frac{\exp K_k}{\exp K_h} = \frac{p_k}{p_h} \cdot \frac{b_k/p_k}{b_h/p_h} = \frac{b_k}{b_h} \quad (13)$$

That is, the formulation of K ensures we obtain the observed proportions at each level in the model. Provided the total number of trips is correct, this implies that we obtain the correct number for each alternative. The utility correction (5) thus satisfies both the requirements to obtain a pivoted model.

Extending to multiple segments

The above formulae apply for any single segment. For pivoting we require to aggregate over segments (person types, purposes) and to apply the correction factors to the aggregate trips. In this case we obtain a more general version of equation (13) for segment s

$$\frac{p_{js}^*}{p_{hs}^*} = \frac{p_{js}}{p_{hs}} \cdot \frac{b_j/p_j}{b_h/p_h} = \frac{b_{js}}{b_{hs}} \quad (14)$$

where $b_{js} = b_j \cdot p_{js}/p_j$, which is the base matrix multiplied by the fraction of flow indicated by the model.

This fraction of flow appears to be the best indication we can get of the fraction of base flow. Moreover, adding up over the segments s gets back to the total base matrix, as required.

What does this mean?

The lengthy calculations above show that it is possible to make a correction to the utility function that allows a tree logit model to match exactly the observed choice proportions. If the total number of trips is correct, of course, this means we match the total observed numbers. The utility correction propagates through the tree, so that it can be added at the bottom level alone and the effect will operate at all levels automatically.

In practice, the model will be run for the base year and the calculations $(K_h + L_{t_h})$ made for all the elementary nodes. Adding these to the utilities will then ensure that the model exactly reproduces the base year observations. Changes to the utilities will then change the forecasts as usual. We can note that the utility correction is entirely fixed, once the base case has been run. The additions to utility are constants and the utility-maximising character of the model is maintained.

Dealing with all eight cases

The formulae above note are possible for cases 8n and 7 only because these are the only places we can calculate $K=\log(B/Sb)$. Suppose for cases 1–6 we set $K=0$. Then we will get the right answer for cases 1–3, 7 and 8n. For the other cases:

- Cases 5 and 6 are forecast consistently with utility because they just add B to Sf and this is like assuming there is some fixed demand – captive, perhaps.
- Case 8e is forecast consistently with utility if $B>Sb$, because we are then just taking a positive demand $X2*(B/Sb-1)$ as being fixed.

Case 4n, which is just set to 0, is also consistent with utility theory but the problem here is that the switch point between case 4n and case 4e depends on Sf and may therefore differ between forecast scenarios. In this case, and in case 8e for $B<Sb$, there is no correct formula, because there is no utility function that will give the results we need, which are linear in Sf but with a negative offset. For this reason there can be no guarantee that we get utility-consistent results.

So a proposal would be to make a two-step pivot. First, do the utility-based pivot, as set out above, with $K=0$ for cases 1–6. This will give us Sf^* , which is utility-consistent for cases 1–3, 5–7 and 8n. For cases 4 and 8e, we then do a second-stage pivot, using Sf^* , to get the final result. For cases 5 and 6 we just have to add B.

Dealing with OD matrices

Some base matrices are defined on an OD basis, rather than on a PA or tour basis, though this does not appear to be a significant problem in the Sydney case. The problem with these matrices is that it is difficult to apply the notion of utility maximisation, because we do not know whether the traveller is choosing the origin or the destination.⁶

The most obvious approach, which seems to be consistent with other procedures discussed in this note, is to allocate OD matrices to PA in the proportions indicated by the model. However, this seems rather onerous. But if we calculate a matrix split in this way, the ‘observed’ number of tours in one direction would be given by

$$B^1 = B \frac{Sb^1}{Sb^1+Sb^2}$$

with the superscripts indicating the direction of the tours and B being the base trip matrix, which is assumed to be symmetric. If we then make a case 8n pivot, for tours in both directions, we would obtain:

$$P = B^1 \frac{Sf^1}{Sb^1} + B^2 \frac{Sf^2}{Sb^2} = B \frac{Sf^1+Sf^2}{Sb^1+Sb^2}$$

which is exactly what we would obtain if we followed a conventional procedure of calculating the synthetic trips, then making the pivot.

⁶ There may also be a problem when dealing with detours (non-home-based) because again we do not necessarily know which trip end is chosen.

Practical procedures

Two alternative procedures seem to be available.

1. Using the utility formulae directly

To apply the note directly we would calculate a matrix of K+L, setting K=0 for cases 1–6, then apply the corrections. Then we make the changes for cases 4, 5, 6 and 8e outside the TRAVDEM. Some thought is needed to do this efficiently.

2. Normalisation

What the note shows is that it is possible to set up a model that matches the base observations exactly and that is consistent with utility. This can be achieved equivalently by factoring or by adding utility (which is the more obvious suggestion in the note). For factoring, we have to work top-down to ensure we have achieved consistency.

This procedure can be compared with the normalisation we have previously done. The main objective of normalisation is to ensure that after factoring the cells, we return to the row or other total we expected. The point we seem to have missed is that the normalisation at different levels has to be consistent. For example, in the UK Long Distance Model, we tested mode then row normalisation before deciding that row then mode normalisation was best because it minimised inconsistency. But we do not seem to have tried normalising the rows (lower level) to give the total given by the normalised mode pivoting.

Moreover, if we consider what happens when we change the utilities for forecasting, based on equation (13)

$$\frac{p_k^1}{p_h^1} = \frac{\exp(V_k + \Delta V_k)}{\exp(V_h + \Delta V_h)} \cdot \frac{\exp K_k}{\exp K_h} = \frac{p_k^*}{p_h^*} \cdot \frac{\exp(\Delta V_k)}{\exp(\Delta V_h)} \quad (15)$$

where p_k^1 is the forecast proportion choosing k and

ΔV_k is the change in utility of alternative k .

which is exactly the formulation used in incremental models, e.g. in the WebTAG advice (2009).

References

Daly, A. J. (1987) “Estimating ‘Tree’ Logit Models”, *Transportation Research*, Vol. 21B, pp 251–67.

WebTAG (2009) “Variable Demand Modelling – Key Processes”, Unit 3.10.3C, paragraph 4.1.9, <http://www.dft.gov.uk/webtag/documents/expert/pdf/unit3.10.3c.pdf>