Airport choice behaviour: findings from three separate studies

Stephane Hess* John W. Polak†

February 25, 2009

Abstract

This chapter summarises the findings of three parallel studies on air travel choice behaviour conducted by the authors in the past few years. The three studies make use of different datasets, with two studies using Revealed Preference data and one study using Stated Choice data. This chapter has three main findings. Firstly, the results give a clear indication of the advantages of Stated Choice data in retrieving significant effects for a number of conceptually important attributes, such as air fares and frequent flier benefits. Secondly, all three studies show that advanced model structures do lead to better performance but that these gains come at the cost of a very significant rise in estimation cost. Finally, all studies show the importance of airport access time in that passengers have a strong preference for departing from the airport closest to their ground level origin.

1 Introduction

As highlighted in the discussions in ***REFERENCE TO HESS CHAPTER ON THEORY***, the number of studies of air travel choice behaviour has increased significantly over recent years, with a particular focus on studies of airport choice behaviour.

Over the last three years, the authors of this chapter have been involved in three separate such studies, making use of Revealed Preference (RP) data collected in the San Francisco Bay (SF-bay) area and Greater London, and Stated Choice (SC) data collected via an internet-based survey in the US. This chapter summarises the findings of these three case studies, with more details on the individual studies being given in the relevant publications; see Hess & Polak (2005*a*,*b*, 2006*a*,*b*) and Hess et al. (2007).

While the overall approach used in the three studies is very similar, there are some differences between the three datasets that change the scope between case

*Institute for Transport Studies, University of Leeds, s.hess@its.leeds.ac.uk

†Centre for Transport Studies, Imperial College London, j.polak@imperial.ac.uk

studies. Detailed description of the three datasets are presented in the respective sections and related papers; the aim of the discussion that follows is simply to describe the approach taken along the various dimensions of choice with the separate datasets.

Although the discussions in ***REFERENCE TO HESS CHAPTER ON THEORY*** have highlighted the fact that air travellers take decisions along a multitude of dimensions, with potentially important interactions between them, in practice, it is generally only possible to look at a subset of these choice dimensions. This is mainly a reflection of the formidable requirements in terms of data that would arise in the case of a study looking jointly at all the choice dimensions. However, it should also be noted that, for an adequate analysis of the interactions between choice dimensions, dynamic model structures would almost certainly be required, in conjunction with repeated choice data.

The three case studies described in this chapter look at the joint choice of airport and airline, where the two RP studies additionally look at the choice of access mode. While the detailed study and modelling of the interactions between choice dimensions is an important avenue for future research, it should be clear that this is a learning process, and that, before attempting such an analysis, there is a requirement to first look at the joint modelling of even a subset of these choices.

More generic choices, such as the decision to travel, the choice of destination, trip timing, and the decision to travel by air, are not modelled in any of the three case studies, primarily for data reasons. A similar reasoning applies for any unmodelled choices along the air travel and ground level choice dimensions, where the issues discussed in ***REFERENCE TO HESS CHAPTER ON THEORY*** apply. Finally, in the case of flight routing, an important difference arises between the two RP case studies and the SC case study. In the former two, the topic is left untreated, and the studies look only at the choice between direct flights, a decision based on the low share of connecting passengers in the data, and the lack of detailed data on connecting flights. In the SC dataset, connecting flights are included in the choice set, allowing for an analysis of the relative valuations of direct and indirect flights, where the choice of air routing is however not modelled.

The remainder of this chapter is organised as follows. Section 2 looks at the findings of the SF-bay study, while Sections 3 and 4 look at the Greater London and SC studies respectively. Finally, Section 5 briefly summarises and compares the findings of the three studies.

2 San Francisco Bay area study

This section describes the case study conducted in the San Francisco Bay area, which is served by three major airports; San Francisco International (SFO), Metropolitan Oakland International (OAK), and Mineta San José International

Dept. Apt.	Burbank, CA	Chicago, OHare, IL	Dallas, Ft. Worth, TX	Denver, CO	Las Vegas, NV	Los Angeles, CA	Ontario, CA	Orange County, CA	Phoenix, AZ	Portland, OR	Reno, NV	Salt Lake City, UT	San Diego, CA	Seattle, WA	Total
SFO	55	89	36	65	57	199	35	37	128	140	1	42	258	213	1,355
SJC	167	58	91	71	163	367	111	247	133	106	156	61	248	169	2,148
OAK	211	1	25	9	68	381	135	177	51	101	39	43	139	208	1,588
Total	433	148	152	145	288	947	281	461	312	347	196	146	645	590	5,091

Table 1: Summary of choice data for SF-bay area case study

(SJC). There is strong geographical captivity, with each of the three airports being in relatively close proximity to one of the main urban centres in the region, something that applies especially in the case of SJC.

The remainder of this discussion is organised as follows. Section 2.1 describes the data used in the analysis, while Section 2.2 discusses model specification. This is followed by a discussion of the results in Section 2.3 and a model validation process in Section 2.4.

2.1 Description of data

Data on passengers' choice behaviour were obtained from the Airline Passenger Survey conducted by the Metropolitan Transport Commission (MTC) in August and October 1995. This contained information on over 21,000 departing air travellers. Passenger interviews were conducted at the three main SF-Bay area airports, as well as at the minor Sonoma County airport (STS), which was not included in the present study. For a detailed description of the survey, see Franz (1996).

A total of 14 destinations served by direct flights from all three airports on every day of the week during the study period were included in the analysis. This includes some destination airports that are themselves located in multi airport regions; this was unavoidable, given the available data, and issues resulting from this are discussed in more detail in Hess & Polak (2006*a*). The final sample contained 5,091 observations, as summarised in Table 1. Here, SJC is over-sampled, while OAK is under-sampled, an issue taken into account at the estimation stage using weighted maximum likelihood. The sample was split into two parts, a dataset used in the actual analysis (4,582 observations), and a 10% sample retained for later validation of the models (509 observations).

The passenger survey data was complemented by appropriate level-of-service data for the chosen and unchosen alternatives need, using air travel data obtained from BACK Aviation Solutions[1], and ground-access data obtained from the MTC in the form of origin-destination (O-D) travel time and cost matrices for the $1,099$ travel area zones (TAZ) used for the SF-Bay area[2].

In the final analysis, a total of eight airlines were included, while, along the access mode dimension, six modes were used, namely car, public transport, scheduled airport bus services, door-to-door services, taxi and limousine.

2.2 Model specification

The final sample contains data on 3 departure airports, 8 airlines, and 6 access modes, leading to 144 distinct triplets of alternatives. The study uses six population segments, grouping travellers into residents and visitors, as well as dividing them across three purpose groups, namely business, holiday, and visiting friends or relatives (VFR).

For each model, attempts were made to include coefficients showing travellers' sensitivity to various attributes of the airports, airlines and access modes. The set of potential explanators used in the specification search included factors such as flight frequency, flight time (block time, which indirectly takes into account airport congestion), fare and aircraft type (jet vs turboprop), as well as access time (in-vehicle), walk time to access mode (e.g. to public transport station), wait time for access mode and access cost. To account for *airport allegiance*, inertia coefficients were included in the utility functions, associated with variables giving the number of flights that a given traveller had taken out of a certain airport over the last twelve months. In addition to direct effects, cross-effects were allowed for by including inertia coefficients associated with SJC and OAK in the utility of SFO, as well as including a coefficient associated with SFO in the utility of SJC. Some potentially important influences, such as carrier loyalty, could not be explored, due to lack of data (e.g. no information on frequent flyer programmes), while, in the presence of national flights only, the notion of allegiance to the national carrier does not apply. Similarly, it was not possible to identify a significant direct effect of the on-time performance of airlines or airports on the respective choice probabilities.

Both linear and various non-linear specifications of the various explanatory variables were tested. The best results were obtained with the use of a logarithmic transform, this however only led to an improvement in model fit when applied to flight frequency and the inertia variables. Finally, attempts were made to segment the population by income, for example in order to show different values of time in different income classes. Three income groups were defined, segmenting the population into low income ($< \$21,000$ per annum), medium income (between $\$21,000$ and $\$44,000$ per annum) and high income (above $\$44,000$ per annum).

[1] www.backaviation.com

[2] www.mtc.ca.gov/maps_and_data/

2.3 Model results

In this section, we summarise the findings of the model fitting analysis. Here, the emphasis is mainly on the substantive findings, where a detailed description of the results is given in Hess & Polak (2006*a*). We first look at the findings in terms of what attributes have an effect on choice behaviour in the various population segments. This is followed by a discussion of the differences in the results across model structures.

2.3.1 Factors affecting choice behaviour

The estimation on the $1,098$ observations for business trips by residents revealed significant effects of walk access time, access cost, in-vehicle access time, flight time and frequency. Additionally, a significant negative effect could be associated with turboprop flights. No meaningful and significant effect of fare could be identified, even after taking into account income. This can mainly be explained by the poor quality of the fare data, but could also signal indifference to fare increases on the part of business travellers. Finally, a low level of fare differentiation for the routes used in the study also plays a role. It was possible to segment the sensitivity to walk time and access cost by income, where the results show lower sensitivity to cost for people with higher income, along with higher sensitivity to increases in walk time. In terms of the airport inertia variables, the estimates show positive direct effects for all three airports.

The estimation dataset contains information on 1,057 business trips by visitors. Just as in the case of resident business travellers, no significant impact of fares could be identified. In-vehicle access time and access cost are again significant, and negative, with increasing sensitivity to in-vehicle access time with higher income and lower sensitivity to cost with higher income. Whereas it was not possible to estimate a significant effect of wait time for resident business travellers, a significant negative effect could be identified for their non-resident counterparts. However, the estimate for flight time is no longer significant at the conventional 95% level, and it was not possible to include an effect of aircraft type, as flights using turboprop planes were never chosen. Also, with this model, no effect could be associated with access walk time, while frequency again has a strong positive effect.

The model estimated on the 831 observations for residents' holiday trips suggests a lower utility for flights using turboprop aircraft, negative impacts by access cost and in-vehicle time, and a positive effect of flight frequency. Finally, for this group of travellers, a negative effect could be identified for fare (although of lower statistical significance) while no effect could be associated with flight time and access walk time.

For the 534 visitors on holiday trips, no significant effect of fare could be identified, and the effect of access cost, although of the correct sign, is not significant at the 95% level. In-vehicle time has a significant negative effect, as has

	Resident			Visitor		
	Business	Holiday	VFR	Business	Holiday	VFR
MNL	0.5861	0.5112	0.5046	0.4379	0.3725	0.5044
NL (by airport)	0.5873	0.5148	0.5088	0.4481	0.3826	0.5038
NL (by airline)	0.5890	0.5141	0.5105	0.4394	0.3770	0.5018
NL (by access mode)	0.5930	0.5207	0.5224	0.4390	0.3774	0.5149
MMNL	0.5871	0.5127	0.5073	0.4389	0.3766	0.5068

Table 2: Model performance on SF-bay area data in terms of adjusted $\rho^2(0)$ measure

flight time, while increases in frequency lead to increases in utility. Finally, the aircraft type coefficient had to be excluded from the model (turboprop flights never chosen), while no effect could be identified for wait time.

The estimates for the model fitted to the sample of 641 residents on VFR trips show significant negative effects of access cost, in-vehicle time and air fare, along with positive effects of flight frequency. Aircraft type could not be included and no effects could be identified for walk time, wait time and flight time, while segmentations by income led to a loss of information in the model.

The final subsample used in the estimation contains information on 421 VFR trips by non-residents. The results show negative impacts of fare in the medium and low income classes (with higher sensitivity in the low income class), while the effect for high-earners was not significant and was dropped from the model. In-vehicle time and flight time have a negative effect, with a positive effect for frequency increases. No effect could be associated with access walk time, wait time, and access cost, and the turboprop coefficient had to be excluded.

2.3.2 Model structure

We next turn our attention to the findings in terms of model structure. In addition to simple Multinomial Logit (MNL) models, the analysis also looked at the estimation of Nested Logit (NL) and Mixed Multinomial Logit (MMNL) structures. For the NL models, only the three possible two-level structures (nesting by airport, airline or access mode) led to satisfactory results, while, for the MMNL models, the random structure of the model was used solely for expressing random taste heterogeneity across respondents[3], and not heteroscedasticity or inter-alternative correlation. In this discussion, we look mainly at the differences in terms of model performance, using the adjusted ρ^2 measure which takes into account differences in the number of parameters. More detailed results are again given in Hess & Polak (2006*a*), while the differences are also highlighted in the comparison of substantive results which follows. The final specification of the utility function was the same across model structures (MNL, NL and MMNL).

The results in terms of model performance are summarised in Table 2. The

[3] Using the Normal distributions for reasons of simplicity.

results show that, with the exception of the first two NL models for non-resident VFR travellers, in each of the six population segments, the advanced model structures outperform the basic MNL model in terms of the adjusted ρ^2 measure. In the three population segments for residents, the best performance is obtained by the NL model using nesting by access mode, while, for non-residents, a less consistent pattern is observed. Overall however, the differences in model performance are very small, especially when taking into account the heightened estimation cost of the more advanced models. This is especially true for the MMNL model. Two very distinct reasons can be given for the relative lack of improvement in model performance obtained with the more advanced models. The advanced models differ from the MNL model in that they explain processes taking part in the unobserved part of utility, either in the form of inter-alternative correlation or in terms of random taste heterogeneity across respondents. The first explanation interprets the similarity in the performance of the various models as an endorsement of the MNL models, suggesting that the (observed) utility specification used in the models captures almost all of the behaviour, reducing the scope for the advanced model to capture any behaviour in the remaining unobserved part of utility. An alternative explanation is based on the reasoning that the advanced structures used are little better that the MNL model in capturing the true phenomena captured by the unobserved component of utility. It is not clear from the empirical results alone which of these potential explanations is most appropriate, especially given the inability to estimate multi-level NL structures or CNL structures (c.f. Hess 2004). The limited success of the MMNL models is also shown by the small amount of random taste heterogeneity retrieved by these models (c.f. Hess & Polak 2006*a*). As such, while the models for resident business travellers and non-resident VFR travellers manage to retrieve random variations in the sensitivity to four attributes, this reduces to a single attribute for the two remaining resident models and the model for non-resident holiday travellers.

Even though the differences in model fit between the five structures are relatively modest, it is conceivable that the actual substantive results vary more significantly across models. To illustrate this, a brief analysis was conducted with the aim of comparing a common trade-off across models, as well as across population subgroups. The only two coefficients that are included in every single model are frequency and in-vehicle time, allowing the computation of the willingness to accept increases in access time in return for increases in flight frequency. The results of this analysis are summarised in Table 3, where, given that frequency enters the utility under a log-transform, K is used to represent the inverse of the current frequency, and where a sign change has been used to represent the willingness to accept *increases* in access time in return for *increases* in frequency. For the closed form models, the trade-off is simply given by the ratio of the two point-estimates, multiplied by K. The same applies in the MMNL model for holiday trips by visitors, where both frequency and access time were treated as fixed coefficients. However, in the remaining five population segments, the variation in the coefficients needs to be taken into account, especially for

MNL

	Resident			Visitor		
	Business	Holiday	VFR	Business	Holiday	VFR
common	$25.28K$	$22.3K$	$29.47K$	-	$14.01K$	$10.38K$
Inc. $< \$21,000$	-	-	-	$26.32K$	-	-
Inc. $> \$21,000$	-	-	-	$15.93K$	-	-

Nesting by airport

	Resident			Visitor		
	Business	Holiday	VFR	Business	Holiday	VFR
common	$24.27K$	$19.81K$	$26.51K$	-	$12.74K$	$10.25K$
Inc. $< \$21,000$	-	-	-	$21.41K$	-	-
Inc. $> \$21,000$	-	-	-	$13.32K$	-	-

Nesting by airline

	Resident			Visitor		
	Business	Holiday	VFR	Business	Holiday	VFR
common	$26.35K$	$23.59K$	$32.98K$	-	$16.81K$	$11.38K$
Inc. $< \$21,000$	-	-	-	$27.58K$	-	-
Inc. $> \$21,000$	-	-	-	$16.74K$	-	-

Nesting by access mode

	Resident			Visitor		
	Business	Holiday	VFR	Business	Holiday	VFR
common	$20.08K$	$19.05K$	$16.52K$	-	$12.45K$	$7.7K$
Inc. $< \$21,000$	-	-	-	$24.69K$	-	-
Inc. $> \$21,000$	-	-	-	$15.16K$	-	-

MMNL

	Resident			Visitor		
	Business	Holiday	VFR	Business	Holiday	VFR
common	$34.13K$	$25.03K$	$38.78K$	-	$12.5K$	$13.6K$
Inc. $< \$21,000$	-	-	-	$28.63K$	-	-
Inc. $> \$21,000$	-	-	-	$17.63K$	-	-

Table 3: Trade-offs between flight frequency and access time (min/flight) in models for combined choice of airport, airline and access mode ($K = 1/f$, with f giving current frequency)

the access time coefficient, which forms the denominator of the trade-off. Here, a basic simulation approach was used, where the aim was simply to produce an estimate of the mean value for the trade-off, such that a censoring approach could be used to deal with the tails of the distribution, hence avoiding extreme values for the trade-off. Given the use of this censoring approach, the estimated standard deviation for the trade-off is unreliable, such that only the mean values are presented in Table 3.

The first observation that can be made from Table 3 is that, overall, the results show a higher willingness to accept increases in access time for residents than for visitors. The differences are especially significant in the case of VFR trips, where the relative value of frequency increases is at its highest for residents, while it is at its lowest for visitors. In terms of purpose-related differences, the results suggest higher relative sensitivity to frequency for business travellers than for leisure travellers in the models for visitors, while for residents, frequency is valued less highly for holiday travellers. To some extent, these conclusions are however potentially influenced by the quality of the access journey level-of-service data.

Some interesting differences also arise when comparing the results across model structures. Here, it is important to note that no major issues with parameter significance arose in any of the models for the coefficients used in the trade-offs, increasing the reliability of the comparisons. The results show that, although there is some overall consistency in the trade-offs, there are also some differences, for example when looking at the results for MMNL, which overall give greater weight to the frequency coefficient, highlighting the effects of accounting for random taste heterogeneity. However, there are also some differences between the MNL model and the NL models, and also across NL models, such as for example in the case of the model using nesting by access mode for resident VFR travellers. Overall, these findings show that, although the differences between models in terms of LL may be relatively modest, the actual substantive conclusions are quite different. This highlights a relative flatness of the LL function, but also shows the impact of model structure on substantive results, making it an important issue in policy-oriented research.

In closing, it is worth noting that, at the average observed flight frequency of 10 flights, the resulting value of K (0.1) leads to very low willingness to accept access time increases in return for increases in flight frequency. This should however be put into context by noting that the average observed access time was only around 30 minutes. Finally, the high values for K in the case of routes with low frequency (e.g. at $f = 2$, $K = 0.5$) imply a willingness to accept significant increases in access time in return for increases in flight frequency on routes with big gaps between individual flight departures.

2.4 Model validation

To complete the analysis, the five sets of models were applied to the validation sample of 519 observations. For each of the models, the final coefficient values produced during the estimation process were used in the apply runs. On the basis of this, the validation approach produces, for every observation, a choice probability for each of the 144 triplets of alternatives; this can be used to calculate the average probability of correct prediction for the actual chosen alternative across respondents. Aside from this probability for the choice of the actual triplet of airport, airline and access mode, it is also of interest to look at the probability of correct choice for just the airport, just the airline, and just the access mode. These probabilities can be obtained through summing the probabilities of the single elementary alternatives falling into the given group.

MNL

	Resident			Visitor		
	Business	Holiday	VFR	Business	Holiday	VFR
Choice	47.13%	30.56%	36.58%	34.33%	27.21%	36.83%
Airport	84.04%	69.58%	80.83%	70.69%	69.53%	73.20%
Access mode	84.04%	67.72%	66.47%	70.18%	63.22%	77.08%
Airline	60.68%	54.93%	60.26%	55.39%	53.31%	60.97%

Nesting by airport

	Resident			Visitor		
	Business	Holiday	VFR	Business	Holiday	VFR
Choice	48.02%	31.39%	36.74%	36.19%	28.97%	36.81%
Airport	83.69%	69.16%	80.07%	70.69%	68.51%	73.13%
Access mode	85.22%	68.91%	67.50%	72.39%	66.41%	77.25%
Airline	61.06%	55.03%	60.08%	55.90%	54.34%	60.73%

Nesting by airline

	Resident			Visitor		
	Business	Holiday	VFR	Business	Holiday	VFR
Choice	47.90%	31.82%	36.50%	35.00%	27.78%	36.93%
Airport	84.18%	70.24%	80.36%	71.21%	68.61%	73.26%
Access mode	84.92%	68.64%	67.26%	71.08%	64.24%	76.96%
Airline	60.30%	54.79%	59.41%	55.27%	51.60%	60.52%

Nesting by access mode

	Resident			Visitor		
	Business	Holiday	VFR	Business	Holiday	VFR
Choice	48.41%	31.38%	39.60%	34.65%	27.78%	37.83%
Airport	85.39%	70.98%	84.97%	71.11%	72.41%	74.46%
Access mode	83.76%	67.29%	66.16%	70.25%	62.11%	76.98%
Airline	61.33%	55.46%	61.36%	55.49%	53.49%	61.04%

MMNL

	Resident			Visitor		
	Business	Holiday	VFR	Business	Holiday	VFR
Choice	47.53%	30.89%	37.45%	34.52%	27.73%	37.59%
Airport	84.21%	69.60%	81.59%	70.71%	68.78%	73.81%
Access mode	84.09%	67.77%	66.80%	70.21%	63.18%	77.15%
Airline	60.85%	55.07%	60.77%	55.52%	53.08%	60.69%

Table 4: Prediction performance on SF-bay area validation data

The results of the validation process are summarised in Table 4. The first observation that can be made from this table is the surprisingly high probability of correct prediction of the actual chosen alternative, where, with an average of 31 available alternatives, the correct prediction probabilities range from 27% to 48%.

In terms of the correct prediction of airport choice, the probabilities range from 68.51% to as high as 85.39%. This performance compares well with the results of other studies, where, as an example, in one of the more recent studies in the SF-bay area, Basar & Bhat (2004) obtain an average correct prediction rate of 72.9% on their validation sample.

The performance in terms of the choice of access mode is also very good, although it is significantly lower than the performance along the airport dimension for residents on VFR trips, while it is also slightly lower for visiting holiday travellers. On the other hand, it is marginally better than the performance along the airport dimension for visiting VFR travellers. The variation in performance does suggest that the choice process is less deterministic in some segments than in others.

The performance of the models in terms of correctly predicting the choice of airline is poorer than that for the choice of airport and access mode; however the values still always exceed 50%, despite the complete absence of a treatment of airline allegiance, and the low quality of the fare data.

In terms of differences between population segments, the best average performance across all choice dimensions is obtained for resident business travellers, where an argument can be made that such travellers behave in a more rational manner (from the modeller's point of view), due to better information. The comparatively poor performance of the models for holiday trips can partly be explained by the fact that at least some of the travellers on such trips have purchased a package holiday or special flight deal, where the choices of departure airport and airline are potentially influenced by factors not included in the models.

Given the relatively modest differences in performance between the five model structures in the actual estimation processes, it should not be surprising that there are no systematic differences in prediction performance on the validation sample. Even though there are some outliers, such as the performance of the NL model using nesting by access mode in the models for residents on VFR trips, the average differences in performance are too small to come to any conclusions in terms of advantages for one of the model structures. This is further reinforced by the fact that it is not directly clear what measure of error should be associated with these probabilities.

3 Greater London study

This section describes the case study conducted in Greater London, an area which has by far the highest levels of air traffic in Europe, with, in 2002, some 117.13 million passengers using the five main airports. The area is dominated by Heathrow (LHR), the world's busiest international airport, and the main hub in Europe. Additionally, a large number of routes are offered from Gatwick (LGW), the world's busiest single-runway airport, while Stansted (STN) and Luton (LTN) act mainly as bases for holiday and low-cost operators. Finally, the centrally located London City (LCY) airport caters primarily to business travellers, and, due to its short runway, is restricted to short haul flights operated by turboprop planes and small jet aircraft.

3.1 Description of data

For the present analysis, data from the 1996 passenger survey were obtained from the Civil Aviation Authority (CAA 1996). The original sample obtained from the CAA contained responses from 47,831 passengers, for 31 destinations (reachable by direct flights from at least two of the five London airports), and 37 airlines. After data cleaning (missing data, compatibility between datasets), a usable sample of 33,527 passengers was obtained. This compares favourably to the sample of 5,091 available for the SF-bay analysis. In the present discussion, we look only at resident business travellers, where 5% of the 8,704 observations were removed from the original sample to use in later model validation. A total of 31 destinations were used in the analysis, all of which are served by a single main airport, avoiding the problem with multi-airport destination areas faced in the San Francisco study. Any competition with high speed rail on short haul routes is not taken into account in the present study, where we work on the basis of an a priori choice of main mode.

Air-side level-of-service data were again obtained from BACK aviation. The main item of information missing from this dataset is that of the fares for the different routes and airlines. Such data were compiled from two sources; the International Passenger Survey (ONS 1996) and the fare supplement of the Official Airways Guide for 1996 (OAG 1996). Information on travel class as well as ticket type (single or return) was taken into account in assembling the data. As was the case with the fare data used in the SF-bay study, the resulting dataset is of highly aggregate nature, leading to similar problems in the estimation of the marginal utility of air fares. Again, no information is available on frequent flier programmes.

For the analysis of the ground-level choice dimension, data from the National Airport Access Model (NAAM) were obtained for the base year 1999 (Scott Wilson Kirkpatrick 1999), and corresponding cost information for 1996 was produced with the help of the retail price index, while assuming that relative travel times have on average stayed constant. Six different modes are considered in the analysis; private car, rental car, public transport (rail, bus, local transport), long distance coach, taxi and minicab (MC). No combinations of modes were considered in the present analysis, and the final mode indicated in the survey was used as the chosen mode. This is a major simplification of the actual choice process, given the high incidence of access journeys using a combination of different modes. However, in the absence of detailed route choice information, this simplification was not avoidable.

3.2 Model structure and specification

With the use of 5 departure airports, 37 airlines, and 6 access modes, a total of 1,110 combinations of airports, airlines and access modes arise. However, not all airlines operate from all airports, and the total number of airport-airline pairs is actually 54, which reduces the number of alternatives (airport, airline, access mode triplets) to 324, compared to 144 in the SF-bay area study. The number of available alternatives for specific individuals in the estimation sample ranges from 6 to 58, with a mean of 31. As was the case in the SF-bay area study, weights were again used in the specification of the log-likelihood function, to account for the quota used in data collection, which are not representative of the population level.

A comprehensive list of variables was used in the initial modelling analysis, including attributes relating to the air journey, such as frequency, fare, flight time, aircraft type, and seat capacity, and attributes relating to the access journey, such as access cost, in-vehicle access time, out-of-vehicle access time, wait time, number of interchanges, and parking cost. The analysis made use of three separate model structures, namely MNL, NL and CNL, where, for NL models, it was again only possible to estimate the three different two-level structures, with any multi-level structures either failing to converge or reducing to a two-level structure.

3.3 Model results

The modelling analysis showed that only a small set of the attributes listed above have a statistically significant impact on choice behaviour, at least with the present sample and model specification. Indeed, no effect could be identified for parking cost, seat capacity, out-of-vehicle access time, wait time, and the number of interchanges. Furthermore, aircraft size, in the form of a dummy for turboprop planes, showed no effect; here however, the highly correlated flight time attribute had a significant negative effect. A significant effect of air fare could not be identified for the present population segment; this is again at least partly due to the poor quality of the data. Finally, the analysis showed that the use of the combined fuel and depreciation cost for car journeys is preferable to the use of fuel cost on its own. For resident business travellers, four attributes were thus found to have a significant effect; access cost, in-vehicle access time

Model	Final LL	Parameters	Adjusted ρ^2
MNL	-14,945.3	55	0.3445
NL by airport	-14,896.1	59	0.3465
NL by airline	-14,870.7	74	0.3469
NL by access mode	-14,816.7	60	0.3499
CNL	-14,603.9	91	0.3578

Table 5: Model performance on London data

(IVT), flight frequency, and flight time, where a log-transform was used for all four attributes, and where the list of significant attributes stays identical across model structures.

Before looking at the substantive results, we first look at the differences in performance across the five structures used in the analysis, as shown in Table 5. This shows that all three NL models outperform the MNL model, with the best performance offered by the model using nesting by access mode. The same is true for the CNL model, which also outperforms all three NL structures, and where the total improvement of the CNL model over the MNL model is bigger than the combined improvements in the adjusted ρ^2 measure for the three NL models. This highlights the fact that the CNL model indeed offers significant improvements over the NL models, and suggests that the combined analysis of the correlation structure along the three choice-dimensions can offer great benefits.

Aside from the differences in model fit, it is again of interest to look at the differences across models in terms of substantive results. Given the use of the log-transform in nominators as well as denominators, the various trade-offs were calculated separately for each individual[4], and summary statistics were then calculated across respondents.

The results of this process are summarised in Table 6. In each case, the tables present the mean value of the respective trade-off, along with the associated standard deviation, and show the values across the five different models estimated. From these results, it can be seen that the first three models produce roughly similar results, while those produced by the CNL model and the NL model using nesting access mode are more extreme (when compared to the three first models), something that is especially true when looking at the value of travel time savings (VTTS) measures for the model using nesting by access mode.

Another observation that can be made for the trade-offs is that the VTTS measures are markedly lower than those reported for example by Pels et al. (2003), although they are still higher than in other contexts, which can be explained partly by concepts of risk-averseness, as discussed for example by Hess & Polak (2005*b*). Travellers are willing to pay for a reduction in the risk of missing their

[4] With $U = \ldots + \beta_1 \ln(x_1) + \beta_2 \ln(x_2) + \ldots$, the ratio of the partial derivatives of U with respect to x_1 and x_2 is given by $\frac{\beta_1}{\beta_2}\frac{x_2}{x_1}$, as opposed to the simple $\frac{\beta_1}{\beta_2}$ ratio used in the case of a linear parameterisation.

	IVT *vs.* access cost (£/hour)	Freq. *vs.* access cost (£/flight)	Freq. *vs.* IVT (hours/flight)	Flight time *vs.* IVT
MNL				
Minimum	1.18	0.02	0.01	0.04
Mean	16.24	1.56	0.11	1.07
Maximum	143.38	231.05	4.06	7.43
Standard deviation	25.44	4.85	0.18	0.7
NL by airport				
Minimum	1.3	0.02	0.01	0.04
Mean	17.85	1.63	0.11	0.97
Maximum	157.65	242.38	3.87	6.72
Standard deviation	27.98	5.09	0.17	0.63
NL by airline				
Minimum	1.29	0.03	0.01	0.04
Mean	17.76	1.79	0.12	1.13
Maximum	156.8	265.11	4.26	7.85
Standard deviation	27.83	5.57	0.19	0.74
NL by access mode				
Minimum	0.99	0.02	0.01	0.04
Mean	13.52	1.11	0.1	1.05
Maximum	119.35	164.74	3.47	7.31
Standard deviation	21.18	3.46	0.15	0.69
CNL				
Minimum	1.16	0.01	0	0.04
Mean	15.96	0.9	0.07	0.95
Maximum	140.89	132.96	2.38	6.59
Standard deviation	25	2.79	0.1	0.62

Table 6: Model results for London data

flight, where this risk clearly increases with access time. While the lower values (compared to the SF-bay studies) could be explained on geographical grounds, it seems more likely that the use of a non-linear specification is the main reason for the lower (and it should be said more realistic) values; indeed, much higher values, together with a lower model fit, were obtained when using a linear specification. Finally, the still high values should also be put into context by noting that the average access journey in this population segment was measured as 57 minutes.

	Elementary alts. (324)	Airport (5)	Airline (37)	Access mode (6)
MNL	16.01%	61.47%	48.01%	39.27%
NL by airport	16.50%	62.88%	48.62%	38.58%
NL by airline	16.17%	61.34%	47.71%	39.54%
NL by access mode	16.03%	61.19%	47.84%	39.51%
CNL	16.48%	62.44%	47.79%	39.23%

Table 7: Prediction performance on London validation data

3.4 Model validation

The final part of the analysis is concerned with model validation. For this, the validation sample of 353 observations was used in application runs using the models presented in the preceding section. The results of this analysis are summarised in Table 7.

The results show relatively little variation between the three model structures, which was to be expected, when comparing the differences in model fit to the base LL. Furthermore, it is not clear a priori what measure of error should be associated with these measures, such that no inferences on differences between models should be drawn on the basis of these results. Without touching on the differences between models, it is of interest to compare the results to those obtained in the SF-bay study. The aggregate average probabilities of correct prediction are well below those obtained in the SF-bay study. This however needs to be put into context by noting that the choice set used in the SF-bay area was considerably smaller (3 airports, 8 airlines and 6 access modes). Furthermore, the exceedingly high market share for car made the analysis of access choice behaviour in the SF-bay area almost trivial. Finally, it seems that airport captivity plays a much bigger role in the SF-bay area than in London, where the levels of competition are much higher. This suggests that the models estimated on this data yield very satisfactory performance, even though they should still only be seen as a first step in the search of an optimal specification. Further gains can be expected by allowing for random taste heterogeneity inside a Mixed GEV framework; this is the topic of ongoing work.

4 Stated Choice study

The two RP studies discussed thus far in this chapter have again highlighted the issues faced with this type of data, notably in terms of the problems with retrieving a meaningful effect of changes in air fares. The study presented in this section aims to illustrate how these problems can be overcome with the use of SC data.

4.1 Description of data

The survey data used in this analysis were collected via the internet in May 2001 from a sample of around 600 individuals who had made a paid US domestic air trip within the twelve months prior to the interview taking place (Resource Systems Group Inc. 2003, Adler et al. 2005).

The survey uses a binary choice set, with ten choice situations per individual. In each choice situation, the respondent is faced with a choice between their recent observed trip, and an alternative journey option, compiled on the basis of the information collected in the RP part of the survey. These two alternatives are hereafter referred to as the *RP alternative* and the *SC alternative* respectively.

Aside from the actual airline and airport names, from which access times can be inferred, the attributes used to describe the two alternatives in the SC 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 flights. Access cost was not included in the surveys (in the absence of an actual specification of the mode-choice dimension), and no choice is given between different travel classes.

The final sample contains data collected from 589 respondents; with 10 choice situations per respondent, a sample size of $5,890$ observations is obtained, split into $1,190$ observations by business travellers, $1,840$ observations by holiday travellers and $2,860$ observations by VFR travellers. Further segmentations, for example by employment status, did not provide additional insights. Given the small sample sizes, especially for the business segment, and the high number of explanatory variables, the decision was taken to include all observations in the estimation process, rather than *waste* some of them on a validation sample.

4.2 Model specification

A large number of variables were included in the initial specification search. Aside from the continuous variables such as flight fare, flight time, access time, on-time performance (in %) and early and late schedule-delay (SDE and SDL), which need no further explanations, a number of discrete variables were also included.

As such, dummy variables were associated with different airports and airlines on the basis of the ranking provided by respondents in initial questioning, where appropriate normalisations were performed. Attempts were also made to estimate a constant associated with the airport closest to the passenger's ground-level origin. Three additional dummy variables were included in the base specification, to account for the effects of frequent flier (FF) membership, where these were associated with *standard* membership, *elite* membership, and *elite plus* membership. Similarly, dummy variables were included for flights with a single or double connection, while dummy variables were also associated with the different types of aircraft included in the survey. Finally, attempts were made to account for respondent inertia or habit formation with the help of a number of variables.

Aside from an alternative specific constant (ASC) for the RP alternative, airport and airline inertia constants were included in the utility of the SC alternative in the case where the RP airport or airline was reused in the SC alternative.

In addition to making use of non-linear transforms (log-transforms) where appropriate, this analysis also makes use of continuous interactions between socio-demographic attributes and taste coefficients, as described in ***REFERENCE TO HESS CHAPTER ON THEORY***. The socio-demographic attributes used in these interactions were income and travel distance (in the form of flight time for the RP alternative). Interactions with other factors, such as trip duration, or party size, were not found to be significant.

4.3 Model results

This section describes the findings of the estimation process. In the current work, only basic MNL structures were used. Nesting structures are not applicable given the nature of the choice set, while the use of mixture models, such as MMNL, was avoided with the aim of attempting to explain taste heterogeneity in a deterministic fashion. A separate analysis which made use of MMNL structures showed little additional gains in model fit, with the main advantage coming in a treatment of the repeated choice nature of the SC data. In this section, we only present a summary of the actual estimation results, with more details given by Hess et al. (2007).

4.3.1 Business travellers

The findings from the analysis using the $1,190$ observations for business travellers revealed effects for all the main continuous variables, including access time, air fare, flight time, and early and late arrival. Except for the early arrival penalty, the analysis showed that the use of a log-transform leads to significant gains in model performance, suggesting decreasing marginal returns for the associated attributes. The results further show positive effects of improvements in on-time performance. Initial results showed a reduced sensitivity to on-time performance on longer flights, but this resulted in problems with significance for the actual on-time performance coefficient. Efforts to use a power formulation for the on-time performance attribute (allowing for a much stronger *dislike* of very late flights) led to problems with parameter significance. In terms of interactions, the estimates additionally suggest a reduced sensitivity to early arrival on longer flights, as well as reduced fare-sensitivity with higher income.

The final part of this discussion looks at the findings for dummy variables. Here, a significant positive ASC was found to be associated with the current alternative, capturing inertia as well as a host of other effects. The estimation further shows a strong effect of frequent flier membership on the utility of an alternative, where, due to insignificant differences, a common factor was used for elite and elite plus membership, where the estimates show this to be over twice as large as for standard frequent flier membership. The fact that none of the airline dummy variables (linked to ranking) was found to be significant suggests that, for business travellers, airline allegiance is primarily limited to membership in frequent flier programmes. In terms of airport allegiance, a significant effect could only be associated with the second and top-ranked airports, where the former one was significant only at the 81% level. The estimated dummy variables for flights with one and two connections were indistinguishable, leading to the use of a common factor, where this can in part be seen as a result of the low incidence of flights with double connections in the data. The final set of dummy variables, associated with aircraft type, show that single-aisle jets are clearly preferred over turboprop planes and regional jets.

4.3.2 MNL model for holiday travellers

The findings from the analysis using the $1,840$ observations for holiday travellers revealed significant effects of access time, air fare and flight time, where a log-transform was again found to be appropriate for all three attributes. The first difference with the business models arises in the treatment for schedule delay, where the use of linear effects was found to be preferable, and where, given the small differences between the effects for early and late arrival, a common coefficient was used. The results again show positive effects of improvements in on-time performance, where the associated interaction term suggests that holiday travellers' sensitivity to on-time performance increases with flight-distance. This can be explained for example by the notion that holiday flights are often pushed to the edges of the off-peak periods, where sensitivity to on-time performance may indeed be greater, and especially so for very long flights.

Other interactions again show a reduced fare-sensitivity with higher income, although the confidence level for the associated term is very low. The interaction terms also show that, for holiday travellers, fare sensitivity increases with flight-distance. It is important to put this into context by remembering that a log-transform is also used on the fare attribute. As such, the results simply suggest that, at a given fare level, increases are valued more negatively in the case of longer flights. A possible explanation for this could be the higher secondary costs associated with longer flights in the case of holiday travellers.

As expected, frequent flier benefits play a much smaller role in this segment of the population, where it was only possible to estimate a common dummy variable for all levels of membership, which in addition only attains a very low level of statistical significance. On the other hand, a significant positive effect is associated with the top-ranked airline.

In this sample, the effect associated with flights with two connections is also significantly larger than for flights with a single connection, and the scale of the difference (factor of 3) supports the decision not to use a linear effect, but to use two separate dummy variables. Finally, for the aircraft type dummies, the results suggest that holiday travellers do not distinguish between single-aisle jets, regional jets, and turboprop planes, with the only aircraft dummy with a modestly significant value being that for wide-body aircraft, which are seemingly given a slight preference over single-aisle jets.

4.3.3 MNL model for VFR travellers

The findings from the analysis using the 2,860 observations collected from VFR travellers reveal an important difference when compared to those for business and holiday travellers. Indeed, while access time and flight fare again enter the utility function under a log-transform, the specification search indicated that it is preferable to treat flight time in a linear fashion. Early and late arrival penalties are treated separately in this model, and both enter the utility in a linear form, where the penalty associated with late arrival is lower, and attains a very low level of statistical significance.

The non-linear interactions retrieved from this data show heightened fare sensitivity on longer flights, along with reduced fare sensitivity with higher income, and lower sensitivity to access time on longer flights, which would support a decision to shift long haul flights to outlying airports, where the issue of point-to-point passengers on the required feeder-flights would however need to be addressed separately.

In this segment, it was not possible to estimate a significant effect associated with frequent flier programmes, while the dummy variables associated with the two most preferred airlines are positive and significant at high levels of confidence. Airport allegiance also seems to play a role, where there is however essentially no difference between the estimates of the dummies associated with the two top-ranked airports. Finally, unlike in the other two population segments, it was also possible to identify a significant positive effect associated with the airport closest to the passenger's ground-level origin.

A common effect was again used for flights with single and double connections, while, in terms of aircraft type, the difference between single-aisle jets and regional jets is significant only at the 87% level, where the results further indicate a significant dislike for turboprop flights, and a significant preference for wide-body jets over single-aisle jets.

4.4 Comparison of results across population segments

The description of the MNL model fitting exercises has highlighted a number of differences between the specifications used in the three population segments, notably in terms of non-linearities and interactions with socio-demographic variables. These differences in model specification need to be borne in mind when comparing the substantive results across the three population segments. The calculation of the trade-offs, and hence the comparison of results across groups, is further complicated by the high number of non-linear terms in the utility functions, where the simple ratio between coefficients is no longer applicable. The

	Business	Holiday	VFR
Reduction in access time (1 hour)	75.40	35.80	35.48
Reduction in SDE (1 hour)	13.27$^{(*)}$	2.61$^{(*)}$	3.68
Reduction in SDL (1 hour)	11.08		2.25$^{(*)}$
On-time (+10%)	10.39	7.02	5.57
FF elite or elite-plus *vs* none	125.24	11.44$^{(*)}$	-
FF standard *vs* none	49.12		-
Top airline *vs* worst	-	25.07	21.06
2^{nd} airline *vs* worst	-	18.16$^{(*)}$	15.27
3^{rd} airline *vs* worst	-	20.09$^{(*)}$	4.77$^{(*)}$
Top airport *vs* worst	83.22	53.97	55.73
2^{nd} airport *vs* worst	30.56$^{(*)}$	41.42	54.63
3^{rd} airport *vs* worst	-	18.54$^{(*)}$	25.89
Airport closest to home	-	-	28.02
No connection *vs* one connection	44.15	19.60	18.98
No connection *vs* two connections		62.21	
Jet *vs* wide-body	29.86$^{(*)}$	-	-
Jet *vs* regional jet	79.51	-	10.59$^{(*)}$
Jet *vs* turboprop	96.94	1.79$^{(*)}$	17.77
Wide-body *vs* jet	-	13.45$^{(*)}$	27.84
Regional jet *vs* jet	-	1.31$^{(*)}$	-

$^{(*)}$ Coefficient used in numerator of trade-off not significant at 95% level

Table 8: MNL trade-offs, part 1: willingness to pay ($)

situation becomes more complicated again in the case of coefficients interacting continuously with income or flight-distance, where any trade-off involving such coefficients will vary across individuals as a function of the associated attribute.

In the present analysis, the comparison was limited to two mains sets of trade-offs, looking at the willingness to accept increases in fare and access time respectively, in return for *improvements* in other determinants of choice. In each case, the trade-offs are presented for the average flight-distance and household income in that population-segment, meaning that the interaction terms become equal to 1.

The results are summarised in Table 8 for the willingness-to-pay indicators, and Table 9 for the willingness to accept increases in access time. In each case, several coefficients used in the trade-offs were not significant at the 95% level, as pointed out in Sections 4.3.1, 4.3.2 and 4.3.3, and this is indicated appropriately in the presentation of the trade-offs.

Consistent with a priori expectations, the results show a much greater willingness to accept higher fares in return for shorter access times for business travellers than for holiday or VFR travellers, by a factor of just over 2. The models also indicate a higher willingness by business travellers to pay for reductions in schedule

	Business	Holiday	VFR
Reductions in fare ($1)	2.14	4.61	4.57
Reduction in SDE (1 hour)	17.38$^{(*)}$	8.25$^{(*)}$	12.24
Reduction in SDL (1 hour)	17.00		7.49$^{(*)}$
On-time (+10%)	13.60	22.16	18.53
FF elite or elite-plus *vs* none	163.97	36.10$^{(*)}$	-
FF standard *vs* none	64.31		-
Top airline *vs* worst	-	79.11	70.08
2^{nd} airline *vs* worst	-	57.31$^{(*)}$	50.81
3^{rd} airline *vs* worst	-	63.40$^{(*)}$	15.88$^{(*)}$
Top airport *vs* worst	108.96	170.29	185.43
2^{nd} airport *vs* worst	40.01$^{(*)}$	130.69	181.78
3^{rd} airport *vs* worst	-	58.49$^{(*)}$	86.14
No connection *vs* one connection	57.81	61.86	63.15
No connection *vs* two connections		196.29	

$^{(*)}$ Coefficient used in numerator of trade-off not significant at 95% level

Table 9: MNL trade-offs, part 2: willingness to accept increases in access time (min)

delay and for improved on-time performance. Interestingly, the models suggest that, except for holiday travellers, respondents are more sensitive to early than to late arrival, a finding that should however be put into context given the small differences, and high associated standard-errors.

The models suggest that business travellers are willing to pay $125 to fly on an airline where they hold an elite frequent-flier account. Even though this figure decreases to $49 in the case of standard membership, the figures are still much higher than for holiday travellers, while no such effects could be identified for VFR travellers. In these latter two groups, the results however show a certain willingness to pay a premium for flying on either of the top-ranked airlines.

In terms of paying a premium for direct flights, the results again suggest a higher willingness for business travellers, although the different treatment in the case of holiday travellers results in a higher value for the trade-off in the case of flights with 2 connections in this group. A difference arises between the three population groups in the trade-offs looking at the willingness to pay for flying on a specific type of aircraft, where the differences in the *most-valued* type of aircraft led to a different base-type.

The findings for the trade-offs looking at the willingness to accept increases in access time do, overall, show a lower willingness for business travellers than for holiday and VFR travellers, which is to be expected. The main exception again comes in the case of frequent-flier benefits, where the results suggest that business travellers are willing to fly out of more distant airports in return for flying on an airline whose frequent-flier programme they are a member of.

Some of the trade-offs presented in this section are very high; this could potentially be a reflection of the well-established notion that in SC studies, there is a tendency for respondents to exaggerate their responsiveness to changes in attributes (e.g. Louviere et al. 2000, Ortúzar 2000).

5 Summary and Conclusions

The results of the three studies discussed in this chapter are not directly comparable, given the geographical differences, the differences in the age of the data, and the differences in survey design as well as data type (RP *vs* SC). Nevertheless, some conclusions can be reached.

From a model structure point of view, all three case studies have shown that the use of more advanced model structures can lead to improvements in model fit. However, although the improvements are statistically significant, they are too small to lead to any major differences in model performance. Nevertheless, the advanced model structures provide further insights into choice behaviour, and there are also differences in the substantive results between the various models.

The main observation that can be made in the comparison of the results across the three studies is the greater ability of the SC models to retrieve significant effects for a range of variables that are generally not well estimated in RP studies, such as air fares, schedule delay and airline and airport allegiance. This is an illustration of the complications that arise with the use of RP survey data in the analysis of air travel choice behaviour, where there are issues of data quality in relation to air fares and availabilities, while information on a number of other attributes, notably the membership in frequent flier programmes, is not generally available in such datasets. On the other hand, it should be remembered that there is a risk of bias when relying solely on the use of SC data, making the joint estimation on RP and SC data an important avenue for further research, as discussed by Algers & Beser (2001).

The one common observation that can be made from the three case studies is that the results do suggest that access time plays a major role in the choice process, with passengers having a strong preference for their local airport. As such, the attractiveness of outlying airports depends heavily on good access connections, unless there are other incentives, such as low air fares. This is reflected in the fact that only low cost carriers find it relatively easy to attract passengers to outlying airports that are not served by convenient and fast ground-level services. It is conceivable that the sensitivity to access time decreases with flight time[5], such that moving long haul services to outlying airports would seem wise; this however causes problems as the associated (and necessary) short haul feeder flights will also carry point-to-point passengers, who will again have a preference for more centrally-located airports.

[5]As suggested by the results for VFR travellers in Section 4.3.3

Acknowledgements

The material described in this chapter is based on work conducted by the main author during his time in the Centre for Transport Studies at Imperial College London. This chapter was partly written during a guest visit in the Institute of Transport and Logistics Studies at the University of Sydney. The support of both institutions is gratefully acknowledged. The authors would also like to thank Thomas Adler for his input to the SC study, and would similarly like to thank the Metropolitan Transport Commission and the Civil Aviation Authority for data support.

References

Adler, T., Falzarano, C. S. & Spitz, G. (2005), *Modeling Service Trade-offs in Air Itinerary Choices*, paper presented at the 84$^{\text{th}}$ Annual Meeting of the Transportation Research Board, Washington, DC.

Algers, S. & Beser, M. (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**(1-2), 28–45.

Basar, G. & Bhat, C. R. (2004), 'A Parameterized Consideration Set model for airport choice: an application to the San Francisco Bay area', *Transportation Research Part B: Methodological* **38**(10), 889–904.

CAA (1996), *UK airports Passenger survey, CAP 677 - Passengers at Birmingham, Gatwick, Heathrow, London City, Luton, Manchester and Stansted in 1996*, The Civil Aviation Authority, London.

Franz, J. D. (1996), *Metropolitan Transportation Commission Airline Passenger Survey: Final Report*, prepared by J.D. Franz Research for the Metropolitan Transport Commission, Sacramento, CA.

Hess, S. (2004), *A model for the joint analysis of airport, airline, and access-mode choice for passengers departing from the San Francisco Bay area*, paper presented at the European Transport Conference, Strasbourg.

Hess, S., Adler, T. & Polak, J. W. (2007), 'Modelling airport and airline choice behavior with stated-preference survey data', *Transportation Research Part E: Logistics and Transportation Review* **43**, 221–233.

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

Hess, S. & Polak, J. W. (2005*b*), 'Mixed logit modelling of airport choice in multi-airport regions', *Journal of Air Transport Management* **11**(2), 59–68.

Hess, S. & Polak, J. W. (2006*a*), 'Airport, airline and access mode choice in the san francisco bay area', *Papers in Regional Science* **85**(4), 543–567.

Hess, S. & Polak, J. W. (2006*b*), 'Exploring the potential for cross-nesting structures in airport-choice analysis: a case-study of the Greater London area', *Transportation Research Part E: Logistics and Transportation Review* **42**(2), 63–81.

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

OAG (1996), *World Airways Guide Fares Supplement 1996*, Reed Elsevier, Dunstable.

ONS (1996), *International Passenger Survey*, The Office for National Statistics, London.

Ortúzar, J. de D. (2000), *Stated Preference Modelling Techniques: PTRC Perspectives 4*, PTRC Education and Research Services Ltd, London.

Pels, E., Nijkamp, P. & Rietveld, P. (2003), 'Access to and competition between airports: a case study for the San Francisco Bay area', *Transportation Research Part A: Policy and Practice* **37**(1), 71–83.

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

Scott Wilson Kirkpatrick (1999), *National airport accessibility model: model status report*, prepared for the Department for Transport by Scott Wilson Kirkpatrick & Co Ltd, Basingstoke.