EXPLORING THE POTENTIAL FOR CROSS-NESTING STRUCTURES IN AIRPORT-CHOICE ANALYSIS: A CASE-STUDY OF THE GREATER LONDON AREA[1]

Stephane Hess

Centre for Transport Studies
Imperial College London
Exhibition Road
London SW7 2AZ
Tel: +44(0)207-594-6105
stephane.hess@imperial.ac.uk

John W. Polak

Centre for Transport Studies
Imperial College London
Exhibition Road
London SW7 2AZ
Tel: +44(0)207-594-6089
j.polak@imperial.ac.uk

ABSTRACT

The analysis of air-passengers' choices of departure airport in multi-airport regions is a crucial component of transportation planning in many large metropolitan areas, and has been the topic of an increasing number of studies over recent years. In this paper, we advance the state of the art of modelling in this area of research by making use of a Cross-Nested Logit (CNL) structure that allows for the joint representation of inter-alternative correlation along the three choice dimensions of airport, airline and access-mode. The analysis uses data collected in the Greater London area, which arguably has the highest levels of inter-airport competition of any multi-airport region; the authors of this paper are not aware of any previous effort to jointly analyse the choice of airport, airline and access-mode in this area. The results of the analysis reveal

---

significant influences on passenger behaviour by access-time, access-cost, flight-frequency and flight-time. A structural comparison of the different models shows that the cross-nested structure offers significant improvements over simple Nested Logit (NL) models, which in turn outperform the Multinomial Logit (MNL) model used as the base model.

1. INTRODUCTION

For a number of reasons, not least of which the precarious financial situation of many of the world's leading airlines, and the long-term horizon of the effects of policy-changes, the analysis of air-travellers' choice-behaviour is a crucial part of travel-demand forecasting.

One area that is of special interest is the analysis of the choices made by passengers departing from major multi-airport regions. The modelling of such choices is not only appealing from a research perspective, due to the complexity of the choice process, but is also of great practical policy importance in many large metropolitan areas. Indeed, the unprecedented increase in air-transport (c.f. IATA, 2002) has led to important problems of congestion in a number of multi-airport regions, leading to urgent needs for capacity expansion, especially given the forecasts for continued strong growth (e.g. Boeing, 2003).

Given the long-term nature of any such expansion work, and the associated financial and environmental constraints, the generation of reliable forecasts of passenger levels (at the airports, as well as in the ground-level access-network) is a prerequisite to any planning process. The generation of such forecasts in turn leads to a requirement for

an understanding of passenger behaviour, as yielded by studies of airport choice behaviour.

Consequently, it is of little surprise that this area of research has seen increased levels of activity in recent years, as detailed in section 2 of this paper. However, the majority of existing studies have been of a fairly basic nature, in their use of restricted model structures, as well as over-aggregated data. Furthermore, studies of airport choice behaviour have generally failed to acknowledge the three-dimensional nature of the choice-set, in that passengers not only make a choice of departure airport, but additionally choose an airline and an access-mode. The joint analysis of these three choice-dimensions, and the interactions between them, can however lead to important gains in model performance and accuracy, as illustrated by Hess (2004) and Hess & Polak (2004, 2005b)

This paper advances the current state of the art in airport choice modelling from a modelling as well as topical perspective. The methodological advance comes in the form of the use of cross-nesting model structures that allow for the joint representation of correlation along the three choice dimensions, without requirements to use a multi-level nesting structure. The other merit of the paper lies in the fact that it presents an analysis of airport choice in the London area, which is arguably the most competitive multi-airport region in the world. The authors of this paper are not aware of any existing independent public-domain study that looks at the choice between the five main airports in this region[2], although some previous studies have included one or more of the London airports (generally Heathrow) in the study of airport choice in wider geographical areas (e.g. Ashford & Bencheman, 1987, and

---

2 A study into the optimal utility specification in Multinomial Logit models for airport choice (ignoring airline and access-mode choice) was carried out by the DfT in 2000, using explanators similar to those used in the present analysis (c.f. DfT, 2000b). This was preceded by the CAA/NATS Second Passenger Allocation Model (SPAM), and other predecessors, all using Multinomial Logit structures.

Brooke et al., 1994). It is worth noting that the present analysis, like most previous studies of airport choice, does not explicitly consider the choice of main mode, a decision that is based on the reasoning that travellers included in the survey have already made the decision to go by air. The analysis of the joint choice of main mode and route is an important avenue for future research. Furthermore, primarily for data reasons, the study looks only at departing passengers, ignoring arriving passengers, as well as connecting passengers. Finally, only passengers on direct flights are included in the analysis.

The remainder of this paper is organised as follows. In the next section, we present a brief review of recent work in the field of airport choice behaviour. The third section describes the airport system in place in the Greater London area, and discusses the data used in the analysis. The fourth section looks at methodology and model specification, while the fifth section presents the results of the analysis. Finally, section six discusses model validation, and section seven summarises the findings of the research.

2. LITERATURE REVIEW

Although research into airport-choice behaviour stretches back to the work of Skinner (1976), the majority of studies in this area date from the mid 1990s to the early 2000s. In turn, the core of this work makes use of data collected in the San Francisco Bay (SF-Bay) area. During the past few years, a number of authors have used these data. Pels et al. (2001) conduct a Nested Logit (NL) analysis of the combined choice of airport and airline and find that that airline choice is nested within the choice of airport, while Pels et al. (2003) again make use of the NL model structure, this time in the joint analysis of airport and access-mode choice, revealing high sensitivity to access time, especially for business travellers. Basar & Bhat (2004) propose the use of

a two-level modelling structure in which the airport choice process is preceded by a choice-set generation stage, thus acknowledging the fact that not all travellers consider all available airports. The results show flight frequency to be the most important factor in choice set composition, with access-time playing this role in the actual choice of airport. Finally, Hess & Polak (2005a) use the SF-bay data in a study aimed at showing the role of random taste heterogeneity in airport-choice behaviour, in turn introducing the Mixed Multinomial Logit model to the area of air-travel choice behaviour (where it had previously only been used in wider mode-choice contexts). The results show that, while a major part of the variation in tastes can be accounted for through a segmentation of the population, a remaining part of variation is purely random. Hess & Polak (2004) further advance the state of the art in airport choice modelling by explicitly modelling the combined choice of airport, airline and access-mode. Finally, Hess & Polak (2005b) adapt this framework to additionally allow for random taste heterogeneity within the three-dimensional choice-process.

To date, three main academic studies of airport access in the United Kingdom (UK) have been undertaken. Ashford & Bencheman (1987) use an MNL model for airport choice at five airports in England (including two London airports), and find that access time and flight frequency are significant factors overall, while flight fares only have an impact for domestic passengers and for international leisure travellers. Thompson & Caves (1993) use an MNL model to forecast the market share for a new airport in North England; access time, flight frequency and aircraft-size are found to be significant. Finally, in an MNL analysis of the distribution of passengers in the Midlands, Brooke et al. (1994) find flight frequency to be most the important factor, where a single London airport, Heathrow, was also included in the choice-set.

Other work in the area of airport choice is of more limited interest to the present study, from a methodological as well as geographical point of view. For a review of such studies, the reader is referred for example to Pels et al. (2001, 2003), Basar & Bhat (2004), and Hess & Polak (2004).

3. DESCRIPTION OF DATA AND CHOICE SCENARIO

The Greater London area has by far the highest levels of air traffic in Europe, with, in 2002, some 117.13 million passengers (departing, arriving and connecting) for the five main airports. The area is dominated by Heathrow (LHR), the world's busiest international airport (number of passengers on international routes). Additionally, a large number of routes are offered from Gatwick (LGW), the world's busiest single-runway airport, while Stansted (STN), Europe's fastest growing major airport, 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. Table 1 shows the annual number of passengers handled at the five London airports between 1997 and 2002 (CAA, 2002). The figures show that LHR is easily the busiest of the five airports, capturing just over half of the total number of passengers. It is worth noting that the total number of passengers handled in the UK in 2002 was 189.1 million, such that the London airports account for over 60% of the total, highlighting the dominating role of London, and by extension the South East.

Forecasts show that air-travel in the United Kingdom can be expected to grow at a very high rate. As an example, forecasts produced by the Department for Transport in 1997 predicted almost a doubling in air passenger numbers by the year 2010 (c.f. DfT, 1997), with a revised version of this forecast, produced in 2000, predicting an increase in passenger numbers to 2.5 times the level observed in 1998, by the year 2020 (c.f. DfT, 2000a). Finally, it has been suggested that, with unconstrained growth, traffic

could rise to around 500 mppa by the year 2030, with around 300mppa for the South East airports, with most of this distributed across the London airports (c.f. DfT, 2002). Given the limits in capacity, unconstrained growth is clearly a purely hypothetical situation and barely reflects the potential in terms of demand. As demand at the London airports, and especially at Heathrow, already exceeds capacity, concerns have been raised that London could lose its status as the main European hub, given the extra capacity available in competing regions, such as Paris, Frankfurt and Amsterdam.

One major airport expansion scheme is already in progress, with the construction of *Terminal 5* at Heathrow. Assuming an annual increase in passenger numbers by around 4% (which is hypothetical given constraints on runway capacity), the new capacity limit of 89mppa could however be reached within 2-3 years of the opening of the new terminal in 2008. There is thus a need for further increases in capacity, and a major consultation has taken place to consider different schemes for expanding airport capacity in London and the South East (c.f. DfT 2002).

The most popular proposal with the main airlines is the construction of a third (short) runway at Heathrow, which would increase capacity to 116mppa; this project is however facing fierce opposition by local communities. Another possibility is the introduction of mixed mode operations at Heathrow, allowing both runways to be used simultaneously for take-offs and landings, which could allow an increase by 10% in the number of take-offs and landings. Public opposition to this scheme is however also very high, primarily on the grounds of extra noise-pollution.

A major problem in the initial search for alternative ways of increasing capacity was the agreement signed in 1979 between the British Airports Authority and West Sussex County Council that a second runway at Gatwick would not be built before 2019,

whereas estimates by DfT (2002) show that capacity could be increased dramatically, were the construction of new runways allowed. Little gains in capacity are possible at London City, while the location and space availability at Luton also limit the possibilities for turning this airport into a major new hub for London. The same reasoning applies in the case of the expansion of secondary airports. Plans to develop a new four-runway airport at Cliffe in North Kent have also been rejected, mainly due to environmental concerns. This has turned increased attention to Stansted. Its current capacity of 15mppa could be increased to 25mppa without new runway development, while capacity could be increased to 82mppa, 102mppa and 129mppa respectively with one, two or three additional runways. A recent UK government White Paper (DfT, 2003) has recommended the construction of a single new runway at Stansted by 2012. There are also plans to extend Heathrow's capacity between 2015 and 2020 with the construction of a new runway and possibly a sixth terminal. At this stage, these plans are however only a recommendation, and deliberations are set to continue, especially since the major airlines, which do not operate at Stansted, are opposed to it being expanded ahead of Heathrow, and are threatening to boycott Stansted. It is thus still of interest to gauge the attractiveness of the different airports, and to analyse how the attractiveness of airports with additional capacity could be improved.

This context makes the London area a prime candidate for a study of airport choice. Furthermore, unlike in many other areas that have been the topic of studies of airport choice, there are very high levels of competition between the different airports, and less captivity by specific geographical areas to a given airport, given their arrangement at roughly equal distances from the centre of London (aside from LCY). Finally, unlike in the case of studies in the US, where the market-share of car can exceed 75% (as in the San Francisco Bay region), the modal split for the access-

journey is very wide, increasing interest in the analysis of choice along this dimension.

For the present analysis, data from the 1996 passenger survey were obtained from the Civil Aviation Authority (CAA, 1996). Although slightly dated, the dataset has the advantage that the effects of the September 11[th] attacks need not be taken into account. On the other hand, the age of the data prevents a detailed analysis of the impact of low-cost carriers on air-travel choice behaviour, given that their operations in 1996 were far more limited than is the case nowadays. It should also be noted that the analysis of the access-mode choice dimension is simplified by the fact that the premium Heathrow Express service only started its operations in 1998. The use of a more recent version of the dataset is an important avenue for further research.

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. This dataset was complemented by air-travel level-of-service data, obtained from BACK aviation[3], and ground-network data obtained from the Department for Transport. After data-cleaning (missing data, compatibility between datasets), a usable sample of 33,612 passengers was obtained. This compares favourably to the sample of 5,091 available to Hess & Polak (2004) in the SF-bay area. In the present paper, only one segment of the population, namely resident business travellers[4], was used, leading to a sample size of 7,059 observations, which was split into an estimation sample of 6,706 observations, and a validation sample of 353 observations. This again compares very favourably to the

---

[3] www.backaviation.com

[4] Respondents were considered as residents if they reside in the Greater London area and immediately adjacent counties for domestic and short-haul European flights, while respondents on medium-distance European flights and intercontinental flights were considered as residents if they reside in the UK.

corresponding sample sizes of 1,098 and 114 observations in the San Francisco Bay-area study of Hess & Polak (2005b).

Of the 31 destinations used in the analysis, 5 are in Great Britain (Aberdeen, Edinburgh, Leeds, Manchester and Newcastle), 1 on the Channel Islands (Guernsey), 3 in Ireland (Cork, Dublin and Shannon), 3 in the Benelux (Amsterdam, Brussels and Rotterdam), 3 in Scandinavia and the Nordic countries (Copenhagen, Gothenburg and Helsinki), 3 in Germany (Düsseldorf, Hamburg and Munich), 3 in Austria and Switzerland (Geneva, Vienna, Zurich), 1 in France (Nice), 3 in Spain (Barcelona, Madrid, Malaga), 2 in the South East of Europe (Athens and Larnaca), 1 in the Middle East (Tel Aviv), and 3 in the United States (Boston, Detroit and Miami). The most popular destination in the sample is Amsterdam, ahead of Edinburgh, Dublin and Brussels. At this point, it is worth noting that all destinations are served by a single main airport, avoiding the problem with multi-airport destination areas described by Hess & Polak (2004,2005a).

The air-travel level-of-service dataset contains daily airline-specific information for all routes used in the analysis, making the analysis more reliable than research relying on the use of weekly or even monthly data. The dataset contains information on flight frequencies, departure times, flight-times (block times, thus taking into account airport congestion), aircraft types used and available seat capacity. The main bit of information missing from this dataset is that of the fares for the different routes and airlines. Such information was compiled from two sources; the International Passenger Survey (ONS, 1996)[5] and the fare supplement of the Official Airways Guide for 1996 (OAG, 1996). Information on travel-class as well as ticket type (single

---

[5] It should be noted that there are potential problems of endogeneity in using data on fares actually paid, given the likelihood that passengers choose the cheapest fares. However, in the face of incomplete listed fare data, this is not avoidable, and the same problems occur with information from the 10% sample, as generally used in studies of airport choice in US multi-airport regions.

or return) was taken into account in assembling the data. As was the case with the fare data used by Hess & Polak (2004 and 2005a,b) and other previous studies of airport choice behaviour, the resulting dataset is of highly aggregate nature[6], giving the average fare paid across travellers on a given route by a given airline, hence ignoring the potential differences resulting from seat and ticket class availability. In the absence of more detailed data, the estimated fare coefficient should thus not be seen as a reliable estimate of the marginal utility of air-fare; the fact that it is often impossible to estimate a significant fare coefficient in airport choice modelling partly reflects the generally low quality of fare data.

Another major problem in airport choice modelling is that there is generally no information on the availability of flights (or indeed ticket classes) on the different routes at a given point in time. This leads to the major, yet necessary assumption, that flights on all possible routes were available at the time of booking. Finally, it is well-known that airline-allegiance plays a major role in air-travel choice behaviour; the absence of any information on frequent flier programmes further limits the power of the models, though some degree of allegiance can be modelled on the basis of traveller nationality. While it is important to stress the potential impact of these limitations on modelling performance, the good performance of existing models (e.g. Basar & Bhat, 2004, Hess & Polak, 2004, 2005b) should nevertheless be seen as an encouraging sign.

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 & Co Ltd, 1999). Corresponding cost information for 1996 was produced

---

[6] The ONS dataset was made available at the individual passenger level for a high number of respondents, such that statistics of the distribution of fares could be calculated. However, given the lack of such detailed information for national flights and some international destinations, where only

with the help of the retail price index, while assuming that relative travel-times have on average stayed constant. This dataset contains level-of-service information for travel between 455 different travel area zones and the five airports. 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). Respondents observed to have used hotel shuttles were excluded from the analysis, given the lack of information on shuttle availability, and the potential correlation between the choice of airport and hotel, and the use of courtesy shuttles. The use of a high level of disaggregation in the non public transport modes alongside aggregated public transport information might be criticised given the continuing focus on competition between premium dedicated airport rail services and other forms of public transport (e.g. Gatwick Express versus local train services); the division used in the present analysis was reflective of the greatest common denominator between the survey and level-of-service datasets, and it is hoped that future work can rely on a higher level of disaggregation.

The NAAM dataset did not contain information on taxi and minicab services; this was produced independently, on the basis of data for the year 2004, with appropriate transformations to obtain usable data for 1996. In terms of availability, taxi and minicab are assumed to be available for all possible ground-level and airport combinations, while the availability of public transport (PT) and long-distance coach (LDC) was determined on the basis of the NAAM data. Finally, rental car was assumed to be available to all travellers above 18 (in the absence of license-holding information), while car was assumed to be available to all travellers (given that only residents were used in this analysis). No combinations of modes were considered in

the present analysis, and the main mode indicated in the survey was used as the chosen mode. For each mode, information was included on travel time, wait time, and the number of interchanges (where appropriate). For the cost information, a fixed one-day charge of £35 was used for hire-cars (in the absence of cost-bearing party information) in addition to marginal running costs (fuel only), while fare information was used for PT, LDC, taxi and minicab. For car, two specifications were retained, one using only the marginal running costs in terms of fuel, with a second also including depreciation. Finally, the dataset was completed by adding parking cost information for the different airports, for short as well as long-term parking.

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 (instead of 185), 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.

4. METHODOLOGY AND MODEL SPECIFICATION

Over the past thirty years, discrete choice models belonging to the family of random utility models (RUM) have become the preferred tool for choice analysis in the area of transportation. In a RUM, the gain that decision-maker n obtains from choosing alternative i from a choice-set of I alternatives is given by the utility U_{ni}. Under the assumption of rational choice behaviour, the alternative with the highest utility is chosen. However, due to modelling and data uncertainty, only part of the utility of an alternative is observed, and U_{ni} is accordingly divided into an observed utility V_{ni}, and

---

consistency.

a remaining, unobserved part of utility (error-term), ε_{ni}, The observed utility V_{ni} is a function of the tastes of the decision-maker, β, and the attributes of the alternatives, grouped in a vector x_{ni}, which can also contain socio-demographic characteristics of the decision-maker. Typically, a linear-in-parameters specification is used, such that $V_{ni} = \beta'x_{ni}$. With the resulting random nature of U_{ni}, the choice becomes probabilistic, with the choice probability of alternative i being given by:

$$P_{ni} = P\big(V_{ni} + \varepsilon_{ni} > V_{nj} + \varepsilon_{nj}, \forall j \neq i\big),$$

where the form of the choice probabilities depends on the distributional assumptions made with regards to the error-terms.

In models belonging to the family of Generalised Extreme Value (GEV) models (McFadden, 1978), the marginal distribution of the individual error-terms (ε_{ni}) is univariate extreme value, where different assumptions about the cumulative distribution of the vector of error-terms $\varepsilon_n = (\varepsilon_{n1}, \ldots, \varepsilon_{nI})$ lead to different model forms. In this paper, we make use of three types of GEV models; the Multinomial Logit (MNL) model (McFadden, 1974), and two of its generalisations, the Nested Logit (NL) model (Williams, 1977, Daly and Zachary, 1978 and McFadden, 1978), and the Cross-Nested Logit (CNL) model (Vovsha, 1997). We will only present a brief description of the three model structures, for detailed explanations and probability functions, the reader is (for example) referred to Train (2003).

In the MNL model, the individual error-terms are assumed to be distributed identically and independently following a type I extreme value (Gumbel) distribution. Although this leads to a convenient form for the choice-probabilities, the MNL model has the disadvantage that all alternatives depend on each other in the same way, leading to unrealistic substitution patterns in the case where correlation exists between the unobserved parts of utility, as in the presence of unobserved attributes shared by

some of the alternatives. Such factors are however likely to play a role in airport choice behaviour (e.g. business airport vs. low-cost base), which would make the MNL model an inappropriate choice of structure.

In more advanced models, correlation across alternatives in the unobserved utility-components is taken into account by specifying a joint distribution of the error-terms with a non-diagonal covariance matrix. The best-known such model is the NL model, which divides the choice set into hierarchical and mutually exclusive nests of alternatives, with increased correlation, and thus higher cross-elasticities, between alternatives sharing a nest. In this model, a structural (logsum) parameter is associated with each nest, say λ_m for nest m. The structural parameters measure the degree of independence between alternatives in the respective nest, with higher λ_m meaning more independence and hence lower correlation (given by $1 - \lambda_m^2$) between the unobserved components of utility of the alternatives contained in nest m. For consistency with utility maximisation, λ_m is generally constrained to lie between 0 and 1, where a value of 1 for all structural parameters leads to the MNL model.

The NL model has been used repeatedly in airport choice modelling, with some examples being the work of Pels et al. (2001,2003) and Hess & Polak (2004). In the analysis of the joint choice of airport, airline and access-mode, three possible one-level nesting structures arise. The first example uses nesting by airport, such that in the case of *K* airports (where *K* is in this case equal to 5), each elementary alternative (triplet of airport, airline and access-mode) is assigned to exactly one airport nest, hence acknowledging correlation in the unobserved utility terms for alternatives sharing the same departure airport. This structure is formally described in Figure 1, where λ_k identifies the nesting parameter associated with the nest grouping together alternatives sharing airport *k*. Only a subset of the composite nests and of the

elementary alternatives is shown in the graph. The corresponding structures for the models using nesting by airline and nesting by access-mode are not reproduced here, being simple analogues of the structure shown in Figure 1.

The NL model can be extended to allow for multi-level nesting; this can be exploited to allow for correlation along multiple choice-dimensions, as used notably by Pels et al. (2001,2003). The main problem with the use of multi-level nesting structures for the analysis of the three-dimensional choice process of airport, airline and access-mode is that this structure can only accommodate correlation along at most two of the three dimensions. As an illustration, if nesting by airport and airline is used, each airport-airline nest would in our case contain 6 elementary alternatives, one for each access-mode. As such, the correlation between alternatives sharing the same access-mode cannot be accommodated. A similar reasoning applies for the joint nesting by airport and airline, or by airline and access-mode. In fact, it can be seen that the full extent of correlation can only be taken into account along one dimension, with a limited amount along the second dimension. Indeed, by nesting the alternatives first by airport, and then by airline, the nest for airline 1 inside the nest for airport 1 will only group together the options on airline 1 for that specific airport. The same reasoning applies for other nests. As such, the model is not able to capture correlation between alternatives using airline 1 at airport 1 and alternatives using airline 1 at airport 2, which is clearly a restriction. This problem also applies in the other multi-level nesting approaches.

These deficiencies of multi-level nesting structures are the motivation for the efforts made in this paper to use cross-nesting structures. In such a CNL model, the allocation of alternatives to nests is fuzzy, with alternative j belonging to nest k with proportion α_{jk}, where the allocation parameters for an alternative sum to 1 over nests.

To the authors' knowledge, there has so far been only one attempt to use a CNL model in airport choice modelling, by Hess (2004). In that analysis, the estimation problems encountered when using the CNL model however meant that the structure failed to outperform some of the one-level NL structures. As an illustration, Figure 2 shows the structure used for the cross-nesting model estimated in this paper. In this model, each alternative belongs to exactly one airport nest, one airline nest, and one access-mode nest, and the model is able to jointly represent the correlation along the three dimensions. Again, only a subset of the composite nests and of the elementary alternatives is shown, and the allocation parameters are not represented in this figure.

A further extension of discrete models comes in the form of structures allowing for random taste heterogeneity, such as the Mixed Multinomial Logit model (c.f. Train, 2003), or Mixed GEV models by extension (c.f. Hess, Bierlaire, Polak, 2005). The analysis of random taste heterogeneity is beyond the scope of the present analysis, whose aim it was solely to investigate the correlation structure in place in the unobserved part of utility. For this, a closed-form GEV approach is not only sufficient, but also most appropriate. As such, the use of an MMNL or Mixed GEV model would not have any advantages, as the treatment of correlation with the help of error-components (as required in the MMNL model) is in fact generally more problematic than the use of a GEV structure (c.f. Hess et al., 2005). For an application of mixture approaches to airport choice modelling, the reader is referred to Hess & Polak (2005a, 2005b). Another advantage of models based on mixing approaches is that they allow for heteroscedasticity; this again is however beyond the scope of the present alternative.

Several important issues relating to model specification deserve some further attention. These relate to the re-weighting of the sample for model estimation, the way

attributes enter the utility function, and the definition of the constants used in the model.

Given that the survey data are choice-based, some form of re-weighting needs to be performed in order for the estimation to represent population-level market shares as opposed to sample-level shares (influenced by survey quotas), thus avoiding biased results. In the present analysis, multiplicative weights were used in the specification of the log-likelihood function, where, for a given respondent, the weight is proportional to the ratio between the population weight and the sample weight for the corresponding group, where group allocation was based on a host of criteria, dominated by route and airline choice.

Another important question arises with regards to the specification of alternative specific constants (ASC), which are used to ensure a mean of zero for the unobserved utility terms. In one-dimensional choice processes, a single ASC is associated with each alternative, with all but one of the constants being estimated (normalisation ensuring identification). In the case of multi-dimensional choice processes, the situation becomes slightly more complicated. Hess & Polak (2004) advocate the use of a set of ASCs in each of the three choice-dimensions, with one normalised ASC in each group. The problem with this approach is that it ignores the potential impact of interactions between the choice-dimensions. To address this deficiency, an alternative specification was attempted in this paper, using a single constant for each airport-airline pair. This increases the total number of airport and airline related constants from 42 (37 airlines at 5 airports) to 54, given that not all airlines operate from all airports. Separate experiments showed that this approach led to very significant gains in model performance, suggesting some interaction between choice dimensions. While it is in theory possible to further improve the specification by using a separate

constant for each airport, airline and access-mode triplet, the gains from this approach are no longer significant, coming at the cost of an increase in the number of constants from 60 to 324. Attempts to use airport-access constants in combination with separate constants in the airline dimension also led to gains in model fit, which were however less significant than those obtained with the airport-airline specification, which was thus retained. Finally, in the population segment used in the present paper, resident business travellers, the number of estimated constants reduced from 58 to 51, due to the specific set of destinations observed for these travellers.

The final point that deserves some discussion is the way in which explanatory variables enter the utility function. Overwhelmingly, research in discrete choice analysis relies on the use of a linear in parameters specification. However, this is not appropriate in the case of attributes with decreasing or increasing marginal returns, where the use of a linear specification can lead to biased estimates. Common examples of such attributes include time, cost and frequency variables. Different transformations can be used to allow for non-linear marginal returns, including the basic logarithmic transform, as well as more advanced transformations such as Box-Cox. In the present analysis, we rely on the use of the log-transform, which was shown to lead to very good performance by Hess & Polak (2005a), and has been used by a number of other authors in the area of airport-choice analysis. A preliminary analysis was conducted to determine which attributes benefited from the use of a non-linear specification. These experiments showed that important gains in model performance could be obtained by using a log-transform for flight frequency, flight time, in-vehicle access-time (IVT), and access-cost, such that this approach was adopted.

5. MODELLING RESULTS

Before the actual discrete choice experiments, a brief analysis was conducted to analyse passengers' stated reasons for choosing their specific departure airport. The results, which are summarised in Table 2, show ground-level origin to be the prime determining factor for resident business travellers, with flight availability and timing unsurprisingly also being major factors; given the lack of appropriate data, these can however only be included in the form of flight frequency information. The importance of geographical proximity is further underlined by the fact that an additional 8.14% of passengers indicate proximity to their workplace as the main deciding factor. Just over a tenth of passengers stated that the decision had been taken by a third party; a model fitting exercise excluding these observations led to very comparable results, such that these observations were retained to increase the overall sample size. This result would suggest that similar choice processes apply for self-bookings and third-party bookings. Finally, the low sensitivity to air-fare by business travellers (at least in 1996) is underlined by the fact that this plays the main role for a smaller share of passengers than a personal penchant for the specific airport; this at least partly explains the difficulty in finding a significant effect of fare in this study, as well as in previous research into airport choice behaviour.

A comprehensive list of variables was used in the initial modelling analysis. These included attributes relating to the air journey (frequency, fare, flight time, aircraft type, seat capacity) and the access journey (access-cost, in-vehicle access-time, out-of-vehicle access-time, wait-time, number of interchanges, parking cost). No treatment of the distribution of departure times was used in the present analysis; this requires the development of specific methodological tools and is the topic of ongoing research. In the absence of frequent flier information, attempts were made to account

for airline allegiance by including a UK-airline dummy variable. No information on past choices was available, such that a treatment of "airport-allegiance" was not possible. Appropriate log-transforms were used, as described in the previous section. Finally, attempts to model further interactions with socio-demographic attributes aside from purpose, such as income, were not successful. The freeware estimation package BIOGEME (Bierlaire, 2003) was used for model calibration and application (validation).

The actual 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 (possibly due to the absence of cost-bearing information), 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 at least partly due to the poor quality of the data, but should be put into context by considering the observations made from Table 2. 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 (IVT), flight frequency, and flight time, where a log-transform was used for all four attributes. The list of significant attributes stays identical across models; the following three sections look at the actual results obtained with the different structures.

5.1. MNL results

The results for the MNL estimation process are summarised in Table 3. For space reasons, only the taste coefficients are shown here. It is worth noting that the majority of the estimated constants were statistically significant at high levels of confidence, which was not the case when using three disjoint sets of constants.

In addition to the estimated coefficients, robust t-statistics and model fit statistics, the table also presents a set of trade-offs. In the presence of logarithmic transforms, the calculation of trade-offs is slightly more complicated than in the case of linear specifications of utility. Given that all concerned attributes enter the utility function under a log-transform in the present analysis, we illustrate the case of a trade-off between attribute z_1 and attribute z_2, when the utility function is of the form:

$$U = \ldots + \beta_1 \ln(z_1) + \beta_2 \ln(z_2) + \ldots$$

The ratio of the partial derivatives of U with respect to z_1 and z_2 respectively is then given by $\beta_1/\beta_2 \cdot z_2/z_1$, as opposed to the simple β_1/β_2 ratio used in the case of a linear parameterisation. In the calculation of the trade-offs, the values for access-time, access-cost and flight frequency on the observed journeys were used, and statistics were calculated for the distribution of these trade-offs over respondents. In the case of the ratio of two coefficients using a log-transform, and in the presence of non-perfectly correlated variations in time, cost and frequency, this approach is preferable to the commonly adopted use of the simple mean values across observations, as it avoids potentially significant levels of bias in the calculation of trade-offs. Furthermore, this approach yields respondent-specific trade-offs, allowing the calculation of a set of statistics for the distribution of the trade-offs, where it should be noted that these variations are an effect of the varying values for the concerned attributes, and do not as such give variations in tastes across respondents, but rather

give an indication of the varying levels of trade-offs under different market conditions.

The results show significant negative changes in utility associated with increases in in-vehicle access-time, access-cost and flight-time, with positive changes in utility associated with increases in flight frequency. The first observation that can be made for the trade-offs is that the implied value of travel-time savings (VTTS), giving the willingness to pay for reductions in access-time, are markedly lower than those reported for example by Pels et al. (2003) and Hess & Polak (2004), although they are still higher than in other contexts, which can be explained partly by concepts of risk-averseness, as discussed by Hess & Polak (2004,2005a). Travellers are willing to pay for a reduction in the risk of missing their 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. While previous research in airport-choice modelling has generally made use of a log-transform for flight-frequency, access-time and access-cost have usually been treated in a linear fashion, which could have caused the high implied VTTS. 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.

The low implied willingness to accept increases in access-cost in return for increases in frequency should be put into context by noting that the average access-cost was £11.55, with an average frequency on the chosen route of 5.2 flights. For the willingness to accept increases in access-time in return for increases in flight

frequency, the implied mean value of around 6 minutes (i.e. just above 10% of the average journey time) seems realistic, though possibly on the low end of the scale. Finally, flight-time and in-vehicle access-time are measured on a similar scale; the ratio of 1.79 between the actual taste coefficients is cancelled out by a very similar ratio between the average flight-times and access-times, with great variations around the mean leading to a wide spread in the trade-off.

5.2. NL results

The next stage of the analysis consisted of fitting the three different NL structures to the dataset. The results of this estimation process are presented in two parts; we first look at the general model fit statistics and marginal utility estimates in Table 4, before moving on to the nesting parameters in Table 5.

The first observation that can be made from Table 4 is that all three nesting approaches lead to statistically significant improvements in model fit over the MNL model. Indeed, for each of the models, the likelihood-ratio test has an associated p-value of 0, as shown by the following calculations:

NL by airport: $-2 \cdot (-14945.3 + 14896.1) =$ $98.4 \sim \chi^2_4 \rightarrow p \approx 0$

NL by airline: $-2 \cdot (-14945.3 + 14870.7)$ $=$ $149.2 \sim \chi^2_{19} \rightarrow p \approx 0$

NL by access-mode: $-2 \cdot (-14945.3 + 14816.7)$ $=$ $257.2 \sim \chi^2_5 \rightarrow p \approx 0$

Though nested likelihood-ratio tests cannot be used to directly compare the performance of the different nesting structures, it can be noted that the biggest improvement in model fit over the MNL model is obtained by using nesting by access-mode, while the performance of the model using nesting by airline is rather disappointing, obtaining only relatively small improvement in model fit over the model using nesting by airport (the least-well fitting one), when taking into account the much higher number of estimated parameters.

24

In terms of the actual estimation results, all the coefficients are of the expected sign, and aside from the frequency coefficient in the model using nesting by access-mode, all coefficients are statistically significant at levels of confidence well above the usual 95% limit. In terms of trade-offs, it can be observed that the first two models yield relatively similar results, which are also close to those produced by the MNL model (aside from a lower than 1 ratio for flight-time over access-time for the model using nesting by airport). The model using nesting by access-mode however leads to a much higher relative estimate of the access-cost coefficient, which manifests itself through significantly lower willingness to accept increases in access-cost in return for reductions in access-time and increases in frequency. This applies to the mean values, as well as the minimum, maximum and standard deviation. The remaining two trade-offs seem to be largely consistent across models. These results show that, although the fit produced by the three nesting structures is relatively similar, the substantive results vary more significantly across models. This would suggest that the likelihood is relatively flat around the actual values at least for the access-cost coefficient.

In terms of the correlation, a total of 5, 37, and 6 structural parameters could be estimated in the models using nesting by airport, airline, and access-mode respectively. A number of these parameters had to be constrained to 1, as they had initially taken on unacceptable values (larger than 1); while such constrained values are reproduced in Table 5 for the models using nesting by airport and access-mode, they are omitted for the model using nesting by airline, due to space constraints. In the model using nesting by airports, the structural parameter for the LHR nest had to be constrained to a value of 1. Furthermore, only two airports, namely LGW and LTN have structural parameters that are statistically different from 1 at satisfactory levels of confidence; this illustrates the limited success of this nesting approach, which is

reflected in the modest gains in model performance when compared to the MNL model. In terms of actual behavioural conclusions, these results suggest heightened correlation between alternatives sharing LTN as their common airport, and to a lesser degree, also for alternatives sharing LGW.

For the model using nesting by airline, a structural parameter with an acceptable value could only be estimated for 19 out of the 37 airlines, and only 3 of those were statistically different from 1 at the usual 95% level of confidence. Though not entirely conclusive, the results seem to suggest a higher level of inter-alternative correlation for flights operated by low-cost carriers; this would support the claims that these carriers induce new demand, and that a non-trivial percentage of their passengers would not travel if services on the low-cost carriers were not available.

The most conclusive results are obtained for the model using nesting by access-mode, which also produces the best model fit. Here, the nesting parameter for public transport needs to be constrained to 1, while the nesting parameter for long-distance coach is not significantly different from 1. The remaining four nesting parameters however are significantly different from 1, and indicate very high levels of inter-alternative correlation (low structural parameters), reflecting high mode-allegiance. Travellers seem to be more unlikely to accept a change of access-mode than a change of airline, or airport.

5.3. CNL results

The final part of the analysis looks at fitting the CNL model described at the end of section 4. Again, the results are presented in two separate tables; Table 6 for the main estimation results, and Table 7 for the structural parameters.

As the CNL model is a generalisation of the NL model, nested likelihood-ratio tests can in theory be used in comparisons of model fit. Let alternative j be defined as the choice of airport $k(j)$, airline $l(j)$, and access-mode $m(j)$, such that:

$$p = j : \begin{cases} 0 \le \alpha_{jk(p)} \le 1 \\ 0 \le \alpha_{jl(p)} \le 1, \\ 0 \le \alpha_{jm(p)} \le 1 \end{cases} \quad and \quad p \ne j : \begin{cases} \alpha_{jk(p)} = 0 \\ \alpha_{jl(p)} = 0 \\ \alpha_{jm(p)} = 0 \end{cases}$$

The CNL model can be seen to reduce to the NL model using nesting by airport, if, after estimation:

$$\alpha_{jp} = \begin{cases} 1, if \ p = k(j) \\ 0, if \ p \ne k(j) \end{cases} \forall j,$$

the NL model using nesting by airline, if, after estimation:

$$\alpha_{jp} = \begin{cases} 1, if \ p = l(j) \\ 0, if \ p \ne l(j) \end{cases} \forall j,$$

and the NL model using nesting by access-mode, if, after estimation:

$$\alpha_{jp} = \begin{cases} 1, if \ p = m(j) \\ 0, if \ p \ne m(j) \end{cases} \forall j.$$

In the present context, a total number of 48 nests were used in the CNL model (5 airports, 37 airlines, and 6 access-modes). Aside from leading to the use of 48 separate nesting parameters (to allow for differential levels of correlation in different nests), this leads, in the presence of a choice-set of 324 elementary alternatives, to a total of 972 allocation parameters (324 along each dimension). As each alternative is associated with exactly one airport, one airline, and one access-mode, only one allocation parameter along each of the three dimension is not constrained a priori to zero for a given alternative. From this, it can also be seen that, given the condition that the allocation parameters for each alternative sum to 1, a total of 648 can be identified. Although this number is reduced somewhat due to availability conditions, this still leads to a very expensive estimation process, and can be seen to result in an

over-parameterised model. Results by Hess (2004) on the San Francisco Bay area data show that, although the estimation of the allocation parameters leads to gains in model fit, these are not statistically significant, given the huge cost in terms of the number of parameters. Similar observations were made in the present study, and accordingly, the decision was taken to constrain all non-zero allocation parameters to a value of 1/3. Efforts to find a more flexible parameterisation of the allocation parameters at an acceptable cost in terms of degrees of freedom are ongoing.

With the use of fixed allocation parameters, it is not possible for the CNL estimation process to yield one of the three NL models during estimation[7]. Given this complication, nested likelihood-ratio tests cannot be used. The models can however still be compared, using the adjusted ρ^2 statistic, which takes into account the cost of a model in terms of the number of parameters. From the values reported in Table 6, it can be seen that the CNL model offers the biggest improvement of any of the nesting models when compared to the MNL model. Furthermore, the improvement in the adjusted ρ^2 measure is in each case bigger than the improvement offered by the respective NL model when compared to the MNL model. Finally, the total improvement of the CNL model over the MNL model is bigger than the combined improvements in the 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.

In terms of substantive results, the coefficients are all of the expected sign, but the frequency coefficient is now only significant at the 87% level. In terms of trade-offs between coefficients, the willingness-to-pay for access-time reductions is higher than

that reported by the NL model using nesting by access-mode, but lower than in the remaining NL models, as well as in the MNL model. Frequency is seemingly valued less highly than in the MNL and NL models, though the high associated standard error should be noted. Finally, as was the case for the NL model using nesting by airport, flight-time is again valued less highly than in-vehicle access-time.

The final part of the results looks at the estimates for the structural parameters in the CNL model, as reported in Table 7. The findings were characterised by very low structural parameters (and hence high correlation) along all three dimensions of choice, but accompanied by high standard errors, especially for the access-mode dimension. While a higher number of parameters took on permissible values in the airline-choice dimension than previously (28 as compared to 19), an additional parameter along the access-mode dimension had to be constrained to 1, namely that for long-distance coach, which was already indistinguishable from 1 in the NL model. A major part in the success of the CNL model in this application is its treatment of correlation along the airline dimension; it retrieves correlation for all airlines where the NL model was able to do so, but also finds correlation for an additional 9 airlines. Overall, the significance levels along this dimension where also higher than in the corresponding NL model.

6. 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 with the models presented in section 5, with results summarised in Table 8. For each observation, the

---

[7] It should be noted that, independent of the allocation parameter, an approximation of the various NL structures can be obtained when all nesting structures except those along the concerned dimension, take on values of 1.

estimation software generates a choice probability for each of the 324 elementary alternatives; these can be summed up appropriately to obtain the choice probability for the different airports (5), airlines (37) and access-modes (6). From this, the probability of correctly predicting a given respondent's choices along each of the three choice dimensions can be retrieved straightforwardly, and averaging over observations yields the "average probability of correct prediction" in each of the choice dimensions. It is important to note that this is different from the unreliable "percentage right" measure, which determines implied choices on the basis of the highest choice probabilities, and calculates the percentage of correct predictions. This latter measure completely misrepresents the random component of utility (c.f. Train, 2003).

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; here the ρ^2 measures as well as the coefficient trade-offs presented in section 5 are of more interest.

Without touching on the differences between models, it is of interest to compare the results to previous research in this area. It should first be noted that the low probabilities for the elementary alternatives must be put into context by remembering that the total number of such alternatives is 324, with an average of 30 available alternatives per individual in the validation sample. The aggregate average probabilities of correct prediction are well below those obtained by Hess & Polak (2004), who obtained 70% to 85% for airport choice, 50% to 60% for airline choice, and 60% to 85% for access-mode choice. 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 in this paper 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.

7. SUMMARY AND CONCLUSIONS

This paper has described an analysis of the combined choice of airport, airline and access-mode for passengers departing from the London area, using a number of standard and more advanced discrete choice models. In common with most previous studies, the analysis shows that access-time is a prime determining factor in travellers' choices of departure airport, while flight frequency, access-cost and flight-time also play a role. At this point, it should be noted that the flight-time coefficient can also be seen as a proxy for smaller aircraft, and for on-time performance, given that the block-time incorporates taxi-time, and hence takes into account congestion. As in many previous studies, it was not possible to estimate a significant effect of air-fare, nor of airline-allegiance, a fact that is down to the general low quality of the level-of-service data for the associated attributes, while for fare, it however also signals general low sensitivity by business travellers, especially in 1996.

In terms of model performance, all attempted nesting approaches lead to significant gains in model performance. When accounting for correlation along just one dimension of choice, the best performance is obtained by the model using nesting by

access-mode. However, the simultaneous treatment of correlation along all three dimensions has clear benefits, and the Cross-Nested Logit model outperforms the three Nested Logit structures; this suggests that this model form can serve as a valuable tool in airport-choice modelling. It should be noted that bigger improvements by the nesting structures would be expected with the use of stated preference (SP) data (as opposed to revealed preference), where the analysis of substitution patterns is made considerably easier by the presence of multiple observations for individual respondents.

A number of avenues for further research can be identified, not least of which the extension to other population segments, but also the use of more advanced model structures, allowing for cross-nesting, continuous deterministic and random taste heterogeneity, and a treatment of flight scheduling. Further refinement of the auxiliary datasets can also be expected to lead to gains in model performance. Finally, aside from only accounting for correlation between alternatives sharing a given airport, airline or access-mode (or a combination thereof), it is also of interest to test for correlation between alternatives at different, yet comparable airports (e.g. LHR & LGW), or different airlines and access-modes.

In closing, it should be noted again that the results from this analysis do suggest that air-passengers, at least those travelling on business, are very reluctant to accept increases in their access-journeys, and are highly captive to 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, 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, which will again have a preference for more centrally-located airports.

ACKNOWLEDGEMENTS

REFERENCES

Ashford, N., Bencheman, M., (1987. Passengers' choice of airport: an application of the Multinomial Logit model. *Transportation Research Record* **1147**, 1-5.

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

Bierlaire, M. (2003), "BIOGEME: a free package for the estimation of discrete choice models", Proceedings of the 3rd Swiss Transport Research Conference, Ascona, Switzerland.

Boeing (2003), Current Market Outlook, Market Analysis, Boeing Commercial Airplanes, Seattle.

Brooke, A. S., Caves, R. E., & Pitfield, D. E. (1994), Methodology for predicting European Short-Haul Air Transport Demand from Regional Airports: An Application to East Midlands International Airport, *Journal of Air Transport Management* **1(1)**, 37-46.

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.

CAA (2002), UK Airport Statistics, The Civil Aviation Authority, London. [Available: http://www.caa.co.uk/erg/erg_stats]

Daly, A. & Zachary, S., (1978) Improved multiple choice models, in D. Hensher and M. Dalvi, eds, *Identifying and Measuring the Determinants of Mode Choice*, Teakfields, London.

DfT (1997), *Air traffic forecasts for the United Kingdom 1997*, Department for Transport, London.

DfT (2000a), *Air traffic forecasts for the United Kingdom 2000*, Department for Transport, London.

DfT (2000b), *SERAS Passenger Allocation Model: An Overview of the Calibration Process*, Scott Wilson Kirkpatrick & Co Ltd, Basingstoke, United Kingdom.

DfT (2002), *The future development of air transport in the UK: South East (2nd edition), The future development of air transport in the UK: a national consultation*, Department for Transport, London.

DfT (2003), *The Future of Air Transport: Government White Paper*, Department for Transport, United Kingdom Government, London.

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, France.

Hess, S., Bierlaire, M. and Polak, J.W. (2005), Capturing taste heterogeneity and correlation structure with Mixed GEV models, in *Applications of simulation methods in environmental and resource economics*, Alberini,A. and Scarpa,R. (eds), Kluwer Academic Publisher, Boston, MA.

Hess, S., Polak, J.W., (2004). On the use of Discrete Choice Models for Airport Competition with Applications to the San Francisco Bay area Airports, paper presented at the 10[th] triennial World Conference on Transport Research, Istanbul.

Hess, S. and Polak, J.W. (2005a), Mixed Logit modelling of airport choice in multi-airport regions, *Journal of Air Transport Management* **11(2)**, 59-68.

Hess, S. & Polak, J.W. (2005b), Accounting for random taste heterogeneity in airport-choice modelling, *Transportation Research Record*, forthcoming.

IATA, (2002). World Air Transport Statistics. International Air Transport Association, Geneva.

McFadden, D. (1974), Conditional logit analysis of qualitative choice behaviour, in P. Zarembka, ed., *Frontiers in Econometrics*, Academic Press, New York, pp.105-142.

McFadden, D. (1978) Modelling the choice of residential location in A. Karlqvist, L. Lundqvist, F. Snickars and J. Weibull, eds, *Spatial Interaction Theory and Planning Models*, North-Holland, Amsterdam, pp.75-96.

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

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

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

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

Skinner, R. E. Jr. (1976). Airport choice: An empirical study, *Transportation Engineering Journal* **102(4)**, 871-883.

Thompson, A., Caves, R., (1993). The projected market share for a new small airport in the North of England. *Regional Studies* **27**, 137-147.

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

Vovsha, P. (1997).. Cross-nested logit model: an application to mode choice in the Tel-Aviv metropolitan area, *Transportation Research Record.* **1607**, 6-15.

Williams, H. (1977). On the formation of travel demand models and economic evaluation measures of user benefits, *Environment and Planning* **A9**, 285-344.

Table 1
Annual passenger counts (departing and arriving, including connecting passengers) at London's main airports, 1997-2002 (in millions)

Year	LHR	LGW	STN	LUT	LCY	Total
2002	63.36	29.63	16.06	6.49	1.60	117.13
2001	60.77	31.18	13.67	6.56	1.62	113.79
2000	64.62	32.07	11.88	6.19	1.58	116.34
1999	62.27	30.56	9.45	5.29	1.39	108.95
1998	60.68	29.17	6.86	4.13	1.36	102.21
1997	58.19	26.96	5.43	3.24	1.16	94.97
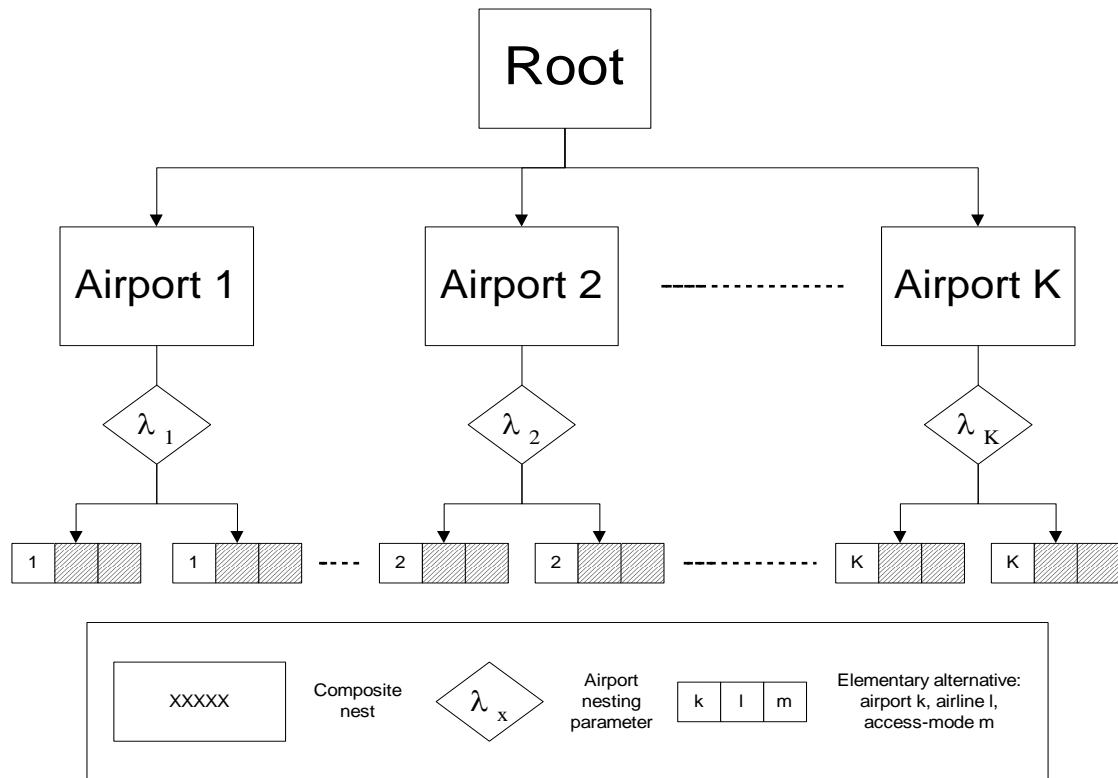
Fig. 1. Structure for models using nesting by airport
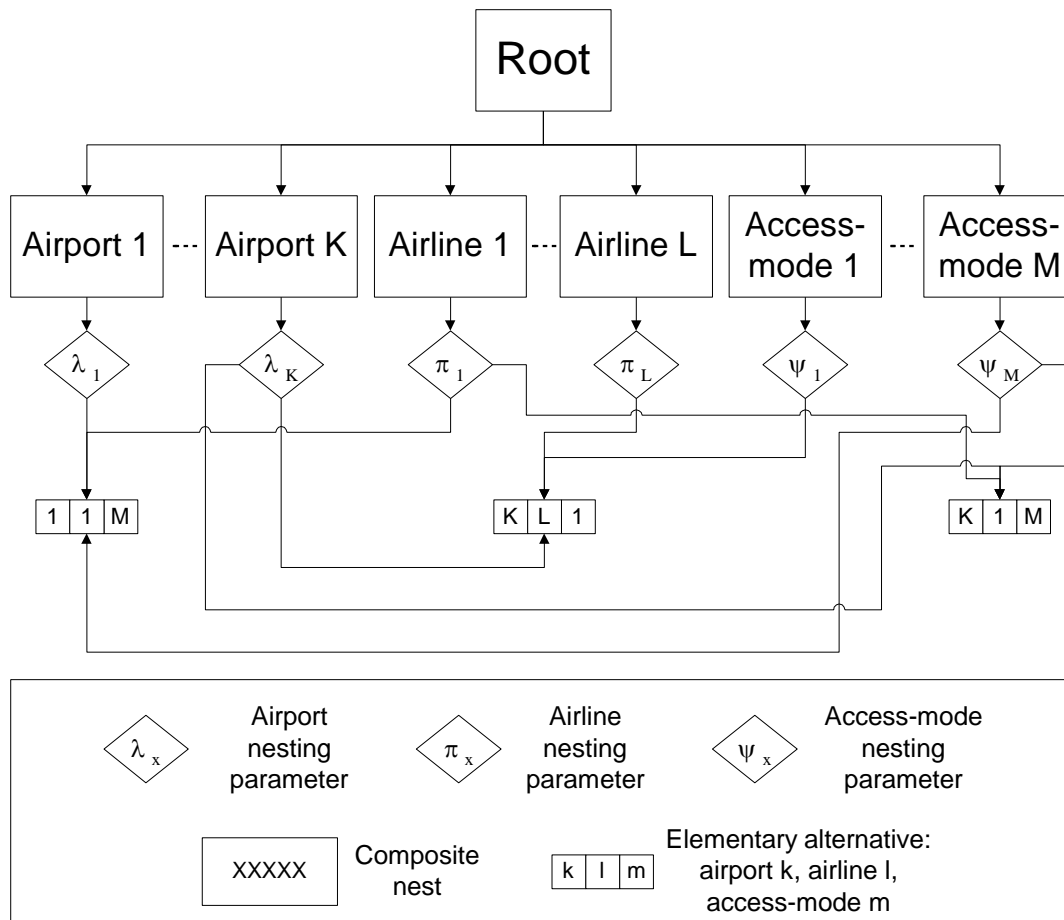
Fig. 2. Structure for cross-nested model

Table 2
Stated reasons for choice of departure airport

Main reasons	Percentage of resident business travellers
Near home	41.90%
Flights available	11.38%
No answer	11.23%
Decision someone else	11.17%
Near business	8.14%
Timing of flights	6.68%
Prefer airport	3.79%
Economic/cheaper	2.71%
Local services inadequate	1.57%
Better surface connections	0.42%
Other	1.01%

Table 3
MNL estimation results

Model statistics

	Observations	Final LL	Adjusted [8] ρ^2	Parameters
	6706	-14945.3	0.3445	55

Estimated parameters

	LN(access-cost)	LN(IVT)	LN(frequency)	LN(flight time)
Value	-1.2830	-1.4440	0.5641	-2.2963
Robust t-statistics	-8.05	-6.21	2.42	-3.17

Trade-offs

	IVT / access-cost (£/hour)	Freq / access-cost (£/flight)	Freq / IVT (hours/flight)	Flight-time / IVT
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	0.70

[8] With N giving the number of parameters of the model, LL(0) giving the log-likelihood at zero, and LL(θ) giving the log-likelihood at convergence with estimates θ, the adjusted ρ^2 measure is given by 1-[(LL(θ)-N)/(LL(0))].

Table 4
Estimation results for NL models, part 1

Model statistics

	Observations	Final LL	Adjusted ρ^2	Parameters
NL using nesting by airport	6706	-14896.1	0.3465	59
NL using nesting by airline	6706	-14870.7	0.3469	74
NL using nesting by access-mode	6706	-14816.7	0.3499	60

Estimated parameters

		LN(access-cost)	LN(IVT)	LN(frequency)	LN(flight time)
NL using nesting by airport	Value	-1.1807	-1.4610	0.5446	-2.1002
	Robust t-statistics	-7.83	-6.75	2.40	-2.91
NL using nesting by airline	Value	-1.1331	-1.3946	0.5716	-2.3415
	Robust t-statistics	-7.04	-6.27	2.50	-3.44
NL using nesting by access-mode	Value	-1.0197	-0.9553	0.3196	-1.49414
	Robust t-statistics	-7.23	-3.27	1.96	-2.95

Trade-offs

		IVT / access-cost (£/hour)	Freq / access-cost (£/flight)	Freq / IVT (hours/flight)	Flight-time / IVT
NL using nesting by airport	Minimum	1.30	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 using nesting by airline	Minimum	1.29	0.03	0.01	0.04
	Mean	17.76	1.79	0.12	1.13
	Maximum	156.80	265.11	4.26	7.85
	Standard deviation	27.83	5.57	0.19	0.74
NL using nesting by access-mode	Minimum	0.99	0.02	0.01	0.04
	Mean	13.52	1.11	0.10	1.05
	Maximum	119.35	164.74	3.47	7.31
	Standard deviation	21.18	3.46	0.15	0.69

Table 5

Estimation results for NL models, part 2, t-statistics calculated with respect to 1

NL using nesting by airport			NL using nesting by airline		
Nest	Estimate	Robust t-statistic	Nest	Estimate	Robust t-statistic
LCY	0.87	0.43	1	0.61	1.36
LGW	0.83	1.90	4	0.89	1.29
LHR	1.00	-	6	0.78	0.44
LTN	0.55	2.14	8	0.63	0.84
STN	0.76	1.27	11	0.72	1.46
			12	0.62	2.02
			14	0.78	0.39
NL using nesting by access-mode			15	0.56	1.48
Nest	Estimate	Robust t-statistic	17	0.90	0.37
Car	0.61	2.56	18	0.61	0.93
Hire	0.37	2.40	19	0.77	0.41
LDC	0.76	0.69	20	0.43	2.50
MC	0.58	2.15	21	0.49	1.77
PT	1.00	-	27	0.72	1.15
Taxi	0.64	2.53	28	0.87	0.76
			30	0.39	2.67
			31	0.66	1.42
			32	0.69	0.97
			37	0.43	0.85

Table 6
Estimation results for CNL model, part 1

Model statistics

	Observations	Final LL	Adjusted ρ^2	Parameters
	6706	-14603.9	0.3578	91

Estimated parameters

	LN(access-cost)	LN(IVT)	LN(frequency)	LN(flight time)
Value	-0.9863	-1.0908	0.2495	-1.5373
Robust t-statistics	-8.97	-6.73	1.52	-3.67

Trade-offs

	IVT / access-cost (£/hour)	Freq / access-cost (£/flight)	Freq / IVT (hours/flight)	Flight-time / IVT
Minimum	1.16	0.01	0.00	0.04
Mean	15.96	0.90	0.07	0.95
Maximum	140.89	132.96	2.38	6.59
Standard deviation	25.00	2.79	0.10	0.62

Table 7
Estimation results for CNL model, part 2, t-statistics calculated with respect to 1

	Airport nests			Airline nests (continued)	
Nest	Estimate	Robust t-statistic	Nest	Estimate	Robust t-statistic
LCY	0.55	0.95	11	0.54	0.92
LGW	0.62	2.67	12	0.29	2.23
LHR	1.00	-	13	0.89	0.27
LTN	0.26	2.35	14	0.10	3.01
STN	0.26	2.10	15	0.57	0.81
			16	0.41	1.57
	Access-mode nests		17	0.28	1.71
Car	0.04	0.74	18	0.34	0.84
Hire	0.31	1.02	19	0.11	3.42
LDC	1.00	-	20	0.10	6.34
MC	0.15	2.35	21	0.48	1.25
PT	1.00	-	24	0.75	0.71
Taxi	0.50	1.25	27	0.23	2.21
			28	0.47	1.77
	Airline nests		29	0.92	0.17
1	0.19	1.03	30	0.24	1.27
2	0.17	1.87	31	0.12	5.81
4	0.48	1.77	32	0.53	2.27
5	0.12	14.81	33	0.99	0.02
6	0.10	4.08	35	0.10	0.09
8	0.48	0.53	37	0.12	1.51
10	0.08	1.28			

Table 8
Prediction performance on validation sample

	Elementary alternatives (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%