Modelling air travel choice behaviour

Stephane Hess*

September 15, 2008

Abstract

The number of studies aiming to model air travel choice behaviour has increased over recent years, as discrete choice methods for analysing disaggregate travel behaviour have become more popular and widely applicable. However, despite this increased interest in this area of research, there is still a general lack of appreciation of the complexity of the choice processes undertaken by air travellers. This chapter aims to provide an in-depth look at the choice processes undertaken by air travellers, as well as discussing ways of modelling these processes. Additionally, the chapter discusses a number of issues that need to be addressed in modelling air travel behaviour, most notably the requirements in terms of data.

1 Introduction

Over the past few years, there has been an increase in the number of studies aiming to model air travel choice behaviour, mainly with the help of discrete choice models. For a comprehensive and up to date review of existing work in this area, see Pels et al. (2003) and Hess, Adler & Polak (2007).

Despite the increased interest in this area of research, a lot of work remains to be done. Indeed, the choice processes undertaken by air travellers are complex, involving decisions along a multitude of dimensions and influenced by a very high number of factors, making the modelling of such behaviour a non-trivial task. While authors are gradually acknowledging this in their work, and while the use of advanced model structures has allowed for more realism, the majority of studies still rely on stringent assumptions that unduly simplify the choice processes. This is partly to blame on the general poor quality of the data available (cf. Section 4.1), but can in many

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

cases also be attributed to the use of inappropriate model structures. Additionally, it should be said that there is still a general lack of understanding of the actual choice processes undertaken by air travellers, a fact that is not helped by the dynamic nature of the problem at hand, as witnessed for example with the advance of low cost carriers.

Rather than presenting yet another empirical analysis, this chapter looks at the issue from a more theoretical perspective, and aims to give an overview of the choice processes undertaken by air travellers, along with discussing ways of modelling these processes on the basis of available data.

The remainder of this chapter is organised as follows. Section 2 provides an overview of the choice processes undertaken by air travellers, with a particular emphasis on the interdependencies between the various choice dimensions. This is followed in Section 3 by a discussion setting targets for practical research in this area while also acknowledging certain simplifications that are necessary in such analyses. Section 4 is concerned with data issues, while Section 5 looks at the question of model structure. Finally, Section 6 provides a brief summary of the chapter.

2 Choice behaviour

2.1 General framework

Broadly speaking, outside a mode-specific context, and without aiming to define the order of choices, travellers can, for a given trip, be seen to take decisions along four main dimensions of choice:

- Destination
- Timing (time & date)
- Mode of travel
- Route choice

It can be seen that the latter two choices are strongly interrelated; in many cases, the choice of a given mode prescribes a specific route, and vice versa, a principle that applies especially in the case of trips involving intermediary stops. Trip timing is clearly influenced by outside factors such as work commitments. Additionally however, the choice set in terms of possible departure times (and in some cases even departure dates) depends on mode-specific attributes for all but self operated modes (i.e. car as driver, walking, cycling). Finally, choice set formation along the mode choice dimension is highly dependent on origin and destination, as well as on socio-demographic characteristics (e.g. driver's license). The choice of destination can clearly be seen to have a very significant impact on the other three dimensions of choice. However, even within this general framework, it can occasionally be argued that the choice of destination is not in fact made a priori, but is itself a function of other choices. As an example, a traveller who takes a decision to rely on public transport (or is forced to do so by socio-demographic characteristics) will be limited in the number of potential destinations. An even stronger example is given by the case of people refusing to travel by air, or by sea. Clearly, such factors come into play mainly in the case of leisure travel, where, depending on the circumstances, they can play a major role.

2.2 Dimensions of choice in air travel

In the case of air travel, the situation becomes significantly more complicated. Indeed, the various dimensions of choice listed above are not only strongly interrelated in the case where air travel is a possibility, but also have several subcategories.

Essentially, on top of the choice of destination and the choice of air as the main mode, the choices made by an air traveller can be divided into three main subcategories (origin side, destination side, and air side) which we will now look at in turn in Sections 2.2.1 to 2.2.3, before discussing the inter-dependencies between the various choice dimensions in Section 2.3.

The choices in this section are described for the outbound leg of a return journey, where no attention is paid to trip duration. In general, for passengers on their return leg, the majority of journey factors are predetermined by the choices made on the outbound leg, although some factors, such as timing and possibly also routing, are determined separately. Finally, for passengers on one way journeys, the choice process is very similar to the one described below for the outbound journey of return passengers, though outside factors and personal priorities may change significantly.

Here, it should be noted that in the remainder of this chapter, we work on the assumption that travellers make a conscious choice along all dimensions of choice. This is clearly a major assumption, as some travellers may consistently ignore certain dimensions of choice. Furthermore, a different rationale might apply in the case of passengers who rely on travel agents for booking their tickets. Such factors are not taken into consideration in the remainder of this discussion, but appropriate simplifications of the framework are possible.

2.2.1 Origin side

The choices on the origin side of an air journey can be divided into two main parts, the choice of a departure airport, and the choices made for the ground level journey to this departure airport.

In many cases, the choice set for the departure airport is very limited, and dominated heavily by the airport closest to the passenger's ground level origin. However, for passengers living near major urban centres, there will often be a choice between a number of airports located at similar distances from a given passenger's ground level origin. In some rare cases, passengers may even be faced with a choice between airports located in separate nearby multi airport regions, such that there is a choice of *departure city*.

Additionally, passengers take multiple decisions along the access journey dimension, which are dominated by the choice of access mode, or combinations thereof. Depending on the mode chosen, and especially in the case of a combination of multiple access modes, there is the additional choice of a route, while, for journeys involving, and especially terminating with, a car journey, there is often a choice to be made between self drive and drop off, and, in the former case, a choice between different parking options. Additionally, passengers do make a choice of departure time, which, although dependent on personal preferences, is highly influenced by the departure time of the actual air journey.

2.2.2 Destination side

In many ways, the destination side choices are the mirror image of those made at the origin side. Aside from the actual choice of ground level destination, these include the choice of destination airport, and ground level transport between this airport and the final destination.

However, there are some subtle differences. Indeed, from the point of view of a passenger on the outbound leg, there is generally an issue of a lower level of knowledge than at the origin side. This relates partly to the geographical location of the different possible destination airports (and distance to the ground level destination), but also to choice set formation along the egress journey dimension, in terms of ground level transport modes, as well as routes. Here, another point needs to be taken into account in that, for the majority of travellers, private car is not an option at the destination end, but is, for at least some of these travellers, replaced by rental car.

Aside from the above discussion about different choice set formation in the ground level dimension, the point about a lower level of information would suggest a less rational behaviour from an outside perspective[1], except for the more regular traveller. Additionally however, it should be noted that the set of priorities at the destination end may be different from those at the origin end. As such, it is possible to imagine a higher reluctance to accept a longer egress journey even in exchange for fare reductions than might be the case for the departure end, given the wish to arrive at the ground level destination more promptly.

2.2.3 Air side

Except for the questions of origin and destination, and ground level transport, the air side category contains all remaining choices describing the journey. Aside from *spontaneous* choices made at different stages of the journey (such as what to do while at the airport, or during the flight), these choices all relate to the actual travel from the origin airport to the destination airport. These in turn can be subdivided into three very much interrelated dimensions of choice.

The first choice is that of an airline operating a route to the chosen destination. In most cases, passengers will travel on a single airline for the duration of their journey. However, on some routings, for example those involving short-haul feeder flights to regional airports, or complex international routings, there is the possibility of a combination of airlines. This situation has in recent years increased in complexity, given that a large number of routes are now operated under *code share* agreements. The choice of an airline is one of the factors that makes air travel different from other areas of transport analysis. Indeed, while in many areas, passengers can be seen to make a choice between different modes of transport, the additional within-mode choice seen in the case of air travel is quite specific, and does for example not apply to the same level in areas such as rail or coach travel.

The choice of an airline or combination of airlines is strongly related to the choice of a routing. The first level along this dimension of choice divides flights into direct and connecting flights, with the possibility of a third category, for flights involving a stopover without a change of aircraft. The second level applies only to connecting flights, and involves the choice between a number of different possible routes, which includes a decision on the number of connections, and the choice of connecting airports. It is also here that the possibility of multi airline journeys arises.

[1]This does not per se suggest irrational behaviour. It simply means that, had the traveller been in possession of all information, he might have been expected to behave differently.

The final dimension of choice for the actual air journey is that of timing, i.e. the choice of a departure time and a departure date. For some passengers, the most important factor will be the departure time, while for others, it will be the arrival time. In practice, this equates to the choice of a specific flight.

There are ways to consider further subdivisions of choice dimensions along the air side category. One example includes the choice of aircraft type. Indeed, some passengers have an inherent dislike for turboprop planes, and use this as a determining factor in their choice process. However, such factors can in fact be seen as an attribute of a specific flight, which is thus accounted for by the other dimensions of choice listed above, and as such, can be included in models as a simple explanatory variable.

2.3 Choice process in air travel

The above discussion has illustrated that the process of putting together a trip from a ground level origin to a ground level destination, with an intermediary air journey, is a complicated one, involving decisions along a multitude of dimensions. What makes the analysis of these choice processes more complicated is the fact that there exists a highly complex structure of interdependencies between the various dimensions of choice, and that the order of priorities amongst dimensions of choice is highly likely to vary across individuals as well as across situations. We will now look at these interdependencies, starting from a very basic situation, and slowly moving up to more detailed interactions.

2.3.1 Base scenario

In the most basic scenario shown in Figure 1, only the most obvious of dependencies are shown for the actual air journey[2]. As such, the choice of access mode(s) and the choice of route clearly depend on the choice of departure airport, with a similar conclusion for the egress journey and the choice of arrival airport. This accounts for the fact that not all airports will be accessible by all ground level modes or combinations thereof. Interactions with individual-specific characteristics, such as car ownership, are not taken into account in this discussion; they are to be treated at an individual-

[2]A link from choice dimension $\mathcal{A}$ to choice dimension $\mathcal{B}$ indicates that the decision taken in choice dimension $\mathcal{A}$ can have an impact on choice dimension $\mathcal{B}$ in two ways; it can lead to changes in the attributes of the alternatives along this dimension, but it can also lead to a change in the composition of the choice set along this dimension.

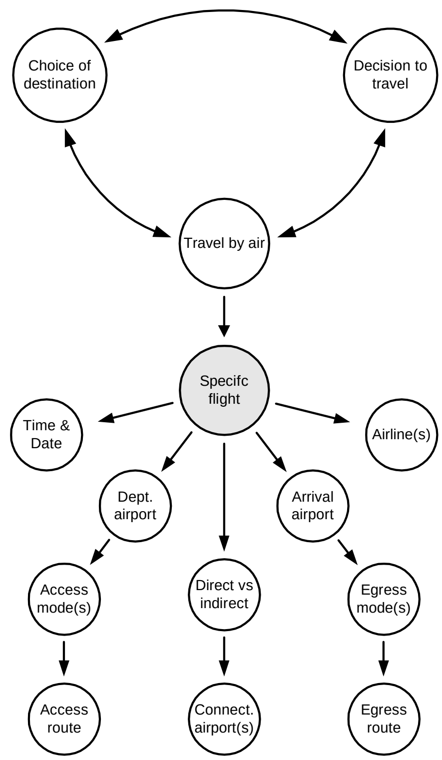

Figure 1: Main choice processes of an outbound air journey

specific level in the models. The other dependency shown for the actual air journey in Figure 1 relates the choice of connecting airports to the question of whether a direct or a connecting flight is chosen. Additionally, it should be noted that the impact of the *choice of destination* (ground level) and the *decision to travel* on flight-specific attributes (e.g. airline, time and date) is allowed for by the link from these nodes through the *travel-by-air* node.

Figure 1 shows an additional set of interactions outside the main air journey, relating the *decision to travel* to the *choice of destination* and the decision to *travel by air*. Each one of these links is shown as being reciprocal, a fact that needs some further explanation.

The *decision to travel* and the *choice of destination* are inherently linked, and the order of decisions depends on the situation at hand. The case where the decision to travel precedes the choice of destination applies mainly in the case of leisure travel, where it needs no further explanation (e.g. holidays). However, it can also be seen to apply in the case of business travel, as for example in the situation where a decision is taken to hold a meeting before deciding on an actual location. Conversely, the case where the choice of destination precedes the decision to travel applies more directly in the case of business travellers, where the a priori determined geographical location of a meeting or conference will have an impact on the decision to attend and hence to travel. For leisure travellers, it is similarly possible to imagine cases where a respondent may make the decision to travel to a certain leisure event (e.g. joining friends on holiday) dependent on the location of the event.

The link between the *choice of destination* and the *decision to travel by air* also acts in both directions. Indeed, the fact whether air travel is a viable alternative obviously depends on the chosen destination, where there are clearly also cases where the choice of a destination effectively eliminates all other modes from consideration. Conversely, a respondent who takes an a priori decision to travel on a specific mode clearly imposes limitations on the set of possible destinations, where the same reasoning applies for travellers who decide to exclude a specific mode from consideration.

The link between the *decision to travel* and the *decision to travel by air* is less straightforward to explain. Nevertheless, it can be seen that a traveller deciding to travel may take the a priori decision to travel by air before choosing a destination, something that applies principally in the case of leisure travel. On the other hand, the link between the decision to fly and the decision to travel in the first place can be explained most readily in conjunction with the choice of destination. As such, if a destination is imposed by outside factors (e.g. meeting, conference), and if this implies a need to fly, some passengers may decide not to travel.

2.3.2 Detailed interactions between choice dimensions

We now turn our attention to the detailed interactions between the various choice dimensions. For this, the choice dimensions are again split into two parts, namely the upper and lower half shown in Figure 1, representing the *upper level* choices and the choices relating to the actual air journey respectively. We will first look at the interdependencies of the various choice dimensions describing the air journey, before turning our attention to the links between these choices and the *upper level* choices.

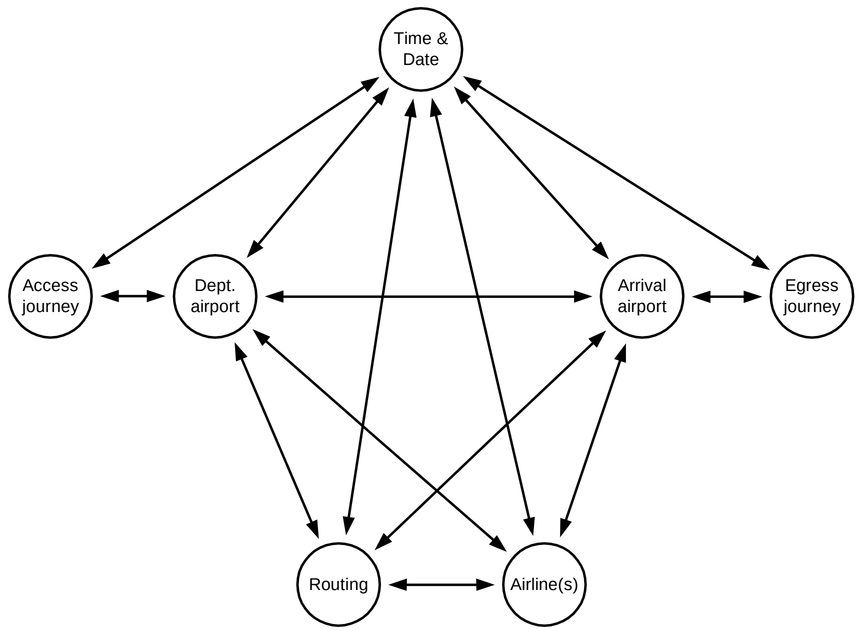

Figure 2: Interactions between air travel choice dimensions

Interactions between air travel choice dimensions

The interactions between the various choice dimensions relating to the actual journey between a ground level origin and destination are shown in Figure 2. Some minor simplifications are used here in that the nodes relating to the access journey (mode and route) in Figure 1 are now replaced by a single *access journey* node, with a corresponding approach for the egress journey. Additionally, the separate nodes for *direct vs indirect* and *connecting airport(s)* are now replaced by a common *routing* node. Finally, the *specific flight* node has been eliminated, to show the direct relationships between individual dimensions without passing through this node.

From Figure 2, it can be seen that, aside from the two ground level journey nodes, all decision processes are connected to each other, in both directions. This shows the high level of inter-dependency between the various choices. We will first turn our attention to the inter-dependencies between the five central nodes, before returning to the access/egress journey nodes to justify their special treatment.

It can easily be seen that the choice of a specific departure airport will reduce or affect the set of options open to the traveller along at least some of the other dimensions of choice, such that it possibly has an impact on the choice of arrival airport, departure time (and to a lesser extent date), airline, and flight routing. The same reasoning applies in the case where a traveller makes his choice dependent on a specific arrival airport. Similarly, it can be seen that a respondent who chooses to travel at a given time and/or on a given date affects his options along the other dimensions of choice, with analogous reasonings in the case of a choice of a specific airline, or routing. As such, it is clear that the links between these five main nodes are justifiably set to act in both directions.

A special treatment applies in the case of the access journey and egress journey nodes. They are linked directly (and reciprocally) only to the respective airport node, and to the *time and date* node. It is clear that the decision to travel at a given time or on a given date will, in some cases, have an impact on the choices made for the access journey. As an example, a very early departure may in some cases remove the possibility to travel by public transport. This link clearly also acts in the opposite direction, in that travellers who are determined (or forced) to use public transport may only be able to travel at certain times or even certain dates. It can also be seen that the choice of a specific departure or arrival airport has an impact on the choice made along the access journey respectively egress journey dimensions. Here, the reciprocal link might be less obvious, but there are clearly cases in which the choices made along the ground level dimensions can take precedence over those choices made in terms of departure, respectively arrival airport (and hence over the three other central dimensions by extension). As such, for passengers who decide to travel on a given mode or combination of modes, the choice set in terms of departure airports may be reduced. While these links between the access and egress journeys and the respective airports as well as trip timing are thus well-defined, it cannot easily be argued that there is a *direct* link (in either direction) between the choice of a specific airline or routing and the choices made for the access or egress journey.

Interactions with upper level choice dimensions

We now turn our attention to the interactions between the *upper level* choices and the choices relating to the actual air journey. We will first concentrate on the *downwards* interactions, before looking at the potential

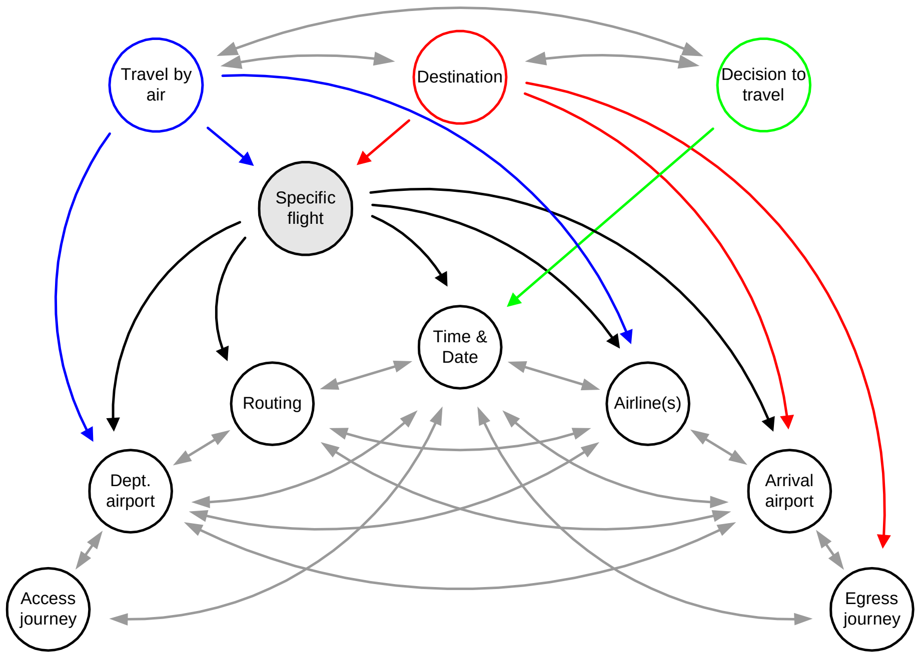

Figure 3: *Downwards* interactions between *upper level* choices and air travel choice dimensions

impact of *lower level* choices on *upper level* decisions.

The downwards interactions are shown in Figure 3, which additionally reproduces the level-specific interactions, as previously shown in Figure 1 and Figure 2. The only justifiable direct impact of the *decision to travel* is on the *time and date* of the journey. Any remaining impacts are indirect, either through the *destination* node, or the *travel by air* node. For the impact of these two nodes, an intermediary node, relating to the choice of a *specific flight*, is introduced again. Through this, the choice of destination and the decision to travel by air have an impact on all the main characteristics of the air journey, and by passing through the appropriate airport nodes, also the access and egress journey. Additionally, the two decisions however potentially also have some direct impacts on lower level choices, which are taken prior to selecting a specific flight, or at least carry special weight in the selection of a specific flight. Such interactions are shown as direct links in Figure 3. For the choice of ground level destination, two such direct links are shown. Indeed, it is clear that the precise geographical location of the chosen ground level destination has a strong direct impact on the choice of arrival airport in the case of a destination served by more than one airport. Similarly, it also potentially has a direct impact on the choice of egress mode and route at the destination end. No direct link is included between the choice of destination and the time and date of travel; this process can pass either through the *decision to travel* node, or the choice of a specific flight. For the decision to travel by air, two main direct impacts are identified, namely to a specific departure airport, and a specific airline, reflecting airport and airline allegiance. All other links are represented indirectly, through the *specific flight* node. It could be argued that the decision to travel by air is in some cases coupled directly to a decision to choose only direct flights, however, this may not always be possible, and it is not clear whether there is a significant share of travellers who would in this case decide not to travel by air. In any case, the indirect link still allows for such effects.

Even though the choices shown in the upper level of Figure 1 are in general seen as taking precedence over those shown in the lower level, this is not always the case. The various possibilities of an *upwards* effect are shown in Figure 4, where it must be said that some of these links apply primarily to leisure travellers, as mentioned at appropriate stages in the following discussion.

The most obvious of the *upwards* effects are the three links associated with the *time and date* node; clearly the a priori choice of a date (and to a lesser extent time) for travelling has an effect on the choice of destination. Similarly, if a date or time for a specific event is imposed on the decision maker through a third party, than this does potentially have an effect on the decision to travel as well as the decision to travel by air[3]. Another obvious effect that needs to be represented is the relationship between the available flight options, in terms of airline, airports, route, and fare, on the decision to travel by air, as well as the decision to travel in the first place. This is represented in Figure 4 by appropriate links from the departure airport, routing and airline nodes to the two travel decision nodes, where the combination of these nodes reflect specific air travel options. The additional dimension of flight fares is not represented here; although there is clearly

[3]The latter applies principally for the case of short-haul trips, where the decision to travel by air as opposed to ground level modes is very much dependent on scheduling flexibility, but also on the advance notification.

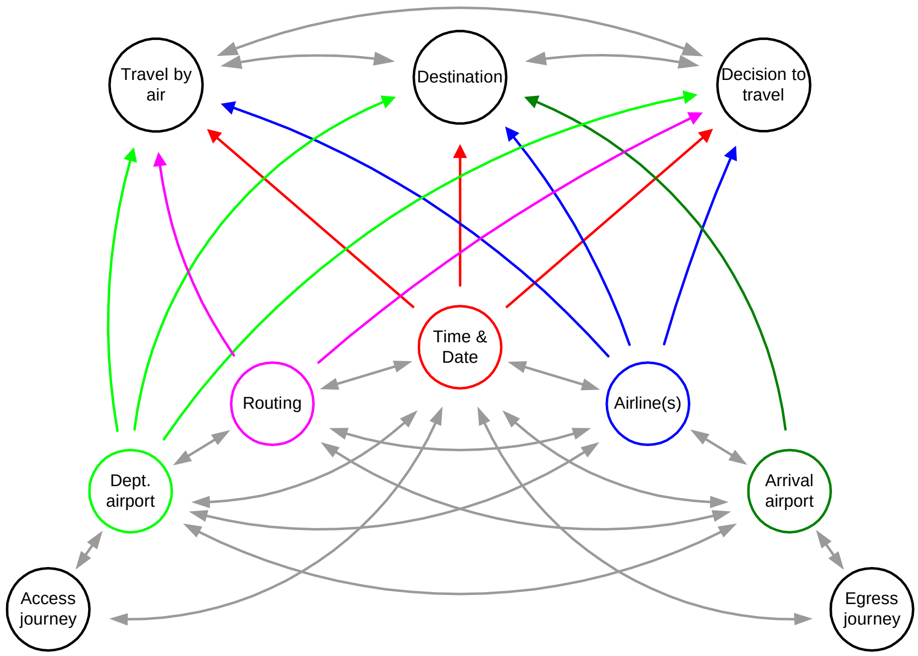

Figure 4: *Upwards* interactions between air travel choice dimensions and *upper level* choices

a choice in terms of class of travel, it can be argued that passengers do not generally make an additional conscious choice of fare-type for a specific flight. The actual effects of differences in fares across different flights are implicit in the links from the various nodes listed above (especially the airline and routing ones).

In addition, some of the links already discussed above, as well as the remaining links shown in Figure 4, can be explained in a different way, which can broadly be seen as a type of airport or airline allegiance, which however applies most directly in the case of leisure travellers. As such, it can be argued that some travellers make an a priori choice of departure airport, before organising the remainder of their journey. This is most easily understood in the context of regional airports. In such cases, the choice of destination clearly depends on the destination on offer from this airport, and if no attractive destinations (or indeed flights) are available, the traveller may decide to travel on another mode, or not at all. A similar phenomenon can be observed in the case of strong allegiance to a given airline, where passengers make the choice of destination, as well as the decision to travel, conditional on what is on offer from this specific airline. While this may in some cases occur as an effect of airline-allegiance in the classical meaning (i.e. frequent flier account), it can be seen to apply even more clearly in the case of low cost airlines; these operators can be seen to induce new demand, such that some of their passengers would not travel at all (or at least not by air), if it wasn't for the presence of the specific airline.

The above discussion suggests that the product offered by the chosen airport, or the chosen airline, cannot only determine the potential destinations, but, if no destination appeals to the passenger, also the decision to travel in the first place. Although these factors are becoming increasingly important with the increased popularity of low cost carriers, they can be seen to also apply in the case of flights offered in conjunction with package holidays, where passengers may be rather flexible in terms of their precise destination. However, as already alluded to earlier on, the principle also applies in the case of business travellers, where the possibilities in terms of departure airport and airline can play a deciding role in the choice for a meeting place.

The final link shown in Figure 4, namely that from the choice of arrival airport to the choice of ground level destination again applies mainly to leisure travellers, and specifically the case where people book a flight before sorting out the detailed ground level destination within the wider surrounding area of the destination airport.

2.4 Discussion

The above description of the choice processes made by air travellers has shown that such journeys not only involve decisions along a multitude of choice dimensions, but that there exist complex inter-dependencies between these various choice dimensions. Given that a large number of the links have been shown to act in both dimensions, it is clearly inappropriate to attempt to model the decision making as a sequential choice process, but rather, that simultaneous analysis is required in the absence of information on the relative level of priorities across travellers[4].

The discussion has also shown that, while the top *upper level* decisions

[4]The use of such information is highly hypothetical in any case, as it is not clear whether there are situations in which it is possible to define a clear sequential choice process involving decisions along all of the above listed dimensions.

have major direct and indirect effects on all *lower level* decisions, these links potentially also act in the opposite direction. This leads us to an important observation. Indeed, the discussion so far, along with the point of view taken in the various diagrams, would suggest that the question of main mode of travel is taken at a separate level. As such, travellers would gauge the overall product offered by the various modes, and then make a choice of main mode before moving on to mode-specific choices (e.g. airline, route, ...). Clearly, for some journeys, a mode choice decision is taken a priori, such that the above discussion holds. This applies most notably in the case of destinations or circumstances where only a single mode is a viable alternative. As such, in the case of aviation, it can be seen that beyond a certain distance, the mode choice question becomes obsolete. The problems arise in the case of short-haul to medium haul distances, where the competition from car and ground level public transport needs to be taken into account, especially with the increase in high speed rail routes. In such cases, it is not necessarily clear whether passenger make an a priori choice between air and rail, before, if applicable, making within-mode decisions in the case of air travel. Rather, it can be imagined that, for at least part of the travelling population, the high speed rail alternative appears on the same level as the various air travel alternatives. As such, the alternatives are evaluated in parallel, with all intra-mode considerations taken into account at the same time as the cross-mode comparison[5].

Although such a parallel analysis does not pose any major problems from a methodological point of view in terms of model structure, it does come at the cost of increased data requirements. As such, it is now necessary to obtain detailed data for ground level modes, in addition to the air travel data, the procurement of which already causes problems on its own, as described in section 4.1. Additionally, it should be noted that, for some routes, the number of possible ground level options (in terms of combinations and routes) is so high that the data requirements can become insurmountable. As such, it should come as no surprise that the majority of studies make an assumption of an a priori decision to travel by air. While this does not per se invalidate the results of the analyses in question, it is important to acknowledge the potential shortcomings, at least in the presence of routes where there is potential high inter-mode competition which is not characterised by

[5]It should be noted that the interactions between the choice dimensions shown in the diagrams do allow for such comparisons, although not described explicitly. Indeed, the fact that the mode-specific characteristics are connected upwards to the decision to *travel by air* does allow for the rejection of an air alternative in favour of an alternative on another mode.

an a priori choice of mode before proceeding to intra-mode decisions.

3 Modelling practice: goals and limitations

This section provides some guidance on several critical points that analysts should pay attention to in modelling air travel choice behaviour, while also looking at some of the simplifications that are required to enable large-scale analysis.

3.1 Goals

3.1.1 Recognise the multi dimensional nature of the choice process

It should be clear from the discussion in Section 2.2 that air travellers take decisions along a multitude of choice dimensions. Not recognising this in practical work can potentially lead to biased results. This is most easily illustrated on the basis of an example, say with a study that does not account for the fact that passengers choose an airline in addition to a departure airport. It is clear that someone with a strong allegiance to airline $\mathcal{X}$ may not be interested in other airlines. As such, in the case where airline $\mathcal{X}$ operates only from an outlying airport with low overall frequency (across all airlines), this respondent can be seen to choose an airport with lower frequency to the desired destination. From this point of view, it is important to work on the basis of frequencies specific to airport-airline pairings, and not airport-specific frequency, as has commonly been the case in existing work. A similar reasoning applies in the case of the other dimensions of choice.

A straightforward approach for the representation of such multidimensional choice processes can in general terms be described as follows. A *combined* alternative chosen by a given traveller is composed of D *elementary* alternatives, one along each dimension of choice. Some of the attributes of the *combined* alternative are specific to a single one of the choice dimensions, while others are specific to a given combination, say for example an airport-airline pairing. The availability of a *combined* alternative is a function of the availabilities of the *elementary* alternatives. Each *combined* alternative counts as a separate alternative, such that the choice process is indeed represented in a simultaneous rather than sequential fashion.

3.1.2 Account for correlation along different choice dimensions

By understanding the multidimensional nature of the choice process, it becomes evident that some of the *combined* alternatives share the attributes of other alternatives along one or more of the choice dimensions. In the context of the combined choice of airport, airline, and access mode, let K, L and M define the number of airports, airlines and access modes respectively, and let us assume that all combinations of airports, airlines and access modes are possible, leading to a total of KLM *combined* alternatives. It can then be seen that a given alternative shares the same airport (and hence the related attributes) with $LM-1$ alternatives, the same airline with $KM-1$ alternatives, and the same access mode with $KL-1$ alternatives. Furthermore, an alternative shares the same airport and airline with $M-1$ alternatives, the same airport and access mode with $L-1$ alternatives, and the same airline and access mode with $K-1$ alternatives. While some of these commonalities can be accounted for through the attributes included in the observed part of utility, it is almost inevitable that there is also some correlation between unobserved utility terms along these dimensions. Here, it is important to account for this correlation, as described in Section 5.2.

3.1.3 Use highly disaggregate level-of-service data

Another major problem with existing studies has been the use of an insufficient level of disaggregation in the level-of-service data. This is highly correlated with the decision in a lot of previous work to use simplifications along a number of choice dimensions. While some aggregation is almost inevitable, for example by working on the basis of average fares across flights on a given day, the use of an excessive amount of aggregation can lead to biased results. This can for example arise in the case of studies making use of weekly (or even monthly) data instead of daily data, hence ignoring the often significant variations in level-of-service attributes between different days of the week, especially in terms of flight frequencies.

3.1.4 Account for differences in behaviour across travellers

One of the most important points to take into account in modelling air travel choice behaviour are the variations in behaviour across respondents. No two respondents are exactly the same, in terms of sensitivities as well as priorities. Not taking such variations in tastes (and hence behaviour) into account will inevitably lead to lower quality results and may even lead to significant bias. This issue is address in more detail in Section 5.3.

3.2 Limitations of practical research

In practical research, it is not generally possible to jointly model all the choice dimensions described in Section 2.2, and some simplifications are required. There are some subtle differences here between RP and SP data. As such, in the case of RP data, respondents did indeed perform a multi-dimensional choice process[6], but the quality of the data often prevents an analysis along all the choice dimensions. On the other hand, in the case of SP data, information is available on all choice dimensions, but respondents are generally only faced with choices along a rather small subset of the dimensions listed in Section 2.2. As such, in both cases, only a subset of the choices will be modelled. We will now briefly look at the various dimensions of choice and discuss what simplifications are generally used.

- The ***choice of destination*** and the actual ***decision to travel*** are not generally modelled, where a reservation applies for the latter in the case of SP surveys giving respondents the option *not to travel*. Research thus relies on the assumption that these decisions are taken at an upper level, prior to putting together the exact details of the journey.

- The ***decision to travel by air*** is also generally treated as having been taken at an upper level. Although this can cause problems in the case of short-haul destinations, this simplification is almost unavoidable, given the heightened data requirements as well as added modelling complexity. As such, it is important to acknowledge that the estimates obtained from such models relate to the part of the population that has decided to travel by air, and are not representative of the overall population.

- The majority of studies of air travel choice behaviour look solely at the choice of ***departure airport*** and ignore the choice of ***arrival airport***, where the latter choice can however be modelled on the basis of non-resident travellers on the return leg of their journey (excluding the possibility of open jaw tickets). Again, this simplification is primarily due to data issues, but also the relatively low number of routes connecting two multi airport regions.

- In studies that model ground level journey related choices in addition to air travel choices, it is generally only possible to model the ***choice of***

[6] The possibility exists for some individuals not to make a conscious choice along some of the dimensions.

access mode to the departure airport on the current leg; an analysis of the ***egress journey*** related choices would lead to a requirement for detailed level-of-service data for the ground level transport network at the destination end. Furthermore, most studies are only able to look at the choice of main mode, ignoring the possibility of trip-chaining, as well as the choice of different routes. The effects of these restrictions are dependent on the geographical context.

- In RP models, ***trip timing*** cannot generally be modelled due to a relative lack of information on preferred departure time and flight availabilities. These issues do not apply in the case of SP models.

- ***Flight routing*** is generally left untreated in RP studies, which often look only at the choice between direct flights; this can have major effects in the case of destinations with a large share of connecting passengers. As with most other factors, SP studies are in theory able to model the choice of routing, where this is again dependent on the design of the SP surveys.

- Given the great role that airline allegiance plays for some travellers, advanced studies increasingly model the ***choice of airline*** in addition to the choice of airport. However, combinations of airlines are generally not allowed in such models, and additional issues can arise in the case of code share flights. Again, the main issue is one of data quality.

4 Data issues

As the above discussion has shown, the analysis of air travel behaviour is a complicated undertaking, and a number of issues need to be dealt with it in the process. While the majority of these issues have been addressed in some form or other in the above discussion, two main remaining issues need to be addressed in more detail. They both relate to the data used in estimation, and are concerned with overall data quality, and the selection of destinations to be used in the analysis. These two issues are now looked at in turn.

4.1 Data quality

Almost certainly the single biggest issue that needs to be faced in the analysis of air travel behaviour is that of the quality of the available data. The issue can be divided into several sub-issues as follows.

4.1.1 Availability and attributes of unchosen alternatives

The first major data problem that needs to be faced is the relative lack of information on the choices that travellers are faced with. Indeed, in an RP context, adequate information will generally be available for the majority of choice affecting attributes for the alternative that was actually chosen. This does however not extend to the unchosen alternatives. While it is generally possible to obtain some information on such alternatives in terms of substantive attributes, the lack of data on availability (of flights as well as of specific fare classes on a given flight significantly complicates the characterisation of unchosen alternatives. A similar issue applies along the access mode dimension, where information on the availability and attributes of unchosen modes is often not available.

4.1.2 Fare data

Another very significant problem, which, although strongly related to the issues of unobserved attributes discussed above, deserves special attention is that of the fare data. Air fares, and especially relative air fares, can be expected to play a major role in air travel choice behaviour. However, in the majority of studies of air travel choice behaviour, it has not been possible to recover a meaningful marginal utility of fare changes. This can almost certainly be explained on the basis of the poor quality of the fare data. Indeed, while, for many attributes of the unchosen alternatives, it may be possible to produce a fairly good approximation to the true value, this does not apply in the case of fare data. In most cases, only the average fare charged by a given airline on a specific route will be available. This clearly involves a great deal of aggregation, as no distinction is made between the fares paid across different travellers (i.e. in terms of travel classes as well as booking classes).

In addition to the obvious loss of information, the aggregation of fares leads to another important problem, given that there is generally also no information on the availability of specific fare classes at the time of booking, as opposed to the mere availability of actual flights. Indeed, this essentially leads to the need for an assumption of equal ticket selling speeds across all flights (routes as well as departure times), which is clearly not necessarily the case. The increasingly dynamic nature of air fares makes the use of aggregate fares even less reliable.

Even though some progress can be made with the help of bookings data, issues of aggregation do remain. In fact, it can be seen that, in RP studies, disaggregate choice data is used in conjunction with aggregate level-of-service data, for at least some of the attributes. While, for some characteristics, this may be acceptable, it does, as described above, create significant problems in the treatment of air fares, and flight availability by extension.

4.1.3 Frequent flier information

It is well known that passengers are heavily influenced in their choice of airline by their membership in frequent flier programmes, either on a personal basis, or as part of a company-wide scheme. Unfortunately, information on frequent flier memberships is generally not collected in passenger surveys. As such, this potentially crucial influence on choice behaviour cannot usually be taken into account in RP case studies (as opposed to SP studies). In the case of datasets including a large number of international flights, operated by a variety of airlines, there is however an alternative way of modelling travellers' loyalty behaviour, by analysing their allegiance to their national carrier.

4.1.4 Influence of other attributes

Aside from the most commonly used attributes such as access time, flight time, frequency and fare, a host of other factors potentially play a role in travellers' choices. The problem is that, much like in the case of frequent flier programmes, the procurement of data on these attributes is often difficult. A possible exception is that of on time performance, which can be expected to play a role especially for experienced travellers. However, again, the data is often not available in a sufficiently disaggregate form. Finally, a number of other attributes that potentially have an influence on choice behaviour are of a highly qualitative nature, making their inclusion in models difficult from a data as well as methodological point of view. Examples include factors such as comfort, quality of the food, and quality of the in-flight entertainment.

4.1.5 Survey design issues and inter-dataset compatibility

An additional problem relates more directly to the nature of the available data. The main input into an analysis of revealed air travel choice behaviour comes in the form of passenger survey data collected at the departure airport. The fact that in many cases, different people (or even institutions) are responsible for the design of the survey and the subsequent modelling analysis creates obvious problems. Furthermore, the survey data need to be complemented by level-of-service data, which can lead to major problems of inter-dataset compatibility.

4.1.6 Stated Preference data

The above discussion has highlighted the issues with using RP survey data in analysing air travel choice behaviour. This in turn could suggest that the way forward would be the use of SP data collected on the basis of tailor-made surveys, designed specifically for the use in advanced modelling analyses. The main advantage of SP surveys in the present context is the fact that the modeller knows precisely what information the respondent was faced with when making his choice, especially with regards to the availability and attributes of unchosen alternatives, something that is not generally the case in RP data.

Given these theoretical advantages of SP data in the present context, it should come as no surprise that studies making use of SP data have been much more successful in retrieving significant effects for factors such as air fares and airline allegiance (cf. Adler et al. 2005, Hess, Adler & Polak 2007). Despite the apparent advantages of SP data over RP data in this context, it should not be forgotten that the use of SP data does pose some philosophical problems, in terms of how the behaviour differs from that observed in RP data. Additionally, important issues need to be faced at the survey design stage, notably in terms of questionnaire complexity.

In closing, it can be said that both approaches have advantages and disadvantages, where the relative merit of each approach needs to be measured on a case by case basis, and depends heavily on the quality of the level-of-service data in the RP context. An interesting approach in this context is to combine RP and SP data, as done by Algers & Beser (2001), hence correcting for the bias inherent to models estimated on SP data. The problem in this case however is one of obtaining compatible RP and SP datasets.

4.2 Choice of destinations

A major issue falling within the wider field of data quality is the question of what destinations to include in a study. This is clearly heavily influenced by the choice date used, and the destinations represented therein[7]. As such,

[7]We ignore the case where the modeller himself is responsible for the collection of the choice data, and hence in a position to define quotas so as to obtain an adequately sized sample for an a priori defined set of destinations.

only destinations with a sufficient number of observations in the choice data can be included in the analysis.

The next point that needs to be taken into account is that only destinations that can be reached from at least two airports in the study area should be included, for obvious reasons. Furthermore, it is of interest to avoid including destinations that are served from more than one airport, but where the frequency at all but one of the airports is negligible. Additionally, if a decision is taken to include only destinations reachable by direct flight, then this clearly reduces the number of eligible destinations further.

The various requirements listed above can in some cases lead to an insufficient number of destinations/observations, depending on the choice data at hand, such that the requirements may have to be relaxed, to guarantee an acceptable sample size. In the case of the analysis of a multidimensional choice process, this may for example be achieved by the inclusion of some destinations that are only reachable from one airport, but where a choice still exists along the other dimensions.

One main complication however still needs to be addressed, namely the treatment of destinations that are themselves located in multi airport regions. Indeed, for such destinations, passengers not only make a choice of departure airport, but must also be expected to make a choice of destination airport, and it is not clear whether the choice in the study area takes precedence. This applies especially in the case of passengers who are on the return leg of their flight.

To conduct a proper analysis, the choice of arrival airport would in this case normally have to be included explicitly in the modelling framework; this however not only leads to a very significant increase in the complexity of the analysis, but also greatly increases data requirements, notably in terms of ground level transport data for the different destination airports. Given these complications, destinations located in multi airport regions should thus ideally be excluded from the analysis if no treatment of the choice of destination airport is to be included in the modelling framework. However, this is not always possible, given the often high representation of such destinations in air passenger surveys. In this case, it is however important to maximise the probability of there being a conscious choice of airport in the study area, for residents as well as visitors. Furthermore, modellers should at least acknowledge the fact that their results may be influenced by the presence of multi airport destination regions.

5 Model structure

One of the most important questions arising in the analysis of air travel choice behaviour is that of what modelling approach to use. Given the important differences across travellers both in terms of behaviour as well as choice context, the use of a disaggregate modelling approach is clearly preferable to an aggregate one. In the area of transport analysis, discrete choice structures belonging to the class of random utility models (RUM) have established themselves as the preferred approach for such studies over the past thirty years. An in-depth discussion of these modelling approaches is beyond the scope of the present chapter, and the readers are referred to the excellent overview provided by Train (2003). In this chapter, we merely discuss the use of discrete models to represent correlation (Section 5.2) and taste heterogeneity (Section 5.3) in the context of air travel behaviour research. This is preceded by a brief introduction of some common notation.

5.1 Basic concepts

In a discrete choice experiment, a *decision-maker* n chooses a single alternative from a choice set C_n, made up of a finite number of mutually exclusive alternatives, where the choice set is exhaustive, and the ordering of alternatives has no effect on the choice process undertaken by the decision-maker. Each alternative $i = 1, \ldots, I$ in the choice set is characterised by a utility $U_{i,n}$, which is specific to decision-maker n, due to variations in attributes of the individuals, as well as in the attributes of the alternative, as faced by different decision-makers. The use of the concept of utility, along with the need for a decision-rule, leads to the single most important assumption in the field of discrete choice modelling, namely that of *utility maximising behaviour* by respondents. As such, respondent n will choose alternative i if and only if $U_{i,n} > U_{j,n} \;\; \forall j \neq i$, with $i, j \in C_n$.

In an actual modelling analysis, the aim is to express the utility of an alternative as a function of the attributes of the alternative and the tastes and socio-demographic attributes of the decision-maker. Here, the limitations in terms of data and the randomness involved in choice-behaviour mean that, in practice, modellers will only be able to observe part of the utility. As such, we have:

$$U_{i,n} = V_{i,n} + \epsilon_{i,n}, \tag{1}$$

with $V_{i,n}$ and $\epsilon_{i,n}$ giving the *observed* and *unobserved* parts of utility respectively. Here, $V_{i,n}$ is defined as $f(\beta_n, x_{i,n})$, where $x_{i,n}$ represents a vector of measurable (to the researcher) attributes of alternative i as faced by decision-maker n[8], and β_n is a vector of parameters representing the tastes of decision-maker n, which is to be estimated from the data. The function $f(\beta_n, x_{i,n})$ is free from any a priori assumptions, allowing for linear as well as non-linear parameterisations of utility. The inclusion of the unobserved utility term, $\epsilon_{i,n}$, means that the deterministic choice process now becomes probabilistic, leading to a random utility model (RUM), with the alternative with the highest observed utility having the highest probability of being chosen.

It can be seen that the probability of decision-maker n choosing alternative i is now given by:

$$P_n(i) = P\left(\epsilon_{j,n} - \epsilon_{i,n} < V_{i,n} - V_{j,n} \;\; \forall j \neq i\right). \tag{2}$$

With the *unobserved* part of utility varying randomly across respondents, the mean of this term can be added to the observed part of utility, in the form of an alternative-specific constant (ASC). The vector $\epsilon_n = \{\epsilon_{1,n}, \ldots, \epsilon_{I,n}\}$ is now defined to be a random vector with joint density $f(\epsilon_n)$, zero mean and covariance matrix Σ, and by noting that the probability of alternative i in equation (2) is the cumulative distribution of the random term $\epsilon_{j,n} - \epsilon_{i,n}$, we can write:

$$P_n(i) = \int_{\epsilon_n} I\left(\epsilon_{j,n} - \epsilon_{i,n} < V_{i,n} - V_{j,n} \;\; \forall j \neq i\right) f\left(\epsilon_n\right) \, \mathrm{d}\epsilon_n, \tag{3}$$

where $I(\cdot)$ is the indicator function which equals 1 if the term inside brackets is true and 0 otherwise. The probability is now given by a multidimensional integral which only takes a closed form for certain choices of distribution for ϵ_n, where the choice of $f(\epsilon_n)$ has a crucial impact on the behaviour of the choice-model. The most basic model structure, the Multinomial Logit (MNL) model, assumes that the error terms are distributed identically and independently, while more advanced model structures allow for correlation between error terms across alternatives and variation in the error terms across individuals, leading to a representation of variable inter-alternative substitution patterns and random inter-agent taste heterogeneity respectively. These two concepts are discussed in Section 5.2 and Section 5.3 respectively.

[8] The vector $x_{i,n}$ potentially also includes interactions with socio-demographic attributes of respondent n.

5.2 Correlation between alternatives

As mentioned in Section 3.1.2, it is clearly a major and probably unwarranted assumption to rule out the presence of heightened correlation in the unobserved utility terms along any of the choice dimensions. In order for such an assumption to be valid, any commonalities between two alternatives sharing a common component along one or more of the choice dimensions would need to be explained in the observed part of utility. This is clearly not possible in general, mainly for data reasons. The likely resulting correlation in the unobserved part of utility makes the use of the MNL model almost surely inappropriate, especially in forecasting.

Three quite different treatments of the correlation in the unobserved part of utility are possible in this context; the use of GEV structures, the use of an error components Logit approach, and the use of Probit structure. While potentially having advantages in terms of flexibility, the latter two are not applicable in the present context, simply on the basis of the added cost of estimation due to simulation requirements. As such, the discussion in this section looks solely at the case of GEV structures. Here, it should be noted that this discussion looks only at the correlation along dimensions, and not at the correlation across dimensions, such as the correlation between different airlines. This remains an important avenue for further research. Finally, the discussion in this section centres on the case of the combined choice of an airport, airline and access mode. Extensions to higher-dimensional choice processes are straightforward.

The Generalised Extreme Value (GEV) family of models, introduced by McFadden (1978), is a set of closed form discrete choice models that are all based on the use of the extreme-value distribution, and which allow for various levels of correlation among the unobserved part of utility across alternatives. This is done through dividing the choice set into nests of alternatives, with increased correlation, and thus higher cross-elasticities, between alternatives sharing a nest. As such, alternatives sharing a nest are more likely substitutes for each other. The use of such a nesting structure means that GEV models are most easily understood in the form of trees, with the root at the top, elementary alternatives at the bottom, and composite alternatives, or nests, in between.

The MNL model is the most basic member of this GEV family, using a single nest of alternatives, resulting in equal cross-elasticities across all alternatives. While in the MNL model, the error terms are distributed *iid* extreme value, in the general GEV formulation, the error terms follow a joint generalised extreme value distribution; the individual error terms follow a univariate extreme value distribution, but the error terms associated with alternatives sharing a nest are correlated with each other. This structure leads us away from the diagonal variance-covariance matrix of the MNL model.

The most basic GEV approach that can be used in the analysis of air travel choice behaviour is a simple two level Nested Logit (NL) model. In this model, alternatives are grouped into mutually exclusive nests, where, for each nest, a structural (nesting) parameter is estimated that relates to the level of correlation in the error terms of alternatives sharing that nest. In the context of the present discussion, three main possibilities arise in this case, nesting together alternatives either by airport, or by airline, or by access mode. The resulting structures allow for correlation along a single dimension of choice, accounting for the presence of unobserved attributes that are specific to a given airport, or, in the two other structures, a given airline respectively a given access mode. In each of the three structures, separate nesting parameters are used with individual nests, to allow for differential levels of correlation across nests.

As an example, the appropriate structure for the NL model using nesting by airport is shown in Figure 5, with K mutually exclusive nests, one for each airport, and where each nest has its own nesting parameter, λ_k. Only a subset of the composite nests and of the triplets of 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 5.

The NL structures described thus far have only allowed for the treatment of inter-alternative correlation along a single dimension of choice. This is clearly a major restriction, and inappropriate in the case where correlation exists along more than one dimension. As such, the two level NL structures should at best be seen as a tool for testing for the presence of correlation along individual dimensions, but, in the case where different structures indicate correlation along different dimensions, the two level structure becomes inadequate for use in forecasting, and potentially even for the calculation of willingness-to-pay indicators.

The NL model can be adapted to allow for correlation along more than one dimension, by using a multi level structure. A common example in the case of air travel is to nest the choice of airline within the choice of airport. It is important to stress that this should not be seen as representing a sequential choice process. Rather, it means that there is correlation between two alternatives that share the same airport, but that the correlation is larger if they additionally share the same airline. The structure of such a model

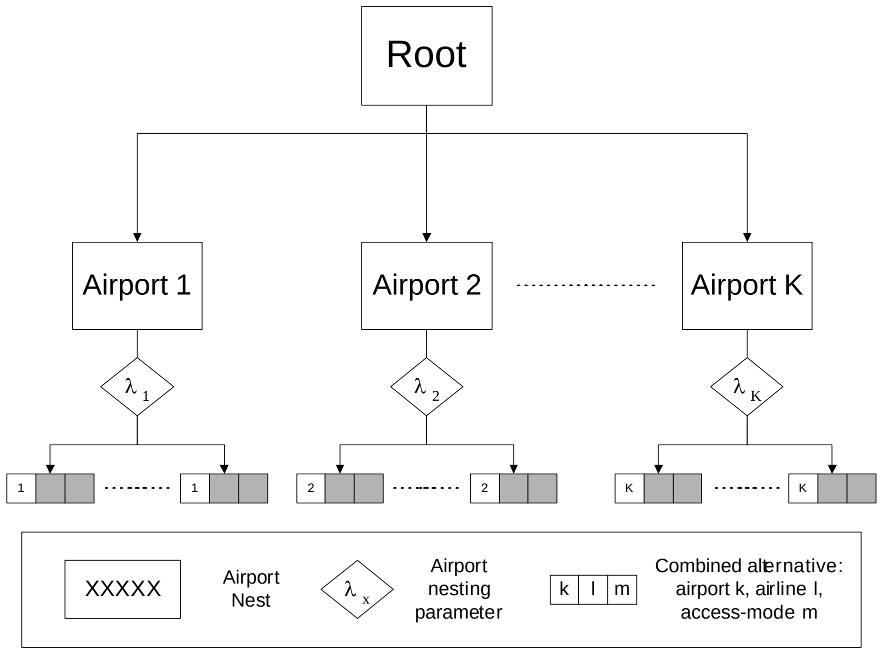

Figure 5: Structure of two level NL model, using nesting along airport-dimension

is illustrated in Figure 6, where λ_k is the nesting parameter associated with airport nest k, and π_l is the nesting parameter associated with airline nest l. Again, only a subset of the composite nests and of the triplets of alternatives is shown.

By noting that a model nesting airport choice above airline choice is not the same as a model nesting airline choice above airport choice, it can be seen that six possible two level structures arise in the present context.

While NL structures can, in this form, thus be used for analysing correlations along two dimensions of choice, it should be noted that multi level NL models have two important shortcomings which limit their potential for the analysis of choice processes of the type described in this chapter.

The main shortcoming in the present context is that the structures can

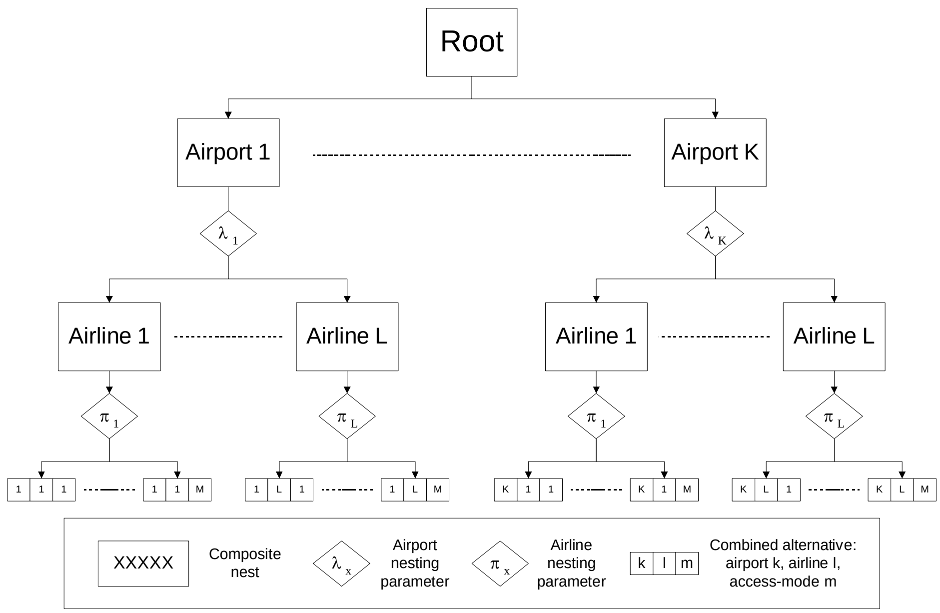

Figure 6: Structure of three level NL model, using nesting along airport-dimension and airline-dimension

be used for the analysis of correlation along at most $D-1$ dimensions of choice, with D giving the total number of dimensions. Indeed, using the example shown in Figure 6, it can be seen that, by adding in an additional level of nesting by access mode below the airline level, each access mode nest would contain a single alternative, as the airline nest preceding the access mode nest would contain exactly one alternative for each access mode. This means that the model could not be used to explore correlation along the access mode dimension. The same principle applies in the case of the other possible four level structures, where, in each case, the lower level of nesting becomes obsolete.

While a three level NL model can in this case be used to analyse the correlation along two out of the three dimensions of choice, the second shortcoming of the structure means that problems arise even with this task. 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 l inside the nest for airport k will only group together the options on airline l for that airport k. The same reasoning applies for other nests. As such, the model is not able to capture correlation between alternatives using airline l at airport k_1 and alternatives using airline l at airport k_2, which is clearly a restriction. This problem also applies in the other multi level nesting approaches. Aside from being a major shortcoming, this is also another reflection of the above comment that the order of nesting matters.

These deficiencies of multi level nesting structures can be addressed through the use of a Cross-Nested Logit (CNL) model, as discussed by Hess & Polak (2006) in an analysis of the combined choice of airport, airline and access mode in Greater London. In the present context, a CNL model is specified by defining three groups of nests, namely K airport nests, L airline nests and M access mode nests, and by allowing each alternative to belong to exactly one nest in each of these groups. As such, the structure addresses both of the shortcomings described above for the three level NL model. The structure is not only able to accommodate correlation along all three dimensions of choice, but does so in a simultaneous rather than sequential fashion. This means for example that the model is able to capture the correlation between all alternatives sharing airline l, independently of which airport they are associated with. At the same time, the correlation will be higher between alternatives that additionally share the same airport.

An example of such a model is shown in Figure 7, where, in addition to the previously defined λ_k and π_l, Ψ_m is used as the structural parameter for access mode nest m. Again, only a subset of the composite nests and of the triplets of alternatives is shown. Additionally, the allocation parameters, governing the proportion by which an alternative belongs to each of the three nests, are not shown in Figure 7.

5.3 Variation in behaviour across respondents

It should be clear from the outset that no two travellers are exactly the same in terms of their sensitivities to changes in attributes defining the alternatives. This applies most obviously in the case of factors such as air fare and access time, but also extends to the willingness to pay for flying on a certain airline or on a certain type of aircraft. On this basis, the assumption of a purely homogeneous population cannot be justified, and efforts need to be made to account for taste heterogeneity in the analysis of air travel choice behaviour, if major bias in the results is to be avoided.

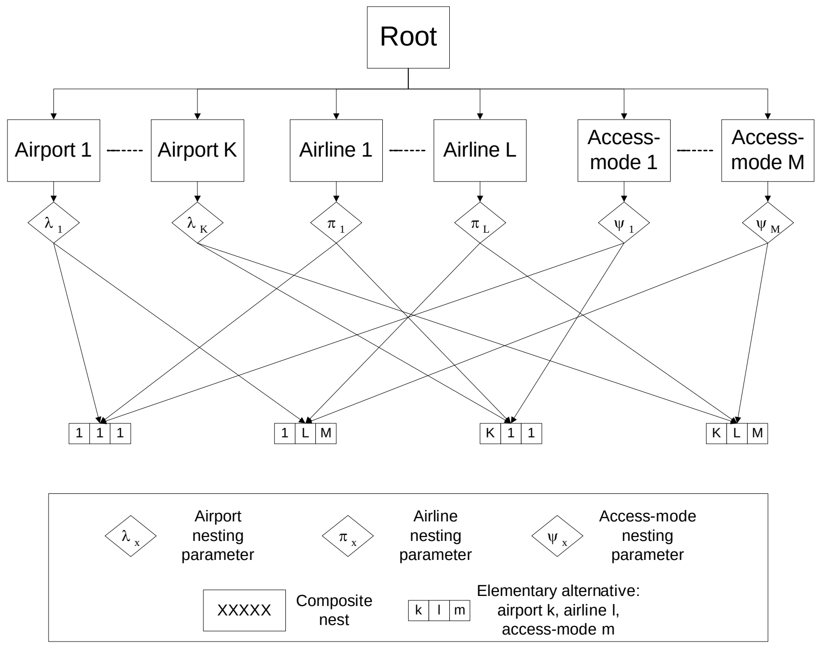

Figure 7: Structure of CNL model for the joint analysis of correlation along the airport, airline and access mode dimensions

Existing treatments of taste heterogeneity fall into three main categories, namely (in order of increasing complexity):

- discrete segmentations,
- continuous interactions, and
- random variations.

Discrete segmentations form the most basic approach for representing variations in sensitivities across respondents, either in the form of separate models for different population segments, or segment-specific coefficients within the same model. In the context of air travel behaviour research, possible segmentations would for example split the travellers by trip purpose, or by residency status (residents vs visitors). While straightforward to apply, they lack flexibility as they still assume homogeneity within groups. Additionally, there is the issue of defining the groups, which can be very arbitrary in the case where continuous measures such as income are split into discrete segments.

Continuous interactions can be used to allow for a variation in a given sensitivity as a function of a continuous socio-demographic measure such as income. In the context of air travel behaviour research, this would for example allow for a continuous interaction between a given traveller's income and his sensitivity to increases in air fares. As such, we would have $V_{i,n} = \ldots + \beta_{fare} \left(\frac{inc_n}{\overline{inc}}\right)^{\lambda_{income,fare}} fare + \ldots$, where inc_n gives the income for respondent n, with $\overline{inc}$ giving the mean income in the appropriate population segment. Here, a negative estimate would be expected for $\lambda_{income,fare}$, indicating reduced fare-sensitivity with higher income. While offering much more flexibility than discrete segmentations, continuous interactions impose greater estimation cost, and are only used sparsely in practice.

The gains in popularity of mixture models such as Mixed Multinomial Logit (MMNL) have meant that modellers increasingly rely on a purely random (as opposed to deterministic) representation of taste heterogeneity. Let's assume that the vector of taste coefficients β follows some multivariate continuous statistical distribution $f(\beta \mid \Omega)$ with vector of parameters Ω. Then, with $P_{i,n}(\beta)$ giving the choice probability of alternative i for individual n, conditional on a given realisation of this vector β, the unconditional choice probability is given by:

$$P_{i,n} = \int_{\beta} P_{i,n}(\beta)\, f(\beta \mid \Omega)\, \mathrm{d}\beta \tag{4}$$

With $P_{i,n}(\beta)$ being of MNL form, Equation (4) gives the corresponding MMNL choice probability, while, in the case where $P_{i,n}(\beta)$ is of a more advanced GEV form, Equation (4) gives the corresponding GEV mixture model. While offering very high flexibility, continuous mixture models are not only costly to estimate, due to the requirement to use simulation (cf. Train 2003), but also lead to a number of specification issues, such as the choice of random distribution (cf. Hess et al. 2005*b*). An alternative approach to the representation of random taste heterogeneity replaces the continuous mixture distributions with a discrete mixture approach (cf. Dong & Koppelman 2003, Hess, Bierlaire & Polak 2007). Here, the choice probabilities are given by:

$$P_{i,n} = \sum_{j_1=1}^{m_1} \cdots \sum_{j_K=1}^{m_K} P_{i,n}\left(\beta = \langle \beta_1^{j_1}, \ldots, \beta_K^{j_K} \rangle\right) \pi_1^{j_1} \ldots \pi_K^{j_K}, \tag{5}$$

where a given taste coefficient β_k has m_k mass points β_k^j, $j = 1, \ldots, m_k$, each of them associated with a probability π_k^j, where we impose the conditions that $0 \leq \pi_k^j \leq 1, \quad k = 1, \ldots, K;\ j = 1, \ldots, m_k$ and $\sum_{j=1}^{m_k} \pi_k^j = 1, \quad k = 1, \ldots, K$. While the issue of a choice of distribution does not arise in the case of discrete mixture models, there is now the problem of deciding on a number of support points to be used. Finally, in both cases, there are big issues in interpretation, as it is, without a posterior analysis, not possible to link the taste heterogeneity to socio-demographic information.

As shown in this section, the various methods for representing taste heterogeneity are characterised by differences in terms of flexibility, ease of implementation/estimation, and ease of interpretation. While a decision of what approach (or combination of approaches) to use needs to be made on a case by case basis, modellers should always be conscious of the fact that the flexibility of the advanced methods comes not only with a hefty price tag in terms of the cost of estimation, but crucially also leads to issues in interpretation.

5.4 Discussion

The modelling approaches described in Section 5.2 and Section 5.3 have quite separate aims; the analysis of inter-alternative correlation, along multiple dimensions, and the representation of deterministic and random variations in choice behaviour. It is quite reasonable to expect that both phenomena play a role in specific datasets. To minimise the risk of confounding between the two phenomena (cf. Hess et al. 2005*a*), it is in this case important to attempt to model the two phenomena jointly, with the help of an advanced GEV mixture model, or, although this is less appropriate in the present context (for numerical reasons), an error components Logit formulation or a Probit model (cf. Train 2003).

6 Summary and Conclusions

The discussion in this chapter has highlighted the complexity of the choice processes undertaken by air travellers, and has shown how they differ from behavioural processes in other areas of transportation. While highlighting the fact that travellers take decisions along a multitude of dimensions, the discussion has also noted that, in practical research, it is almost inevitable to use some simplifications of the choice process, partly because of modelling complexity, but mainly because of data issues. In this context, the *advantages* of SP data make the use of such datasets an important avenue for further research, potentially in conjunction with compatible RP data.

In closing, it is worth offering some guidance to modellers. On the basis of the discussions in this chapter, researchers are advised to:

- explicitly model the multi dimensional nature of the choice process,
- allow for deterministic, random, and continuous variations in choice behaviour,
- use advanced structures for correlation along all dimensions of choice, and
- avoid over-aggregation in level-of-service data.

Acknowledgements

This chapter was partly written during a guest visit in the Institute of Transport and Logistics Studies at the University of Sydney. The author would also like to acknowledge the input of John Polak in earlier stages of this work and would like to thank John Rose for useful feedback and comments.

References

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

Algers, S. & Beser, M. (2001), 'Modelling choice of flight and booking class - a study using stated preference and revealed preference data', *International Journal of Services Technology and Management* **2**(1-2), 28–45.

Dong, X. & Koppelman, F. S. (2003), *Mass Point Mixed Logit Model: Development and Application*, paper presented at the 10th International Conference on Travel Behaviour Research, Lucerne.

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

Hess, S., Bierlaire, M. & Polak, J. W. (2005*a*), *Confounding between substitution patterns and random taste heterogeneity*, CTS Working Paper, Centre for Transport Studies, Imperial College London.

Hess, S., Bierlaire, M. & Polak, J. W. (2005*b*), 'Estimation of value of travel-time savings using mixed logit models', *Transportation Research Part A: Policy and Practice* **39**(2-3), 221–236.

Hess, S., Bierlaire, M. & Polak, J. W. (2007), 'A systematic comparison of continuous and discrete mixture models', *European Transport* **36**.

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

McFadden, D. (1978), Modelling the choice of residential location, *in* A. Karlqvist, L. Lundqvist, F. Snickars & J. W. Weibull, eds, 'Spatial Interaction Theory and Planning Models', North Holland, Amsterdam, chapter 25, pp. 75–96.

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

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