

Accounting for random taste heterogeneity in airport-choice modelling

Stephane Hess

Centre for Transport Studies

Imperial College London

Exhibition Road

London SW7 2AZ

Tel: +44(0)20 7594 6105

Fax: +44(0)20 7594 6102

stephane.hess@imperial.ac.uk

John W. Polak

Centre for Transport Studies

Imperial College London

Exhibition Road

London SW7 2AZ

Tel: +44(0)20 7594 6089

Fax: +44(0)20 7594 6102

j.polak@imperial.ac.uk

Word count: 6,718 words + 3 Tables = 7,468 words

Keywords: airport-choice, Mixed Logit, taste heterogeneity, San Francisco Bay area

ABSTRACT

In this paper, we present the findings from a disaggregate analysis of the choice of airport, airline and access-mode for business travellers living in the San Francisco bay area. Aside from formulating the multi-dimensional choice process, the main objective is to explore random taste heterogeneity amongst decision-makers in their evaluation of the attractiveness of the different alternatives. The results indicate that this is present in tastes relating to in-vehicle access-time, access-cost and flight frequency. Accounting for this heterogeneity leads to gains in model fit, but more importantly, leads to important insights into the differences in behaviour across decision-makers, and avoids the bias introduced into trade-offs when using fixed coefficients in the presence of significant levels of heterogeneity. In terms of substantive results, the models also reveal a significant impact of changes in out-of-vehicle access-time, indicate a preference for jet-services over turboprop-flights, and show that past experience plays an important role in air-travel choice behaviour.

1. INTRODUCTION

Discrete choice models have established themselves as the preferred tool for the analysis of travel-behaviour. However, while the most advanced and flexible types of models are increasingly being used in some areas of transportation research, such as mode-choice, route-

choice, and time-of-day modelling, in the area of aviation, there is still a heavy reliance on the most basic and assumption-bound models, such as Multinomial Logit (MNL).

In this paper, we advance the state of the art in one dimension of air-transport research, the analysis of passengers' choices in multi-airport regions. With forecasts predicting a return to pre-2001 growth-levels in global air-traffic (1), pressure to expand capacity exists at many of the world's busiest airports (c.f. 2, 3). The fact that a considerable portion of these airports are part of a multi-airport region complicates the decision-making process for transportation planners, who are no longer simply interested in the traffic at an individual airport, but are also concerned with the distribution of passengers across the different airports in the region. Furthermore, the effects of any changes in policy and/or infrastructure on congestion in the ground-level network need to be taken into account, and appropriate provisions need to be made for increases in traffic to and from any airport for which capacity is increased. Given the volatility of the air-travel business, and the high-cost and long-term nature of developments in air-transport infrastructure, detailed and reliable forecasts of passenger-behaviour thus form an important part of any risk-minimising transport planning process in multi-airport regions.

It should be recognised that, aside from the a priori decisions to travel and to use air as the main mode, passengers not only make a choice of airport, but also additionally choose an airline and an access-mode. The joint analysis of these three choice-dimensions, and the interactions between them, can lead to important gains in the accuracy of any analysis of air-travel choice-behaviour. While some studies have looked at the combined choice of airport and airline (4), or the combined choice of airport and access-mode (5), few studies have jointly examined the three dimensions of airport, airline and access-mode choice, with a notable exception being the work of Hess & Polak (6).

One area of discrete choice analysis that has seen increased activity over recent years is the use of models that allow for random variations in behaviour across individuals. Accommodating such random taste heterogeneity not only often leads to important gains in model performance, but can also lead to substantively different results (c.f. 7, 8). In a recent analysis, Hess & Polak (9) have shown that random variations in tastes play a role in airport choice-behaviour in the San Francisco Bay (SF-Bay) area. This research was however restricted to the single dimension of airport-choice, and only included a very limited number of explanatory variables; in the present paper, we extend the analysis to the additional dimensions of airline and access-mode choice, and explore the use of additional explanatory variables.

The remainder of this paper is organised as follows. In the next section, we present a brief review of previous work in the area of airport-choice modelling. This is followed by a description of the data used (Section 3), and an overview of the modelling methodology used in this research (Section 4). Section 5 discusses the findings of the research, Section 6 presents the model validation exercises, and Section 7 gives a summary and conclusions.

2. LITERATURE REVIEW

In this section, we present a brief review of previous work in the area of airport choice modelling; for other reviews on this topic, see for example Pels et al. (5) and Hess & Polak (6).

One of the first studies of airport choice was conducted by Skinner (10), using an MNL model for airport choice in the Baltimore-Washington DC area. The research identifies flight frequency and ground accessibility to be the main determining factors, with travellers being more sensitive to the latter. These results were later repeated by Windle & Dresner (11), using an MNL model on a more recent version of the choice-data. This analysis also shows an important impact of past

experience, with travellers being relatively more likely to choose an airport they have chosen in the past.

The SF-bay area has been the subject of a substantive number of studies of airport choice, mainly thanks to the availability of very good data. An early example is given by the MNL-analysis of Harvey (12), who finds access-time and flight frequency to be significant determinants of choice. More recently, Pels et al. (4) have used a Nested Logit (NL) model in an analysis of the combined choice of airport and airline, finding that for both business and leisure travellers, it is preferable to nest airport-choice above airline-choice. In a later NL study, Pels et al. (5) show that airport-choice should similarly be nested above access-mode choice. To recognise the fact that the different airports are not necessarily all seen as viable alternatives by all travellers, Basar & Bhat (13) use a two-level modelling structure in which airport-choice is preceded by a choice-set generation stage. The results indicate that flight frequency is the most important aspect in choice-set composition, while access-time is the dominating factor in the actual choice of airport. In a predecessor to the analysis presented in this paper, Hess & Polak (9) have shown that random taste heterogeneity plays a role in airport-choice behaviour, most notably in the sensitivity to access-time. Finally, in an analysis of the joint choice of airport, airline and access-mode, Hess & Polak (6) have recently shown important differences across population groups in the correlation structures in place in the choice-set of alternatives, and have noted that, in general, the highest level of correlation exists between alternatives sharing the same access-mode. This study has also highlighted the advantages of explicitly accounting for the three-dimensional nature of the choice-process.

Another area that has seen repeated interest in terms of studies of airport choice is the United Kingdom. In an MNL analysis of the choice between five airports in England, Ashford & Bencheman (14) find access-time and flight frequency to be significant factors, with flight fares having an impact only for domestic passengers and international leisure travellers. In other studies, Ndoh et al. (15) find that the NL model outperforms the MNL model in a study of passenger route choice in central England, while, in an MNL analysis of the projected market share for a new airport in North England, Thompson & Caves (16) find access-time, flight frequency and aircraft-size (comfort) to be significant factors. In an MNL analysis of the distribution of passengers between airports in the Midlands, Brooke et al. (17) find flight frequency to be the most important factor. Finally, in a study of airport choice in London, Hess & Polak (18) have repeated findings from their San Francisco Bay area work highlighting the benefits of using a three-dimensional choice-structure, and have also shown the advantages of cross-nested model structures for accounting for correlation along each of the three choice-dimensions.

In studies from around the world, Ozoka & Ashford (19) use an MNL model to forecast the changes in market shares following the construction of a third airport in a multi-airport region in Nigeria. The results show high sensitivity to access-time, making the choice of location and the provision of good ground-access facilities important determinants in the planning process. Innes & Doucet (20) use a binary logit model for airport-choice in Canada, notably showing that travellers prefer jet services to turboprop flights. In a rare combined study of the choice of departure and destination airport, Furuichi & Koppelman (21) estimate an NL model that shows significant effects for access-time, access-cost and flight-frequency. Finally, the Integrated Airport Competition Model of Veldhuis et al. (22) shows that passenger behaviour is represented most appropriately by a sequential structure that models the choice of main mode followed by

the combined choice of airport and air-route, and finally the choice of access-mode at the chosen airport.

3. DATA

The multi-airport region used in the present analysis is the San Francisco Bay area, served by three major airports; San Francisco International (SFO), San Jose Municipal (SJC) and Oakland International (OAK). Of these, SFO is the largest, with, in 1995, some 15 million passengers (~55.8%), followed by OAK, with some 7.7 million passengers (~28.6%), and SJC with 4.2 million passengers (~15.6%). Forecasts by the Metropolitan Transport Commission (23) predict significant increases in traffic, which will lead to problems with capacity, and different expansion schemes are currently under consideration (2).

The dataset used in the present analysis was assembled by Hess & Polak (6) from three separate sources. Data on individual travellers' choices of airport and airline were obtained from the 1995 Airline Passenger Survey conducted by the MTC in August and October 1995, containing information on over 21,000 departing air-travellers (24). From this dataset, Hess & Polak (6) produce a usable dataset of 5,091 passengers, using 14 separate destinations in the continental United States. In common with most previous studies, only departing passengers were included in this dataset, and passengers on connecting flights and flights with a stopover were also excluded from the analysis. For this paper, a subsample of 1,212 resident business travellers was used; this was divided into an estimation sample of 1,098 observations, and a validation sample of 114 observations. For both subsamples, weights were calculated to correct for the sampling bias introduced during data-collection (c.f. 6); these weights were taken into account in model estimation and validation by an appropriate specification of the log-likelihood function. The data used are summarised in Table 1, which highlights the sampling bias in the survey data, with SJC carrying more weight than SFO, despite the fact that actual traffic levels at SFO were over three times higher than at SJC in 1995. The specific choice of destinations had an effect on the market shares for the different airlines, and this was taken into account in the calculation of weights. Special care is required in the presence of destinations that are themselves located in multi-airport regions; for a discussion of this issue in the context of the destinations used in this survey, see Hess & Polak (9).

The air-passenger survey dataset contains information on the observed choices; this needs to be complemented by information on the unchosen alternatives, along each of the three choice dimensions. Two main datasets were used to this extent; air-travel level-of-service data were obtained from *BACK Aviation Solutions* (www.backaviation.com), while ground-transport level-of-service information was obtained from the MTC in the form of origin-destination travel time and cost matrices for the 1099 travel area zones (TAZ) used for the SF-Bay area. The air-travel dataset contains disaggregate daily information for each operator serving the selected routes in August and October 1995; eight airlines were used in the analysis, where not every airline operates from each airport to all 14 destinations used. The dataset contains the daily frequencies for the different operators (weekday-dependent), in addition to flight times and the type of aircraft used. Information is also available on the average fares paid on a given route operated by a given airline; this clearly involves a great deal of aggregation, as it makes no differences between fare classes and essentially makes the assumption of similar booking rates on all routes (c.f. 9). Given the poor quality of the available fare data, an important avenue for further research is the modelling of fare availability, as described by Battersby (25). Finally, information on the on-time performance of the different airlines, as well as the overall on-time performance of

airlines at the three airports (available from the Bureau of Transportation Statistics) was appended to the dataset. The ground-transportation dataset contains information on travel distance, travel time and tolls for car travel, under peak and off-peak conditions, and for varying car-occupancy (which has an impact on tolls and the eligibility to use car-pool lanes). Similarly, the dataset contains information on access-time, wait time, travel time, egress time and fares for public transport journeys. All information is available for peak and off-peak conditions, and this was taken into account in the specification of the choice-set for the different individuals included in the sample. Corresponding values for other modes, such as taxi, limousine and special airport bus services were calculated separately, based on current prices and the changes in the Consumer Price Index for California from August and October 1995 to September 2003. Due to data limitations (c.f. 6), no difference could be made between rental cars and private cars, and parking-cost could not be taken into account. Attempts to explicitly account for the marginal running costs for car journeys were unsuccessful, such that only toll-costs were included in the modelling analysis. Six access-modes were used in the analysis; car, public transport (transit), scheduled airport bus services, door-to-door services, taxi and limousine, where it was assumed that car, taxi and limousine are available for each origin, while the availability of the remaining three modes was location-dependent.

The final dataset thus contains information for three departure airports, eight airlines, and six access-modes, leading to 144 distinct triplets of alternatives. Given the three-dimensional choice set, any given alternative shares the attributes of 73 other alternatives along exactly one dimension of choice, and 14 alternatives along exactly two such dimensions. For more details on the assembly of the final dataset, and the treatment of availabilities, see Hess & Polak (6).

4. THE MIXED MULTINOMIAL LOGIT MODEL

In a random utility model, the gain that a decision-maker can expect to obtain from choosing a specific alternative i from a choice-set of I alternatives is given by a variable U_i . In theory, the alternative with the highest utility (gain) is always chosen. However, due to modelling uncertainty, only part of the utility of an alternative is observed, and the variable U_i is accordingly divided into an observed utility V_i , and a remaining, unobserved part of utility, ε_i . With the resulting random nature of U_i , the choice now becomes probabilistic, with the alternative with the highest observed utility having the highest probability of being chosen. The observed utility V_i is a function of the tastes of the decision-maker, β , and the attributes of the alternatives, grouped into a vector x_i , which can also contain socio-demographic characteristics of the decision-maker. Typically, a linear-in-variables specification is used, such that, for decision-maker n , $V_{ni} = \beta' x_{ni}$. Different assumptions about the distribution of ε_i lead to different model forms. In the basic Multinomial Logit (MNL) model (26), the individual error-terms ε_i ($i=1, \dots, I$) are assumed to be distributed identically and independently following a type I extreme value distribution. This leads to a convenient form for the choice-probability for alternative i for decision-maker n given by:

$$P_{ni}(V_n) = \frac{e^{V_{ni}}}{\sum_{j=1}^I e^{V_{nj}}}, \quad \dots [1]$$

where this probability is independent of the unobserved utility terms. The MNL model has the disadvantage that all alternatives depend on each other in the same way, leading to unrealistic substitution patterns (c.f. 27). In more advanced models, correlation across alternatives in the

unobserved utility-components is taken into account, as for example in the NL model (28, 29, and 30).

In the Mixed Multinomial Logit (MMNL) model (c.f. 31, 27), the vector V_i ($i=1, \dots, I$) is itself a random variable. The choice probabilities are rewritten as:

$$P_{ni} = \int_{V_n} P_{ni}(V_n) dV_n, \quad \dots[2]$$

and the elements in the vector V_n can be rewritten as $V_{ni}=g(\beta, x_{ni})+\zeta_{ni}$, where ζ_{ni} is an additional error-term in the model (aside from the extreme-value term). In the present analysis, we exploit this formulation to allow for a random distribution of tastes across decision-makers, leading to a model form generally referred to as the Random Coefficients Logit model. For details on the mathematically identical, yet conceptually different Error-components Logit formulation, see Walker (32). In the random-coefficients formulation, β is assumed to be distributed according to $f(\beta|\Omega)$, where Ω is a vector of parameters of the distribution $f()$, giving for example the mean and standard deviation of the individual elements in β across decision-makers. This notation allows for some elements of β to be fixed; only a single parameter in the vector Ω will be associated with such taste-coefficients. The choice probabilities in the RCL formulation of the MMNL model are given by:

$$P_{ni} = \int_{\beta} P_{ni}(\beta, x_{ni}) f(\beta|\Omega) d\beta = \int_{\beta} \frac{e^{g(\beta, x_{ni})}}{\sum_{j=1}^I e^{g(\beta, x_{nj})}} f(\beta|\Omega) d\beta \quad \dots[3]$$

Except in the case of a trivial distribution function for β , the integrals representing the choice-probabilities do not have a closed-form solution, and numerical techniques, typically simulation, need to be used in the estimation and application of the model, such that the choice-probabilities P_{ni} are replaced by the simulated choice probabilities (c.f. 27). Despite improvements in computer performance, the reliance of standard simulation methods on the use of a very high number of draws can still make the estimation of complex models computationally expensive and impractical. Improvements in simulation performance can be obtained by using alternatives to the pseudo-random numbers that Monte-Carlo simulation processes are typically based on. These quasi-random number sequences offer improved coverage of the area of integration, thus allowing for stable estimation with a lower number of draws, leading to reductions in the computational overhead in estimation and application (see for example 33).

A second issue with the MMNL model is the choice of distribution to be used for the random taste coefficients, especially in the case where an *a priori* assumption exists about the sign of a given coefficient. For a discussion of this issue, see for example (8, 7, 34). In the present analysis, the Normal distribution was used for all randomly distributed taste coefficients, mainly for reasons of simplicity, but also because problems with slow convergence were encountered when using bounded distributions. It should however be stressed again that the Normal distribution is not fully appropriate in the case of coefficients with an *a priori* sign assumption (e.g. time and cost coefficients), given that, by being unbounded, it can, depending on the parameter estimates, indicate a high probability of a counter-intuitively signed coefficient even in the case where such values do not actually exist in the population. In the case where the main aim is to explore the prevalence of taste heterogeneity, without making policy recommendations based on the tail behaviour of the distribution, the Normal distribution may however safely be used, as it generally leads to acceptable performance in terms of recovering the true mean and

standard deviation of the coefficients. Nevertheless, researchers should in this case be mindful not to infer any conclusions based on the implied tail behaviour, particularly in the case of results indicating significant shares of individuals with counter-intuitively signed coefficients. In the present analysis, the probability of counter-intuitively signed coefficients was always at an acceptably low level, such that the Normal could be used.

5. EMPIRICAL ANALYSIS

The specification used for the utility function is based on previous work by Hess & Polak (6), using the same data. In that research, considerable effort went into finding interactions between taste coefficients and socio-demographic characteristics such as income, with the aim of describing taste heterogeneity in a deterministic way. This approach however only led to limited success. In the present analysis, we re-use the same explanatory variables, but attempt to explain any taste heterogeneity in a purely random way. The results of this research can then be used to facilitate the search for deterministic variations in taste. Indeed, for reasons of interpretation, a deterministic treatment of taste heterogeneity is always preferable, and as such, the MMNL model can be seen as an exploratory tool for finding significant levels of taste heterogeneity.

The set of coefficients for which a significant effect could be identified for resident business travellers includes access-cost, in-vehicle access-time, out-of-vehicle access-time, and flight-frequency. In addition, past experience was found to play a role, and a significant effect was identified for a dummy variable for turboprop flights. No significant effect could be identified for airfares; this could be a result of the poor quality of the fare-data, but may also be seen as a sign of the relative insensitivity of business travellers to fare increases, especially in 1996. For a discussion of the relative importance of airfare compared to other explanatory variables, see Hess & Polak (18). Other factors for which no significant effect could be identified include block time (which indirectly takes into account airport congestion), wait-time (for access-modes), and the on-time performance of the different airlines and airports. No information on the specific timing of flights was used in the models; in conjunction with an assumption of a relatively even spread of flights across the day, the frequency variable can thus be seen as a proxy for the expected difference between the desired departure-time and the best-possible departure-time. Finally, it was not possible to explore the impact of airline-allegiance programmes on travellers' preferences, for data reasons.

A coefficient that deserves some further explanations is that associated with past experience. Information was available on the number of flights a given traveller had taken from each of the three SF-bay airports in the twelve months preceding the survey. For each of the airports, a coefficient was included in the model that shows the marginal utility of past experience at the given airport. To account for the potential impact of past experience on another airport's utility, cross-effect coefficients were similarly included in the models, where, of the six possible cross-effect coefficients, three had to be constrained to zero for identification reasons (c.f. 6). As expected, the inclusion of these six coefficients led to very important gains in model fit, significantly improving the explanatory power of the models. It should be noted that the inclusion of these coefficients could lead to potential problems with endogeneity, as past choice can be seen to be highly correlated with other explanatory variables, as well as unobservables. This could lead to important problems in the case where the model was used in forecasting under hypothetical choice scenarios. As this was not the purpose of the present analysis, and as the values of the remaining coefficients remained largely unaffected, the inclusion of these experience terms was deemed not to have any detrimental effect on the validity of the models.

It should be recognised that several of the explanatory variables used in the model potentially interact with utility in a non-linear fashion. This is especially the case for flight-frequency and past experience, where decreasing marginal returns should be expected. For both of these variables, a log-transform was used, leading to important gains in model performance. A linear specification was used for all remaining factors.

First, a simple MNL model was fitted to the data, using the utility specification described above. The results of this estimation are summarised in the first part of Table 2. Due to space constraints, the 17 alternative-specific constants (ASC) estimated on the data (2 airports, 7 airlines and 5 access-modes) are not reproduced in the table; detailed results are available from the first author on request. Given the high number of ASCs, it could be argued that the model is overloaded with constants; however, their inclusion led to very significant improvements in prediction performance, with no significant effects on substantive results, such as willingness-to-pay indicators.

As expected, the results show negative impacts of increases in access-cost and access-time (in-vehicle and out-of-vehicle) along with a positive effect of increases in flight frequency. The results further show a negative value for the dummy associated with turboprop flights, showing that passengers prefer jet-services. In terms of past experience, the models show positive direct effects for all three airports, where the value for SFO is markedly lower, which could be a reflection of the high congestion at that airport. The cross-effects of experience at SJC and OAK on the choice-probability of SFO are positive, while the cross-effect of past experience at SFO on the utility of SJC is not statistically significant. At this point, it should also be noted that the estimates for the out-of-vehicle access-time coefficient and the cross-effect of SJC-experience on SFO are only significant at the 91% and 93% levels respectively. The implied values of time are very high, especially so for the out-of-vehicle access-time. Such high values have however also been reported for example by Pels et al. (5); this could be a reflection of risk-averseness, in that travellers associate longer access-journeys with a higher risk of missing their flight. The values should also be put into context by noting that the average access-journey durations were around 30 minutes of in-vehicle time, and under 5 minutes of out-of-vehicle time (due to the high market-share for car).

To assess the relative value of increases in flight frequency, the trade-offs between such increases and corresponding increases in access-time were calculated for one additional flight, at different base frequencies. The results clearly show the decreasing marginal returns of increases in frequency with higher base frequencies. Similar decreasing returns are observed for the willingness to accept increases in out-of-vehicle access-time. The implied willingness to accept access-time increases may seem rather low, especially for out-of-vehicle time, but should be taken into context by remembering the average access-journey times reported above.

In the MMNL model, significant random taste-heterogeneity was identified for in-vehicle access-time, access-cost and flight frequency. Initially, significant heterogeneity was also identified for the out-of-vehicle access-time coefficient; this however led to unrealistically high mean values, and a very high implied willingness to pay for out-of-vehicle access-time reductions. This could again be seen as a result of the low market share for modes with a significant out-of-vehicle access-time component. A possible remedy to this problem would be to identify an out-of-vehicle component for the car-mode (such as walking-time from the car park); this was however not possible with the present dataset. Due to the biased results described above, it was decided to keep the out-of-vehicle travel-time coefficient fixed in the model results reported in this paper.

Finally, no significant heterogeneity could be identified for any of the past experience coefficients.

The results from the estimation are summarised in the second part of Table 2. The results show that, with only three additional parameters, the MMNL model leads to a statistically significant (yet modest) improvement in log-likelihood by 30.95 units (giving a LR test-value of $61.9 \sim \chi_3^2$, with a 95% critical value of 7.82, and a p-value of $2 \cdot 10^{-13}$) when compared to the MNL model. The standard deviations for the access-cost and in-vehicle access-time coefficients are significant at high levels of confidence, while the standard deviation of the log-frequency coefficient is significant only at the 90% level. The same level of confidence applies for the fixed out-of-vehicle access-time coefficient. The results confirm that the effects of using the Normal distribution are benign in this application, with probabilities of *wrongly* signed coefficients of 1%, 1% and 1.5% for the access-cost, in-vehicle access-time and log-frequency coefficients respectively. In the calculation of trade-offs involving randomly distributed coefficients, 1,000,000 random draws were used per coefficient, where the lower and upper two percentiles were removed to reduce the impact of outlying values. Correlation between coefficients was minimal, so that there was no need to incorporate it in the calculations. In terms of the willingness-to-pay for access-time reductions, Table 2 shows that the use of the MMNL model leads to lower mean values of time for in-vehicle travel, and especially for out-of-vehicle travel. In both cases, the results show high standard deviations, reflecting the differences across people in their sensitivities to access-time and access-cost changes. For the willingness to accept increases in access-time in return for increases in frequency, the results show slightly higher mean trade-offs than in the MNL model, again reflecting the bias that can arise when not accounting for random taste heterogeneity. Overall, the results from the MMNL analysis show that modest gains in model performance can be obtained by accounting for random taste heterogeneity. More importantly however, this approach can significantly reduce the risk of biased trade-offs, besides giving some insight into the distribution of such trade-offs across individuals.

6. MODEL VALIDATION

Perhaps the most important indicator of a model's performance is its ability to correctly predict the choices in a sample not used in the actual model calibration process. For this purpose, the sample of 114 observations retained for model validation was used in model application runs, producing, for each observation, choice probabilities for each of the 144 elementary alternatives. These probabilities were used to calculate the average choice probability across individuals for the actual chosen alternative, giving the "average probability of correct prediction". In the context of the application described in this paper, it is of interest to also give the average probability of correctly predicting the choice of airport, the choice of airline, and the choice of access-mode. These three separate measures can be obtained by combining appropriate sets of elementary choice probabilities. Additionally, the predicted market shares for the different airports, airlines and modes were calculated, and compared to the observed market shares; using the root mean squared error (RMSE) over the different alternatives in each group. These processes were repeated for the MNL and MMNL models presented in section 5, and the results are summarised in Table 3. The results show that the differences between the MNL and MMNL model are minor, as could be expected on the basis of the difference in log-likelihood reported in Table 2. Overall, the MMNL model performs slightly better, but is outperformed marginally by the MNL model in the mode-choice dimension. However, it is not clear a priori what measure of

standard error should be associated with these differences, so that, given their scale, no reliable inferences can be made about the differences in model performance.

For both models, the results show very good prediction performance, with an average probability of correct prediction for airports and access-modes over 80%. The correct prediction rate for airlines is lower; this can be seen as a result of the absence of airline-specific attributes, such as data on frequent-flier programmes. The correct prediction rates are consistent with previous experience using the same data (c.f. 6) and illustrate the high explanatory power of the attributes included in the present research. The results also compare favourably to the rate for correct airport-choice prediction of 74.9% obtained recently by Basar & Bhat (13). Overall, the rates obtained by Hess & Polak (6) were marginally higher; this can be seen as an effect of the more explicit population segmentation used in that research.

Finally, the prediction rates are significantly higher than those reported by Hess & Polak (18) for the London area, using similarly advanced models, and even more detailed data. This is a reflection of the higher levels of competition in the London area, which make the choice process less deterministic. This is specifically caused by a wider modal-split and the fact that the major airports are all at similar distances from the main urban centre.

7. SUMMARY AND CONCLUSIONS

In this paper, we have conducted a detailed analysis of the combined choice of airport, airline and access-mode for resident business travellers in the San Francisco Bay area. The results have confirmed results from earlier research showing the high explanatory power of the access-time and flight frequency attributes. Furthermore, the analysis has repeated the results of Windle and Dresner (11), showing a significant effect of past experience on current choices. Significant effects were also identified for access-cost, out-of-vehicle access-time and a turboprop dummy variable. It was not possible to estimate a significant impact of fare; this is consistent with similar problems faced in earlier work by Hess & Polak (6) when using the same population segment. Similar problems were experienced by Basar & Bhat (13), while in some research it was possible to estimate a significant fare coefficient only for some subgroups of the population (e.g. 6, 18).

The overall results suggest very high sensitivity to access-time increases, which could be seen as a sign that travellers associate long access-times with an increased risk of missing a flight. This is also reflected in the low willingness to accept increases in access-time in return for increases in flight frequency. Even higher values of in-vehicle access-time are reported by Pels et al. (5), while a very low willingness to accept increases in access-time in return for increases in flight frequency is obtained by Basar & Bhat (13). At this point, it should be stressed again that the estimated coefficients are likely to capture the effects of a range of correlated, yet unmodelled attributes; this is however not avoidable, given the quality of the data.

The main aim of the application presented in this paper was to test for the prevalence of random taste heterogeneity in a sample population of air-travellers. Significant heterogeneity could be identified for the in-vehicle access-time coefficient, the flight frequency coefficient, and the access-cost coefficient. While allowing for such variation leads only to marginal (yet significant) gains in model fit, it avoids the bias in trade-offs resulting from the use of fixed coefficients in the MNL model. It also leads to important insights into the differences in choice-behaviour across individuals. The fact that the gains in model fit are smaller than expected should be seen as a sign of the good performance of the MNL model, rather than poor performance of the MMNL model, this being a result of the use of a detailed specification of the utility function.

Several important avenues for further research can be identified. First, it would be of interest to attempt to explain the random taste heterogeneity retrieved in the present analysis in a deterministic way. As such, the MMNL model can be seen as an exploratory tool. Next, the present analysis should be extended to other population segments. Additionally, the use of other statistical distributions should be attempted, as described in section 4. Finally, a host of other explanatory variables potentially have an impact on airport-choice behaviour; the analysis of the impact of such variables (e.g. parking cost, frequent flier membership) is a further important avenue for future research; in many cases, this will however lead to a requirement for more detailed passenger surveys and better auxiliary datasets.

ACKNOWLEDGEMENTS

The authors would like to thank the MTC, especially Chuck Purvis, for making the SF-bay area passenger survey data available for this research. The authors would also like to thank John Weber of BACK aviation for offering support on the air-level of service data.

REFERENCES

1. Boeing (2003), Current Market Outlook, Market Analysis, Boeing Commercial Airplanes, Seattle.
2. RAPC (2000), Regional Airport System Plan Update 2000, Regional Airport Planning Committee, Oakland, CA.
3. DFT (2003), The Future of Air Transport, Government White Paper, UK Department for Transport, London.
4. 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*, vol. 35, no. 1, pp. 1-9.
5. Pels, E., Nijkamp, P., & Rietveld, P. (2003), "Access to and competition between airports: a case study for the San Francisco Bay area", *Transportation Research*, vol. 37A, pp. 71-83.
6. 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 10th triennial WCTR Conference, Istanbul, 2004.
7. Hensher, D. & Greene, W. H. (2003), "The Mixed Logit Model: The State of Practice", *Transportation* 30(2), pp. 133-176.
8. Hess, S. & Polak, J.W. (2004), "Mixed Logit estimation of parking type choice", paper presented at the 83rd Annual Meeting of the Transportation Research Board, Washington, DC.
9. Hess, S. & Polak, J.W. (2005), "Mixed Logit modelling of airport choice in multi-airport regions", *Journal of Air Transport Management*, 11(2), pp.59-68.
10. Skinner, R. E. Jr. (1976), "Airport choice: An empirical study", *Transportation Engineering Journal*, vol. 102, no. 4, pp. 871-883.
11. Windle, R. & Dresner, M. (1995), "Airport choice in multi-airport regions", *Journal of Transportation Engineering*, vol. 121, no. 4, pp. 332-337.
12. Harvey, G. (1987), "Airport choice in a multiple airport region", *Transportation Research*, vol. 21A, no. 6, pp. 439-449.
13. Basar, G. & Bhat, C. R. (2004), "A Parameterized Consideration Set model for airport choice: an application to the San Francisco Bay area", *Transportation Research* 38B, pp. 889-904.
14. Ashford, N. & Bencheman, M. (1987), "Passengers' choice of airport: an application of the Multinomial Logit model", *Transportation Research Record*, vol. 1147, pp. 1-5.
15. Ndoh, N. N., Pitfield, D. E., & Caves, R. R. (1990), "Air transportation passenger route choice: a Nested Multinomial Logit analysis," in *Spatial Choices and Processes*, M. M. Fisher, P. Nijkamp, & Y. Y. Papageorgiou, eds., Elsevier Science Publishers B.V., North Holland, pp. 349-365.
16. Thompson, A. & Caves, R. (1993), "The projected market share for a new small airport in the south of England", *Regional Studies*, vol. 27, no. 2, pp. 137-147.

17. 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*, vol. 1, no. 1, pp. 37-46.
18. Hess, S. & Polak, J.W. (2005), "Exploring the potential for cross-nesting structures in airport-choice analysis: a case-study of the Greater London area", CTS Working paper, Centre for Transport Studies, Imperial College London.
19. Ozoka, A. I. & Ashford, N. (1989), "Application of disaggregate modelling in aviation systems planning in Nigeria: a case study", *Transportation Research Record*, vol. 1214, pp. 10-20.
20. Innes, J. D. & Doucet, D. H. (1990), "Effects of access distance and level of service on airport choice", *Journal of Transportation Engineering*, vol. 116, no. 4, pp. 507-516.
21. Furuichi, M. & Koppelman, F. S. (1994), "An analysis of air traveler's departure airport and destination choice behaviour", *Transportation Research*, vol. 28A, no. 3, pp. 187-195.
22. Veldhuis, J., Essers, I., Bakker, D., Cohn, N., & Kroes, E. (1999), "The Integrated Airport Competition Model, 1998", *Journal of Air Transportation World Wide*, vol. 4, no. 2.
23. MTC (2000), *Aviation Demands Forecast: Executive Summary*, Metropolitan Transportation Commission, Oakland, California.
24. MTC (1995), "Metropolitan Transportation Commission Airline Passenger Survey: Final Report", prepared by J.D. Franz Research, Oakland, CA.
25. Battersby, B. (2003), "Quasi-maximum likelihood estimation of censored multi-class air-travel demand in regional Australia", paper presented at the *2003 Australian Conference of Economists*, October 2003, Canberra.
26. McFadden, D. (1974), "Conditional logit analysis of qualitative choice behaviour", in P. Zarembka, ed., *Frontiers in Econometrics*, Academic Press, New York, pp.105-142.
27. Train, K. (2003), *Discrete Choice Methods with Simulation* Cambridge University Press, Cambridge, MA.
28. Williams, H. (1977), "On the formation of travel demand models and economic evaluation measures of user benefits", *Environment and Planning A*9, 285-344.
29. 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, 1978.
30. 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.
31. McFadden, D. & Train, K. (2000), "Mixed MNL Models for discrete response", *Journal of Applied Econometrics*, vol. 15, pp. 447-470.
32. Walker, J. L. (2001), "Extended Discrete Choice Models: Integrated Framework, Flexible Error Structures, and Latent Variables", PhD thesis, MIT, Cambridge, MA.
33. Hess, S., Train, K., & Polak, J.W. (2005), "On the use of a modified latin hypercube sampling (mlhs) approach in the estimation of a Mixed Logit model for vehicle choice", *Transportation Research B*, *forthcoming*.
34. Hess, S., Bierlaire, M. & Polak, J.W. (2005), "Estimation of value of travel-time savings using Mixed Logit models", *Transportation Research A*, 39(2-3), pp. 221-236.

		Destination airport														
		BURBANK, CA	CHICAGO, O'HARE, 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
<i>Estimation sample</i>																
Departure Airport	SFO	16	21	7	15	7	47	3	11	24	11	1	4	39	29	235
	SJC	37	16	29	25	26	79	30	64	27	20	35	24	59	20	491
	OAK	65	1	6	4	14	88	33	40	11	21	5	8	36	40	372
Total		118	38	42	44	47	214	66	115	62	52	41	36	134	89	1098
<i>Validation sample</i>																
Departure Airport	SFO	0	0	0	1	2	7	1	0	1	3	0	1	1	1	18
	SJC	3	1	7	2	5	5	2	5	1	7	8	5	9	1	61
	OAK	6	0	2	0	0	14	0	2	1	3	0	0	4	3	35
Total		9	1	9	3	7	26	3	7	3	13	8	6	14	5	114

Table 1: Data used in the analysis

		MNL		MMNL	
		Estimate	t-test	estimate	t-test
Access-cost	μ	-0.0354	-2.93	-0.0677	-5.63
	σ	-	-	0.0293	3.75
In-vehicle access-time	μ	-0.0587	-9.68	-0.0832	-7.54
	σ	-	-	0.0384	3.04
Out-of-vehicle access-time	μ	-0.1406	-1.71	-0.1389	-1.65
Log-frequency	μ	1.3365	7.55	1.6229	4.74
	σ	-	-	0.7498	1.65
Turboprop	μ	-7.8828	-9.68	-9.4282	-7.88
<i>Effect of past experience</i>	OAK on OAK	1.7041	7.72	1.8362	6.78
	SFO on SFO	0.9978	4.70	1.1316	5.03
	SJC on SJC	1.9755	6.70	2.1221	5.87
	OAK on SFO	0.6200	2.82	0.6923	2.55
	SJC on SFO	0.5939	1.84	0.6428	1.72
	SFO on SJC	-0.0397	-0.19	-0.0328	-0.14
<i>Value of time (\$/hour)</i>					
In-vehicle access	μ	99.49		93.27	
	σ	-		80.67	
Out-of-vehicle access	μ	238.30		155.70	
	σ	-		109.93	
<i>Trade-off between frequency and in-vehicle time increases (min/flight), standard deviations in brackets</i>					
One flight at base frequency of 1		22.77		23.61 (18.23)	
One flight at base frequency of 5		4.55		4.72 (3.65)	
One flight at base frequency of 10		2.27		2.36 (1.82)	
<i>Trade-off between frequency and out-of-vehicle time increases (min/flight), standard deviations in brackets</i>					
One flight at base frequency of 1		9.51		11.69 (4.80)	
One flight at base frequency of 5		1.90		2.34 (0.96)	
One flight at base frequency of 10		0.95		1.16 (0.48)	
Observations		1,098		1,098	
Log-likelihood		-1576.31		-1545.36	
ρ^2		0.5702		0.5787	

Table 2: Estimation results for resident business travellers

	MNL	MMNL
	Probability of correct prediction	
Elementary alternative	43.77%	44.78%
Airport	81.73%	82.51%
Airline	58.86%	59.52%
Access-mode	82.52%	82.33%
	RMSE between observed and true market shares	
Airport	0.0681	0.0619
Airline	0.0544	0.0522
Access-mode	0.0221	0.0232

Table 3: Model validation