Posterior analysis of random taste coefficients in air travel behaviour modelling

Stephane Hess[*]

Institute for Transport Planning and Systems, ETH Zürich, and Institute of Transport and Logistics Studies, The University of Sydney

Abstract

Increasing use is being made of random coefficients structures, such as Mixed Logit, in the analysis of air travel choice behaviour. These models have the advantage of being able to retrieve random variations in sensitivities across travellers. An important issue however arises in the computation of willingness to pay indicators, such as the valuation of travel time savings, on the basis of randomly distributed coefficients. Indeed, with the standard approach of using simulation of the ratios across random draws, major problems can be caused by outliers, leading to biased trade-offs, which in turn lead to major issues in policy analyses. Here, a different approach is explored, making use of individual-specific draws from the random distributions, conditioned on the observed sequence of choices for each respondent. An analysis making use of stated preference data for airport and airline choice confirms the advantages of the approach using conditional draws, producing much more realistic distributional patterns for a range of willingness to pay indicators.

KEYWORDS: air travel choice behaviour, mixed logit, taste heterogeneity, individual-specific taste coefficients

1 Introduction

Air travel behaviour research has seen a flurry of activity over recent years, with analysts increasingly making use of advanced discrete choice methods when representing complex air travel choice processes. These mathematical models not only enable researchers to explicitly recognise the multi-dimensional nature of the choice processes but also allow for a representation of the correlation along and between these various dimensions of choice. Furthermore, the most recent batch of models also allow for a representation of random variations in behaviour across respondents[1].

---

[*]Corresponding author: stephane.hess@ivt.baug.ethz.ch

[1]For an in-depth discussion of discrete choice methods, see Train (2003).

Existing applications range from the choice of air as a mode of travel (González-Savignat, 2004), to the choice of airport in multi airport regions (Pels *et al.*, 2001, 2003; Pathomsiri *et al.*, 2004; Basar and Bhat, 2004), the choice of airline or fare class (Proussaloglou and Koppelman, 1995; Chin, 2002), and the choice of access mode Monteiro and Hansen (1997); Psaraki and Abacoumkin (2002). A number of authors have explicitly recognised the multi-dimensional structure of the choice process, looking for example at the joint choice of airport and airline (Bondzio, 1996) or even the choice of an airport, airline and access mode triplet (Hess and Polak, 2006a). Increasingly, researchers also make use of the more advanced model structures available, for a representation of random taste heterogeneity across travellers, or the multi-dimensional correlation between alternatives sharing sub-choices along some of the choice dimensions (Hess and Polak, 2006b).

While the majority of studies of air travel choice behaviour make use of revealed preference (RP) data, an increasing number of analyses are now carried out on stated preference (SP) data[2]. While posing certain problems in terms of response quality (Louviere *et al.*, 2000), studies using SP data have the advantage of being based on accurate records of all information presented to respondents, which is not generally the case with RP data. As such, it should come as no surprise that SP studies are generally more successful in retrieving significant effects for crucial factors such as air fares and frequent flier benefits.

One point of interest in studies of travel behaviour is the representation of variations in choice behaviour across travellers in the form of different sensitivities to changes in explanatory variables, such as air fares and access time. Given the limitations of a purely deterministic approach (e.g. segmentation), modellers increasingly rely on a random representation of these variations in tastes. The mixed multinomial logit (MMNL) model (Train, 2003) is increasingly being used in studies of travel behaviour, including in the area of aviation (Hess and Polak, 2005b,a).

Despite their popularity, important issues arise with the use of random coefficients models, such as MMNL. Not only are they far more expensive to estimate and apply (Bhat, 2001; Hess *et al.*, 2006), but there is a need to make an a priori choice of mixing distribution for each random coefficient, where the majority of applications rely exclusively on the normal distribution. The choice of distribution not only has potentially significant impacts on model performance and behaviour, but also leads to issues in interpretation, especially in the context of trade-offs between two randomly distributed coefficients (Hensher and Greene, 2003; Hess *et al.*, 2005). Indeed, the distribution of such trade-offs generally needs

---

[2]Examples of this include the work of Bradley (1998), Algers and Beser (2001), Adler *et al.* (2005) and Hess *et al.* (2007).

to be obtained using simulation methods, producing a large number of pairs of draws, and calculating the average of the ratio between draws across all pairs. Even when taking into account correlation between the two randomly distributed coefficients, the presence of outliers (very large or small values) for one or both of the coefficients will lead to extreme values in the simulation of the ratio. This can lead to an overestimation of the range for the distribution of the ratio, while potentially also biasing the mean value.

Here, we take another look at the interpretation of results obtained from MMNL models in the context of air travel behaviour research, making use of conditional rather than estimated distributions[3].

2 Methodology

The MMNL model assumes that tastes vary randomly across respondents according to some pre-specified distributions. Here, we let β be a vector of taste coefficients that are jointly distributed according to $f(\beta \mid \Omega)$, where Ω is a vector of distributional parameters to be estimated. Let Y_n give the sequence of observed choices for respondent n, and let $L(Y_n \mid \beta)$ give the probability of observing this sequence of choices with a specific value for β. Then it can be seen (Train, 2003) that the probability of observing the specific value of β given the choices of respondent n is:

$$K(\beta \mid Y_n) = \frac{L(Y_n \mid \beta) f(\beta \mid \Omega)}{\int_\beta L(Y_n \mid \beta) f(\beta \mid \Omega) \, \mathrm{d}\beta} \tag{1}$$

We replace the continuous formulation by a discrete approximation using summation over a very high number of draws. A mean for the conditional distribution for respondent n is then obtained as:

$$\widehat{\beta_n} = \frac{\sum_{r=1}^R [L(Y_n \mid \beta_r) \beta_r]}{\sum_{r=1}^R L(Y_n \mid \beta_r)}, \tag{2}$$

where β_r with $r = 1, \ldots, R$ are independent multi-dimensional draws with equal weight from $f(\beta \mid \Omega)$ at the estimated values for Ω.

With the help of $\widehat{\beta_n}$, it is possible to calculate a single value for each trade-off per respondent, and distributional statistics across respondents can be obtained straightforwardly. Here, we compare the findings obtained using this approach

---

[3]For applications of this approach in other areas of travel behaviour research, see for example Revelt and Train (1999), Sillano and Ortúzar (2004) and Greene *et al.* (2005).
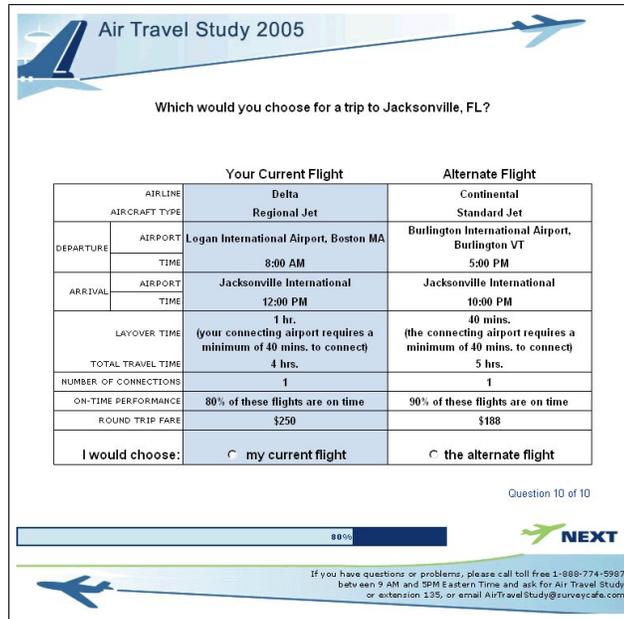
Figure 1: Example screen-shot for SP survey

with those when using simulation across random draws to approximate the distribution of the ratio. Even with this approach, some risk of bias remains. Indeed, by calculating an individual-specific trade-off on the basis of a ratio of means of the two individual-specific conditional distributions, we disregard any information on the variance of the distributions for each individual. A more accurate approach would make use of simulation of the ratio across the conditional distributions. This, however, would again lead to problems with outlying values.

3 Data

The analysis makes use of SP data collected via the Internet by Resource Systems Group in the US (Resource Systems Group Inc., 2003). Specifically, we make use of the 2005 version of the survey, with a sample of 4,136 observations collected from 517 travellers.

Prior to the SP survey, information was collected on a traveller's most recent air trip, along with detailed socio-demographic information. Each traveller is then faced with 8 binomial choices, where in each case, a choice is offered between the current, or RP alternative, and an alternative option, the SP alternative. While

the attributes of the RP alternative remain fixed across the 8 choice sets, those of the SP alternative are varied according to an experimental design. The airports and airlines used for the SP alternatives are selected on the basis of information gathered from respondents in terms of a ranking of the airports and airlines available to them.

Aside from the airport and airline names, from which access times can be inferred, the attributes used to describe the alternatives in the SP survey include flight time, the number of connections, the air fare, the arrival time (used to calculate schedule delays), the aircraft type, and the on-time performance of the various services. Access cost was not included (in the absence of an actual specification of the mode choice dimension), and no choice is given between different travel classes; this can be regarded as an upper-level choice, taken before the actual air journey choices. An example of one choice situation is shown in Figure 1.

Certain doubts as to the response quality arise in the context of SP studies (Louviere *et al.*, 2000). However, in the face of problems with the quality of RP data in air transport, any SP disadvantages are outweighed by the quality of the data on availabilities and explanatory variables. Furthermore, the design used here is simple enough so as not to place too big a cognitive burden on respondents, while still being complex enough to approximate a real-world choice process.

4   Main modelling analysis

Models were jointly estimated on all $4,136$ observations, as opposed to using a segmentation along a socio-demographic dimension such as trip purpose. This is mainly motivated by the poor results obtained when segmenting the data, with very low levels of statistical significance for some of the parameters. Furthermore, no significant interactions with income or other continuous socio-demographic attributes were observed.

4.1   Model specification

All attributes are specified to enter the utility function in a linear fashion, such that the observed utility for the RP alternative is given by:

$$\begin{aligned}
U_{\mathrm{RP}} = \beta_{\mathrm{current}} & \\
& + \beta_{\mathrm{access\ time}} \cdot \mathrm{access\ time}_{\mathrm{RP}} + \beta_{\mathrm{air\ fare}} \cdot \mathrm{air\ fare}_{\mathrm{RP}} + \beta_{\mathrm{flight\ time}} \cdot \mathrm{flight\ time}_{\mathrm{RP}} \\
& + \beta_{\mathrm{OTP}} \cdot \mathrm{OTP}_{\mathrm{RP}} + \beta_{1\ \mathrm{connection}} \cdot \delta_{1\ \mathrm{connection,RP}} + \beta_{2\ \mathrm{connections}} \cdot \delta_{2\ \mathrm{connections,RP}} \\
& + \beta_{\mathrm{standard\ FF}} \cdot \delta_{\mathrm{standard\ FF,RP}} + \beta_{\mathrm{elite\ FF}} \cdot \delta_{\mathrm{elite\ FF,RP}} + \beta_{\mathrm{closest\ airport}} \cdot \delta_{\mathrm{closest\ airport,RP}}
\end{aligned}$$

(3)

where β parameters are to be estimated from the data. The meaning of the first four entries in Equation 3 should be clear. The fifth parameter, β_{OTP}, relates to the on-time performance (in percentage points) of an alternative. The two dummy variables $\delta_{1 \text{ connection,RP}}$ and $\delta_{2 \text{ connections,RP}}$ are set to 1 for flights with one and two connections respectively, while $\delta_{\text{standard FF,RP}}$ and $\delta_{\text{elite FF,RP}}$ are set to 1 if the respondent holds standard or elite frequent flier (FF) membership with the current airline. Finally, $\delta_{\text{closest airport,RP}}$ is set to 1 if the airport used in the RP alternative is that closest to the respondent's home. The utility function for the SP alternative is specified in a similar fashion, with the absence of the RP constant (β_{current}), and with SP as opposed to RP values for the various attributes and dummy variables. The specification used is by no means complete in terms of attributes included, as well as in their treatment (purely linear).

Two types of model were estimated on the data, a basic MNL structure, and a more advanced MMNL model[4]. In the MMNL model, the likelihood is specified with the integration carried out over sequences of choices for the same respondent as opposed to individual choices (Train, 2003), leading to an assumption of constant tastes across replications for the same respondent. To further account for SP effects in terms of serial correlation across observations for the same respondent, an individual-specific SP factor is included in the utility functions[5]. With $V_{\text{n,i}}$ giving the observed utility for alternative i and respondent n, we have:

$$U_{\text{n,RP}} = V_{\text{n,RP}} + \varepsilon_{\text{n,RP}} + \varphi\,\xi_{\text{RP}}$$
$$U_{\text{n,SP}} = V_{\text{n,SP}} + \varepsilon_{\text{n,SP}} + \varphi\,\xi_{\text{SP}}, \tag{4}$$

where $\varepsilon_{\text{n,RP}}$ and $\varepsilon_{\text{n,SP}}$ are the usual type I extreme value terms, distributed identically and independently over alternatives and observations. The two additional terms ξ_{RP} and ξ_{SP} are normally distributed random variables with a mean of zero and a standard deviation of 1, distributed independently across alternatives and individuals, but not across observations for the same individual. In conjunction with the multiplication by φ, this specification allows for an individual-specific effect that is shared across alternatives. The inclusion of this term can in general be expected to lead to an upwards correction of the standard errors (Ortúzar et al., 2000).

---

[4]Both are coded in Ox 4.2 (Doornik, 2001), where specific code was also written for the calculation of the means of the conditional distributions.

[5]The author would like to thank Andrew Daly for this suggestion.

4.2 Results

The results for the two models are summarised in Table 1[6]. All coefficients in the MNL model are statistically significant and of the expected sign. In the MMNL model, the estimate for φ is highly significant, indicating the presence of a significant individual-specific effect. Additionally, significant levels of random taste heterogeneity are retrieved for four coefficients; $\beta_{\text{access time}}$, $\beta_{\text{air fare}}$, $\beta_{\text{flight time}}$ and $\beta_{\text{closest airport}}$. A Normal distribution was used for all four coefficients. This is a major assumption (Hess *et al.*, 2005), but is consistent with the overwhelming majority of other studies, hence allowing us to provide general conclusions in comparisons between the results obtained with the estimated and conditional distributions. Levels of correlation between different random taste coefficients were negligible, such that in the final estimation, the coefficients were treated as independent.

In terms of model fit, the MMNL model obtains a very significant improvement in log-likelihood (LL) of 187.77 units over the MNL model, at the cost of just 5 additional parameters. As expected, there is a drop in significance levels for almost all parameters when compared with the MNL models, as a result of the inclusion of the individual-specific error component. Nevertheless, all estimates still attain high levels of statistical significance.

---

[6]During the review stage of this paper, both anonymous referees raised concerns about the use of excessive precision in the presentation of the results, specifically with a view to using four decimal places. It seems important to clarify something at this point. It should be clear that the use of four decimal places when working with large values is different from the use of four decimal places when working with small values. The average absolute value for the MNL estimates in Table 1 is only 0.3727. With all coefficients estimated very robustly, the use of four decimal places is in this case warranted. Indeed, by reducing the precision for individual coefficients to three decimal places, the valuation of savings in flight time from the MNL model would reduce from \$16.92/hour to \$15/hour, a drop by more than 11%. A reduction to two decimal places is not even possible, with the estimate for some coefficients being smaller than $1 \cdot 10^{-2}$. The small parameter estimates are a results of the comparatively large values for the associated attributes. An alternative approach would have consisted of rescaling the attribute values by a factor of $\frac{1}{10}$ or $\frac{1}{100}$, which would have enabled the use of a lower level of precision in the parameter estimates, as they would have been rescaled upwards by the same factor. However, the results would have been identical, and preference was given to the approach used here such that the air fare coefficient for example relates to an attribute expressed in dollars and not cents.

		MNL		MMNL	
Null LL:		-2,866.86		-2,866.86	
Final LL:		-1,436.19		-1,248.42	
adj.ρ^2		0.4955		0.5593	

		est.	asy.t-rat.	est.	asy.t-rat.
β_{current}		0.7562	12.07	1.0546	8.11
$\beta_{\text{access time}}$ (min)	μ	-0.0072	-8.07	-0.0112	-6.22
	σ	-	-	0.0085	2.69
$\beta_{\text{air fare}}$ (\$)	μ	-0.0156	-21.47	-0.0330	-15.39
	σ	-	-	0.0165	8.54
$\beta_{\text{flight time}}$ (min)	μ	-0.0044	-8.56	-0.0129	-8.16
	σ	-	-	0.0109	8.93
β_{OTP} (%)		0.0084	6.35	0.0150	5.23
$\beta_{\text{1 connection}}$		-0.6608	-6.78	-0.7197	-3.77
$\beta_{\text{2 connections}}$		-0.8956	-3.74	-1.0981	-2.50
$\beta_{\text{standard FF}}$		0.2254	2.22	0.4827	2.71
$\beta_{\text{elite FF}}$		0.7248	3.08	1.0315	2.51
$\beta_{\text{closest airport}}$	μ	0.4287	5.12	0.9533	5.84
	σ	-	-	0.9099	2.80
	φ	-	-	0.7052	5.51

Table 1: Estimation results for MNL and MMNL models

5 Estimated distributions versus distribution of conditional means

5.1 Distribution of marginal utility coefficients

The mean of the conditional distribution for each of the four random taste coefficients ($\beta_{\text{access time}}$, $\beta_{\text{air fare}}$, $\beta_{\text{flight time}}$ and $\beta_{\text{closest airport}}$) is obtained for each respondent, using the approach set out in Equation 2. Summary statistics across these 517 values are then calculated for each of the four coefficients, and these are compared with the corresponding statistics for the estimated distributions in Table 2.

With the use of the Normal distribution, the range for the estimated distribution is unbounded for all four coefficients. This is not the case for the distribution of the conditional means, and the much narrower range is reflected in the significantly lower standard deviations, especially for $\beta_{\text{access time}}$ and $\beta_{\text{closest airport}}$. The mean, on the other hand, is almost exactly the same whether working with the conditional or estimated distributions. Finally, differences arise in the probability of counter-intuitively signed coefficients. With the unbounded nature of the Normal distribution, we obtain significant probabilities of counter-intuitively

	β	air fare	access time	flight time	closest airport
estimated	mean	-0.0330	-0.0112	-0.0129	0.9533
	std.dev	0.0165	0.0085	0.0109	0.9099
	Prob(<0)	97.71%	90.64%	88.21%	14.74%
	Prob(>0)	2.29%	9.36%	11.79%	85.26%
cond. means	min.	-0.0525	-0.0214	-0.0260	-0.0967
	max.	-0.0059	-0.0052	0.0162	1.7632
	mean	-0.0331	-0.0112	-0.0129	0.9596
	std.dev.	0.0099	0.0023	0.0059	0.2868
	Perc(<0)	100.00%	100.00%	96.71%	0.19%
	Perc(>0)	0.00%	0.00%	3.29%	99.81%

Table 2: Summary statistics for estimated distributions and distributions of conditional means for four randomly distributed taste coefficients

signed values for $\beta_{\text{access time}}$, $\beta_{\text{flight time}}$ and $\beta_{\text{closest airport}}$. But, when working with the conditional means, a non-zero share of counter-intuitively signed coefficient values is only obtained for $\beta_{\text{flight time}}$, where this is however small enough to be ignored. The fact that the distributions of the conditional means thus offer no conclusive evidence of counter-intuitively signed coefficient values supports the argument of Hess *et al.* (2005) that estimates showing high shares of such values are often affected by the distributional assumptions. The findings in the table are reflected in a graphical representation of the cumulative distribution functions for the four coefficients (Figure 2), showing a much narrower range when working with the conditional means, along with a much lower incidence of sign violations.

In the presence of individual-specific draws for each of the four randomly distributed coefficients, it is possible to test for correlation between the distributions of conditional means for these coefficients[7]. The results of this process are summarised in Table 3, that shows the correlations for the various pairs of coefficients, along with the associated p-values. It can be seen that significant levels of correlation exist for all pairs of coefficients, except between $\beta_{\text{access time}}$ and $\beta_{\text{flight time}}$. This is a striking result because no meaningful correlation between the distributions for individual coefficients could be retrieved in the unconditional estimation.

The results show negative correlation between $\beta_{\text{air fare}}$ and $\beta_{\text{access time}}$, indicating that respondents with a higher fare sensitivity have a lower access time sensitivity. This is consistent with intuition, and explains why low-cost airlines are able to attract passengers to outlying airports. A similar situation of negative correlation is found between $\beta_{\text{air fare}}$ and $\beta_{\text{flight time}}$, suggesting that respondents

---

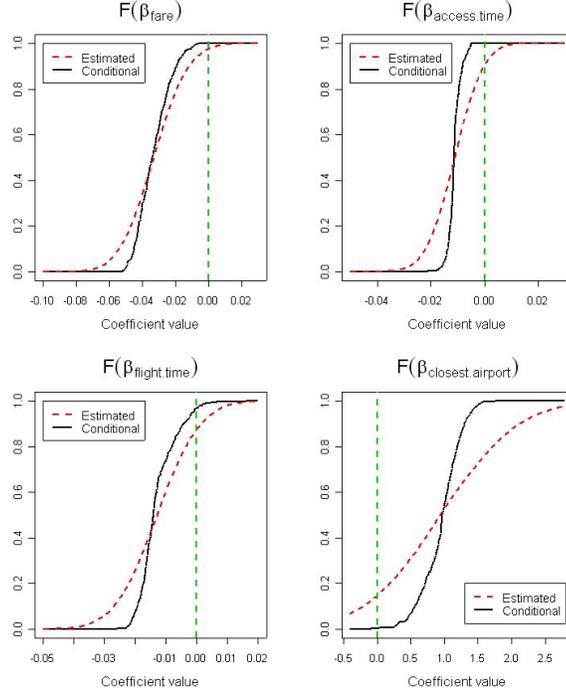[7]The author would like to thank John Rose for this suggestion.

Figure 2: Estimated distributions and distributions of conditional means for four randomly distributed taste coefficients

		$\beta_{\text{access time}}$	$\beta_{\text{air fare}}$	$\beta_{\text{flight time}}$	$\beta_{\text{closest airport}}$
$\beta_{\text{access time}}$	corr.	1	-0.35	0.07	-0.52
	p-val.	-	0.00	0.10	0.00
$\beta_{\text{air fare}}$	corr.	-0.35	1	-0.32	0.43
	p-val.	0.00	-	0.00	0.00
$\beta_{\text{flight time}}$	corr.	0.07	-0.32	1	-0.10
	p-val.	0.10	0.00	-	0.02
$\beta_{\text{closest airport}}$	corr.	-0.52	0.43	-0.10	1
	p-val.	0.00	0.00	0.02	-

Table 3: Correlation in distribution of conditional means of marginal utility coefficients

with greater spending power are more concerned about flight time than respondents with low spending power. For pairwise comparisons involving $\beta_{\text{closest airport}}$, the positive sign needs to be taken into account. As such, the negative correlation between $\beta_{\text{access time}}$ and $\beta_{\text{closest airport}}$ shows that, as expected, respondents with a higher access time sensitivity also have a greater preference for the closest

	mean.	std.dev.	P(>0)	P(<0)
$\beta_{\text{access time}}$ vs $\beta_{\text{air fare}}$	17.53	891.56	88.66%	11.34%
$\beta_{\text{flight time}}$ vs $\beta_{\text{air fare}}$	43.19	1,296.74	86.82%	13.18%
$\beta_{\text{1 connection}}$ vs $\beta_{\text{air fare}}$	36.01	1,098.26	97.74%	2.26%
$\beta_{\text{2 connections}}$ vs $\beta_{\text{air fare}}$	54.94	1,675.76	97.74%	2.26%
$\beta_{\text{closest airport}}$ vs $\beta_{\text{air fare}}$	44.61	1,722.63	83.63%	16.37%
$\beta_{\text{standard FF}}$ vs $\beta_{\text{air fare}}$	24.15	736.55	97.74%	2.26%
$\beta_{\text{elite FF}}$ vs $\beta_{\text{air fare}}$	51.61	1,574.12	97.74%	2.26%
β_{OTP} vs $\beta_{\text{air fare}}$	75.08	2,290.00	97.74%	2.26%
$\beta_{\text{air fare}}$ vs $\beta_{\text{access time}}$	-2.62	476.20	88.66%	11.34%
$\beta_{\text{standard FF}}$ vs $\beta_{\text{access time}}$	-60.97	6,922.34	90.58%	9.42%
$\beta_{\text{elite FF}}$ vs $\beta_{\text{access time}}$	-130.30	14,794.14	90.58%	9.42%
$\beta_{\text{1 connection}}$ vs $\beta_{\text{access time}}$	-90.91	10,321.76	90.58%	9.42%
$\beta_{\text{2 connections}}$ vs $\beta_{\text{access time}}$	-138.70	15,749.34	90.58%	9.42%

Table 4: Summary statistics for estimated distributions for various trade-offs

airport. Similarly, the positive correlation between $\beta_{\text{air fare}}$ and $\beta_{\text{closest airport}}$ indicates that respondents with more spending power have a greater preference for the closest airport. Finally, there is low negative correlation between $\beta_{\text{flight time}}$ and $\beta_{\text{closest airport}}$, suggesting that respondents with a higher flight time sensitivity have a greater preference for the closest airport (i.e., higher access time sensitivity).

5.2 Distribution of trade-offs

We now move on to the distribution of trade-offs between the various coefficients. Thirteen trade-offs are used, looking at the willingness to accept increases in air fare or access time in return for improvements along some other dimension. The results of this calculation are shown in Table 4 for the estimated distributions, and Table 5 for the distribution of the conditional means. When working with the estimated distributions, the ratios are calculated using simulation over $100,000$ independent draws from the various distributions.

A question arises in the simulation of the ratios over draws produced from the Normal distribution. Indeed, the infinite range of the distribution means that positive as well as negative numbers are used in the simulation, while values very close to zero are also included. The latter leads to problems for the coefficient used in the denominator, resulting in extreme values and an inflated range for the distribution of the ratio. A possible approach is to censor the distribution of the coefficients by removing the upper and lower percentile points, to guarantee a unique sign and the absence of values very close to zero (Hensher and Greene, 2003). This approach is however not only very arbitrary, but also leads to a loss of

	mean	std.dev.	Perc(> 0)	Perc(< 0)
$\beta_{\text{access time}}$ vs $\beta_{\text{air fare}}$	23.81	14.51	100.00%	0.00%
$\beta_{\text{flight time}}$ vs $\beta_{\text{air fare}}$	28.07	22.45	96.71%	3.29%
$\beta_{\text{1 connection}}$ vs $\beta_{\text{air fare}}$	24.96	12.81	100.00%	0.00%
$\beta_{\text{2 connections}}$ vs $\beta_{\text{air fare}}$	38.09	19.55	100.00%	0.00%
$\beta_{\text{closest airport}}$ vs $\beta_{\text{air fare}}$	34.50	22.25	99.81%	0.19%
$\beta_{\text{standard FF}}$ vs $\beta_{\text{air fare}}$	16.74	8.59	100.00%	0.00%
$\beta_{\text{elite FF}}$ vs $\beta_{\text{air fare}}$	35.78	18.36	100.00%	0.00%
β_{OTP} vs $\beta_{\text{air fare}}$	52.05	26.71	100.00%	0.00%
$\beta_{\text{air fare}}$ vs $\beta_{\text{access time}}$	3.18	1.47	100.00%	0.00%
$\beta_{\text{standard FF}}$ vs $\beta_{\text{access time}}$	45.45	11.92	100.00%	0.00%
$\beta_{\text{elite FF}}$ vs $\beta_{\text{access time}}$	97.13	25.47	100.00%	0.00%
$\beta_{\text{1 connection}}$ vs $\beta_{\text{access time}}$	67.77	17.77	100.00%	0.00%
$\beta_{\text{2 connections}}$ vs $\beta_{\text{access time}}$	103.40	27.11	100.00%	0.00%

Table 5: Summary statistics for distributions for various trade-offs using conditional means

information by artificially reducing the true standard deviation of the individual coefficients. While acceptable in the case of just a few percentile points, such as for $\beta_{\text{air fare}}$, the need to possibly remove over 10 percentile points from both ends of the distribution for $\beta_{\text{access time}}$, $\beta_{\text{flight time}}$ and $\beta_{\text{closest airport}}$ to maintain balance means the resulting trade-offs underestimate the variation to such an extent that they are themselves unreliable. As such, all draws were included in the simulation.

The effects of the extreme values on the simulation of the trade-offs are clearly visible from Table 4, with a hugely inflated range for the trade-offs when compared with the ratios obtained with the conditional means (Table 5). Furthermore, there are issues in terms of sign violations, especially for WTP measures involving a random numerator, or trade-offs using the access time coefficient in the denominator. The biggest problems arise for the willingness to accept increases in access time. Here, the inclusion of positive draws for $\beta_{\text{access time}}$, along with draws close to zero, distorts the distribution of the trade-off to such an extent that the mean values of the trade-offs are negative even though around 90% is in fact positive.

The above discussion highlights the problems that arise when working with the estimated distributions as opposed to making use of means from the conditional distributions. The latter problems only arise for a single trade-off, namely when looking at the WTP for reductions in flight time, where there is a 3.3% probability of a negative WTP. The 0.2% probability of a negative WTP for using the closest airport is negligible.

Comparing the results from Table 4 and Table 5, we can see that the differ-
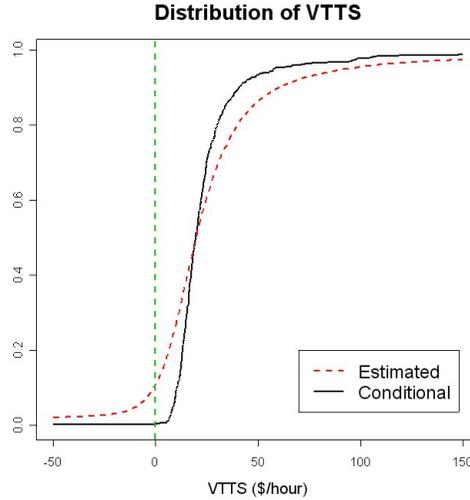
Distribution of VTTS

Figure 3: Distribution of the valuation of travel time savings (VTTS) with unconditional estimates and conditional mean estimates

ences are not restricted to sign violations or the range of the distribution, but do also extend to the mean values. As a further illustration of the differences between the approaches using the estimated distributions and conditional means, Figure 3 compares the cumulative distribution functions for the WTP for reductions in access time. The plot not only shows the much narrower range when making use of the conditional means, but also highlights the sign violations resulting from using the estimated distributions.

The evidence thus far shows that trade-offs calculated on the basis of conditional means are more reliable, even though some risk of bias remains. This approach is adopted for the remainder of the paper. As a first step, we look at a graphical representation of the distribution of the trade-offs, with density functions for the various WTP measures in the first two rows in Figure 4, while the last row in the figure shows the density functions for the various trade-offs using access time in the denominator. For all 13 trade-offs, there is clear evidence of significant levels of heterogeneity across respondents. All trade-offs have a longer tail to the right, where the degree of asymmetry varies across trade-offs.

In a direct comparison of related trade-offs, we observe a slightly higher WTP for reductions in flight time than for reductions in access time, along with a longer tail to the right. Similarly, there is a higher WTP for a reduction from two to no connections, than the corresponding reduction from one connection to a direct flight, along with a longer tail. The ratio however is not of the order of 2 : 1,
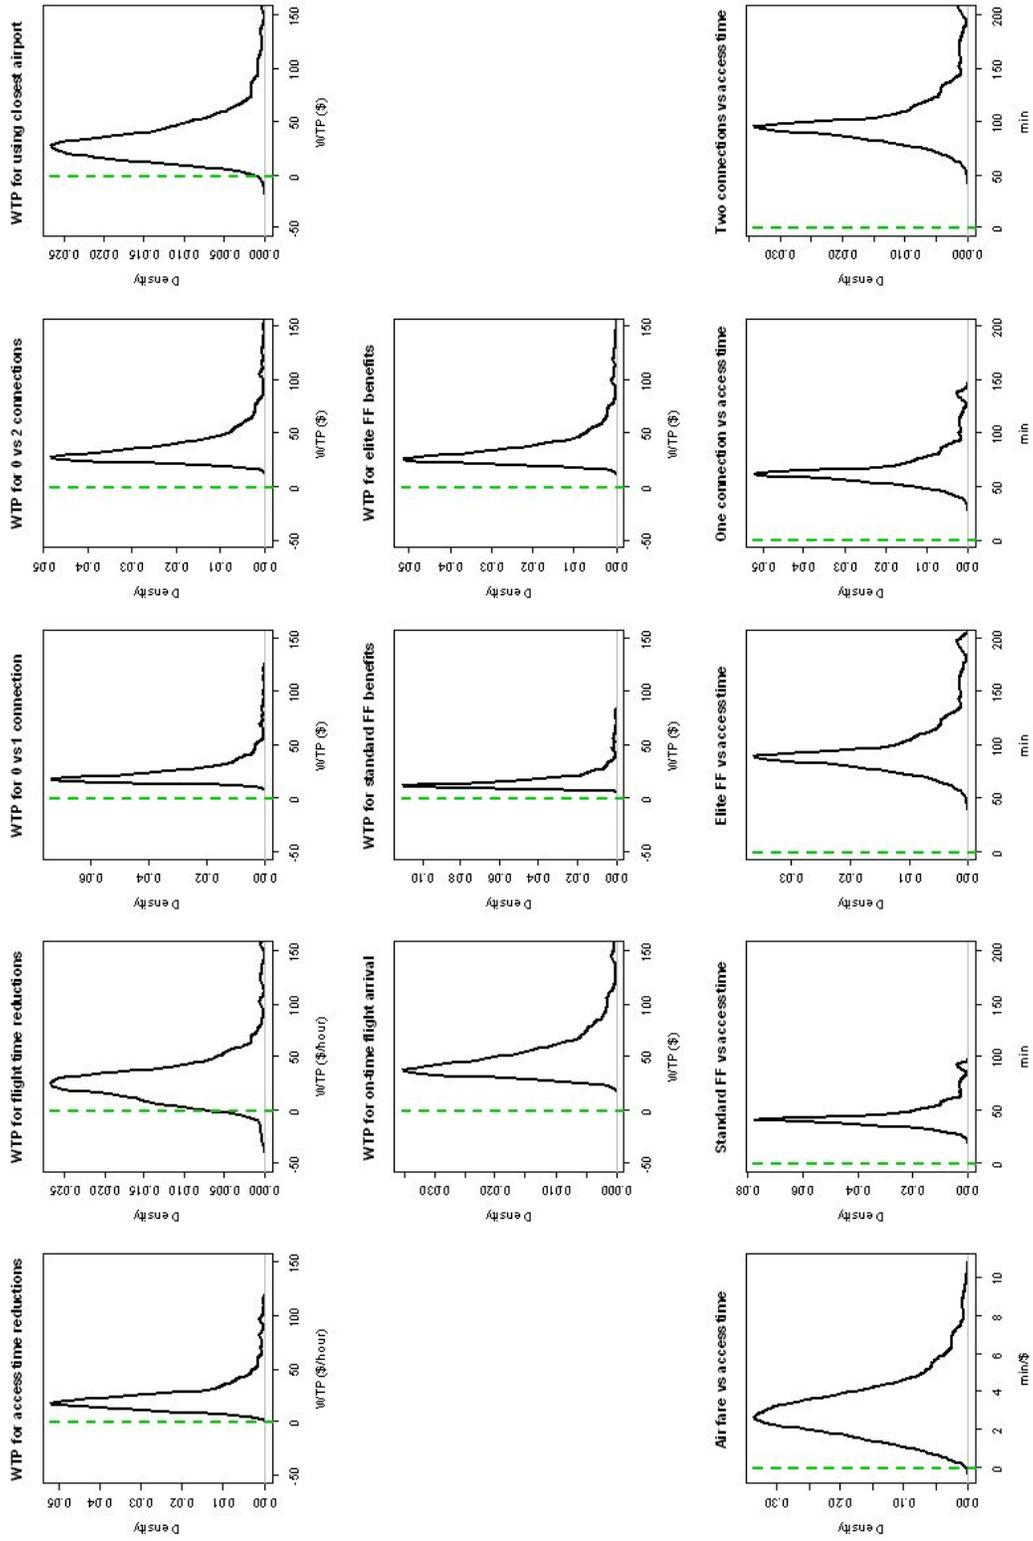
13

Figure 4: Density functions for trade-offs calculated with conditional means

14

		Business	Vacation	VFR[i]
$\beta_{\text{access time}}$ *vs* $\beta_{\text{air fare}}$ (\$/hour)	min.	6.49	6.63	6.05
	max.	113.70	63.17	99.17
	mean	29.69	21.62	23.01
	std.dev.	21.90	9.79	12.98

[i]VFR = visiting friends or relatives

Table 6: Distribution of VTTS by purpose segment

justifying the use of two separate coefficients for the two levels of connection. Finally, there is a higher WTP for elite FF benefits than for standard FF benefits, along with a longer tail for the former.

When looking at the willingness to accept increases in access time, we observe a significant level of heterogeneity for the willingness to travel to more distant airports in return for lower fares, with travellers at the upper end of the distribution being part of the target market for low-cost airlines operating from regional airports. Respondents are also willing to travel to more outlying airports in return for frequent flier benefits, where the willingness is greatest for elite members, along with greater variation. The same applies in the case of respondents' willingness to travel to outlying airports in return for direct flights.

For reasons of interpretation, it is preferable to link the variations in tastes and choice behaviour to socio-demographic attributes of the respondents. Extensive attempts were made to establish a link between the distribution of the various trade-offs (as well as individual coefficients) and attributes such as trip purpose, respondents' income and trip distance. Initially, regression models were estimated; however, while they indicate some form of interaction, the levels of significance of the estimated parameters are below any reasonable level of confidence. The most convincing results are obtained by a very basic posterior analysis that segmented the population according to trip purpose. The main findings show a higher access time sensitivity for business travellers, along with a lower air fare sensitivity. Table 6 shows the differences between the three population segments for the distribution of the VTTS. Along with a higher mean value, there is a greater spread in the business segment, with the narrowest distribution being for holiday travellers.

5.3 Reestimation of models with conditional means

A reestimation of the models imports the conditional means for the four random taste coefficients[8], so that only the remaining six coefficients, along with φ, are

---

[8]I.e., using the most likely values of the coefficients for each respondent.

estimated[9]. To allow for the potential differences in scale between the original model and the reestimated model, a rescaling parameter λ was associated with the imported coefficients. The utility function for the RP alternative in the reestimated model is given by:

$$
\begin{aligned}
U_{\mathrm{RP}} = {} & \beta_{\mathrm{current}} + \beta_{\mathrm{OTP}} \cdot \mathrm{OTP}_{\mathrm{RP}} + \beta_{1\ \mathrm{connection}} \cdot \delta_{1\ \mathrm{connection,RP}} + \beta_{2\ \mathrm{connections}} \cdot \delta_{2\ \mathrm{connections,RP}} \\
& + \beta_{\mathrm{standard\ FF}} \cdot \delta_{\mathrm{standard\ FF,RP}} + \beta_{\mathrm{elite\ FF}} \cdot \delta_{\mathrm{elite\ FF,RP}} + \lambda \cdot \tilde{\beta}_{\mathrm{access\ time}} \cdot \mathrm{access\ time}_{\mathrm{RP}} \\
& + \lambda \cdot \tilde{\beta}_{\mathrm{flight\ time}} \cdot \mathrm{flight\ time}_{\mathrm{RP}} + \lambda \cdot \tilde{\beta}_{\mathrm{air\ fare}} \cdot \mathrm{air\ fare}_{\mathrm{RP}} + \lambda \cdot \tilde{\beta}_{\mathrm{closest\ airport}} \cdot \delta_{\mathrm{closest\ airport,RP}}
\end{aligned}
\tag{5}
$$

The individual-specific values for $\tilde{\beta}_{\mathrm{access\ time}}, \tilde{\beta}_{\mathrm{air\ fare}}, \tilde{\beta}_{\mathrm{flight\ time}}$ and $\tilde{\beta}_{\mathrm{closest\ airport}}$ are imported, while λ is estimated in addition to $\beta_{\mathrm{current}}, \beta_{\mathrm{OTP}}, \beta_{1\ \mathrm{connection}}, \beta_{2\ \mathrm{connections}}, \beta_{\mathrm{standard\ FF}}$ and $\beta_{\mathrm{elite\ FF}}$.

The results are summarised in Table 7. When compared with the unconditional model in Table 1, there is a remarkable increase in model fit by 565.82 units, showing the greater level of accuracy that is obtained when working with individual-specific coefficients rather than using integration over the sample-wide distribution. The aim behind this is however not simply to reveal this improvement in model fit, but to test for the potential impact on the remaining coefficients. All six β coefficients remain statistically significant, and still take the correct sign. Similarly, φ is still statistically significant, and the additional scale parameter λ is significantly larger than 1.

When using these new estimates in a recalculation[10] of the trade-offs from Table 5, significant differences are observed for those trade-offs involving at least one fixed coefficient, as a result of the reduction in the estimated values for the five relevant coefficients. Table 8 shows results for the MNL model alongside those obtained with the conditional draws for the simple MMNL model (MMNL$_{\mathrm{C1}}$), and those obtained with the conditional draws for the MMNL model estimated with imported conditional draws (MMNL$_{\mathrm{C2}}$). The differences between MMNL$_{\mathrm{C1}}$ and MMNL$_{\mathrm{C2}}$ are especially significant for the trade-offs involving $\beta_{1\ \mathrm{connection}}$ or $\beta_{2\ \mathrm{connections}}$, where there is a clear decline when compared to the old values. Finally, it can be seen from Table 8 that there are also important differences between the MNL trade-offs and the MMNL trade-offs, showing the effects of not allowing for random taste heterogeneity, where, for example, in the MNL models, and contrary to what is observed with the MMNL models, the WTP for reductions in access time is greater than that for reductions in flight time.

---

[9]There is little gain from using conditional draws for the serial correlation terms $\varphi \xi_{\mathrm{RP}}$ and $\varphi \xi_{\mathrm{SP}}$, with the resulting draws being essentially no different from a value of zero.

[10]The recalculation also accounts for the differences in scale retrieved through the estimation of λ.

	est.	asy.t-rat.
Null LL:	-2,866.86	
Final LL:	-682.60	
adj.ρ^2	0.7591	

	est.	asy.t-rat.
δ_{current}	1.0961	11.35
$\beta_{\text{access time}}$	conditionals	
$\beta_{\text{air fare}}$	conditionals	
$\beta_{\text{flight time}}$	conditionals	
β_{OTP}	0.0168	7.21
$\beta_{1 \text{ connection}}$	-0.5342	-4.51
$\beta_{2 \text{ connections}}$	-0.7745	-2.30
$\beta_{\text{elite FF}}$	1.0843	3.02
$\beta_{\text{standard FF}}$	0.4873	3.10
$\beta_{\text{closest airport}}$	conditionals	
φ	0.3634	3.21
λ	1.3449	21.43

Table 7: Reestimation of model with imported conditional means

6   Conclusions

This paper has discussed the issue of the computation of trade-offs such as willing-ness to pay indicators in the analysis of air travel choice behaviour. Specifically, it looked at the scenario where some taste coefficients follow a random distribution. The analysis confirms that when using simulation of these ratios over random draws from the appropriate distributions, the presence of extreme values in the draws can lead to biased estimates of the mean and spread of the trade-offs. The problems caused by these outliers are so severe that the resulting trade-offs have little practical use for policy making.

The study shows that the problems with biased trade-offs largely disappear when calculating the ratios on the basis of individual-specific coefficient values conditioned on given travellers' observed choices. These findings are consistent with those of Greene *et al.* (2005). Furthermore, unlike with the unconditional distributions, it was possible to retrieve significant and meaningful patterns of correlation between the distributions of conditional means for the four randomly distributed coefficients.

	MNL	MMNL$_{C1}$	MMNL$_{C2}$
$\beta_{\text{access time}}$ *vs* $\beta_{\text{air fare}}$ (\$/hour)	27.57	23.81 (14.51)	
$\beta_{\text{flight time}}$ *vs* $\beta_{\text{air fare}}$ (\$/hour)	16.89	28.07 (22.45)	
$\beta_{\text{1 connection}}$ *vs* $\beta_{\text{air fare}}$ (\$)	42.46	24.96 (12.81)	13.78 (7.07)
$\beta_{\text{2 connections}}$ *vs* $\beta_{\text{air fare}}$ (\$)	57.54	38.09 (19.55)	19.98 (10.25)
$\beta_{\text{closest airport}}$ *vs* $\beta_{\text{air fare}}$ (\$)	27.54	34.5 (22.25)	
$\beta_{\text{standard FF}}$ *vs* $\beta_{\text{air fare}}$ (\$)	14.48	16.74 (8.59)	12.57 (6.45)
$\beta_{\text{elite FF}}$ *vs* $\beta_{\text{air fare}}$ (\$)	46.57	35.78 (18.36)	27.96 (14.35)
β_{OTP} *vs* $\beta_{\text{air fare}}$ (\$)	54.14	52.05 (26.71)	43.32 (22.23)
$\beta_{\text{air fare}}$ *vs* $\beta_{\text{access time}}$ (min/\$)	2.18	3.18 (1.47)	
$\beta_{\text{standard FF}}$ *vs* $\beta_{\text{access time}}$ (min)	31.52	45.45 (11.92)	34.12 (8.95)
$\beta_{\text{elite FF}}$ *vs* $\beta_{\text{access time}}$ (min)	101.36	97.13 (25.47)	75.92 (19.91)
$\beta_{\text{1 connection}}$ *vs* $\beta_{\text{access time}}$ (min)	92.41	67.77 (17.77)	37.4 (9.81)
$\beta_{\text{2 connections}}$ *vs* $\beta_{\text{access time}}$ (min)	125.25	103.4 (27.11)	54.23 (14.22)

Table 8: Comparison of trade-offs between models (std. dev. in brackets)

References

Adler, T., C. S. Falzarano and G. Spitz (2005) *Modeling Service Trade-offs in Air Itinerary Choices*, paper presented at the 84[th] Annual Meeting of the Transportation Research Board, Washington, DC.

Algers, S. and M. Beser (2001) Modelling choice of flight and booking class - a study using stated preference and revealed preference data, *International Journal of Services Technology and Management*, **2**, 28–45.

Basar, G. and C. R. Bhat (2004) A Parameterized Consideration Set model for airport choice: an application to the San Francisco Bay area, *Transportation Research B*, **38**, 889–904.

Bhat, C. R. (2001) Quasi-random maximum simulated likelihood estimation of the mixed multinomial Logit model, *Transportation Research B*, **35**, 677–693.

Bondzio, L. (1996) Modelle für den Zugang von Passagieren zu Flughäfen (Models for the Passengers' Access to Airports), Ph.D. Thesis, Ruhr-University, Bochum.

Bradley, M. A. (1998) Behavioural models of airport choice and air route choice, in J. de D. Ortúzar, D. A. Hensher and S. R. Jara-Diaz (eds.), *Travel behaviour research: updating the state of play*, Elsevier, Oxford.

Chin, A. T. H. (2002) Impact of frequent flyer programs on the demand for air travel, *Journal of Air Transportation*, **7**, 53–86.

Doornik, J. A. (2001) *Ox: An Object-Oriented Matrix Language*, Timberlake Consultants Press, London.

González-Savignat, M. (2004) Competition in air transport: The case of high speed train, *Journal of Transport Economics and Policy*, **38**, 77–108.

Greene, W. H., D. A. Hensher and J. M. Rose (2005) Using classical simulation-based estimators to estimate individual wtp values, in R. Scarpa and A. Alberini (eds.), *Applications of Simulation Methods in Environmental and Resource Economics*, Springer, Dordrecht.

Hensher, D. A. and W. H. Greene (2003) The Mixed Logit Model: The State of Practice, *Transportation*, **30**, 133–176.

Hess, S., T. Adler and J. W. Polak (2007) Modelling airport and airline choice behavior with stated-preference survey data, *Transportation Research E*, **43**, 221–233.

Hess, S., M. Bierlaire and J. W. Polak (2005) Estimation of value of travel-time savings using mixed logit models, *Transportation Research A*, **39**, 221–236.

Hess, S. and J. W. Polak (2005a) Accounting for random taste heterogeneity in airport-choice modelling, *Transportation Research Record*, **1915**, 36–43.

Hess, S. and J. W. Polak (2005b) Mixed logit modelling of airport choice in multi-airport regions, *Journal of Air Transport Management*, **11**, 59–68.

Hess, S. and J. W. Polak (2006a) Airport, airline and access mode choice in the san francisco bay area, *Papers in Regional Science*, **85**, 543–567.

Hess, S. and J. W. Polak (2006b) Exploring the potential for cross-nesting structures in airport-choice analysis: a case-study of the Greater London area, *Transportation Research E*, **42**, 63–81.

Hess, S., K. Train and J. W. Polak (2006) On the use of a Modified Latin Hypercube Sampling (MLHS) method in the estimation of a Mixed Logit model for vehicle choice, *Transportation Research B*, **40**, 147–163.

Louviere, J. J., D. A. Hensher and J. Swait (2000) *Stated Choice Models: Analysis and Application*, Cambridge University Press, Cambridge.

Monteiro, A. B. F. and M. Hansen (1997) Improvements to airport ground access and behavior of multiple airport system: BART extension to San Francisco International Airport, *Transportation Research Record*, **1562**, 38–47.

Ortúzar, J. de D., D. A. Roncagliolo and U. C. Velarde (2000) Interactions and independence in stated preference modelling, in J. de D. Ortúzar (ed.), *Stated Preference Modelling Techniques: PTRC Perspectives 4*, PTRC Education and Research Services Ltd, London.

Pathomsiri, S., H. Mahmassani and A. Haghani (2004) *Airport Choice Model for Baltimore-Washington Region*, paper presented at the 83[rd] Annual Meeting of the Transportation Research Board, Washington, DC.

Pels, E., P. Nijkamp and P. Rietveld (2001) Airport and airline choice in a multi-airport region: an empirical analysis for the San Francisco bay area, *Regional Studies*, **35**, 1–9.

Pels, E., P. Nijkamp and P. Rietveld (2003) Access to and competition between airports: a case study for the San Francisco Bay area, *Transportation Research A*, **37**, 71–83.

Proussaloglou, K. and F. S. Koppelman (1995) Air carrier demand: an analysis of market share determinants, *Transportation*, **22**, 371–388.

Psaraki, V. and C. Abacoumkin (2002) Access mode choice for relocated airports: the new Athens International Airport, *Journal of Air Transport Management*, **8**, 89–98.

Resource Systems Group Inc. (2003) *Air Travelers 2003: The New Realities?*, annual Air Survey project report, White River Junction.

Revelt, D. and K. Train (1999) *Customer-specific taste parameters and Mixed Logit: Households' choice of electricity supplier*, Working Paper No. E00-274, Department of Economics, University of California, Berkeley.

Sillano, M. and J. de D. Ortúzar (2004) Willingness-to-pay estimation from mixed logit models: some new evidence, *Environment & Planning A*, **37**, 525–550.

Train, K. (2003) *Discrete Choice Methods with Simulation*, Cambridge University Press, Cambridge, MA.