Airports Commission Discussion paper 1: Aviation demand forecasting – a response by Birmingham Airport

Contents

1 Executive summary

The Department for Transport's (DfT) forecasts confirm that there will continue to be sizeable growth in passenger demand for aviation, and additional capacity will need to be brought online to cater for this. However, the DfT's forecasts should not form the sole basis for deciding on the location of future airport capacity. The primary purpose of the DfT's demand and allocation model forecasts is to 'inform long-term strategic aviation policy'. Its findings are highly contestable representations of theoretical constructs of future demand based on past behaviour. It lacks the granularity to engage with the aspirations of individual airports, their catchment areas and the future growth of the catchment economies.

Chapter one argues that the literal results of the DfT's model are only useful if the Airports Commission views airport capacity as a siloed decision about demand management. If the Airports Commission decision is based on what is best for the future UK economy, it will build a model that allows airports to plug the discrete regions of the UK into global wealth.

Chapter two explores analysis by York Aviation, commissioned by Birmingham Airport, into the DfT's forecasting model – Appendix 1. It reveals that there are underlying weaknesses in the assumptions, inputs, rationality and outputs underpinning the DfT's model. These are summarised in table 1:-

Table 1 – A summary of the caveats within the DfT's demand forecasting models

Caveat	Explanation
Exaggeration of the past	The DfT model is calibrated on past behaviour and represents an exaggerated version of the status quo. The historical data may no longer reflect current market interactions, it is not necessarily reflective of future behaviour, ignores the underlying potential of airport catchment areas and any future route development activities.
Lack of granularity	The economic growth inputs used by the National Air Passenger Demand Model (NAPDM) lack granularity, assuming uniform growth across the UK. There is no segmentation of inputs to reflect the UK's economic geography.
Allocation is very misleading	The current failure to recognise the limitations of the National Air Passenger Allocation Model (NAPAM) is the greatest failure of the DfT's approach to forecasting, notably the impact of the 'ballooning effect' from traffic spilling out of constrained airports and its inability to consider how airports compete with overseas airports for hub traffic.
Modelling is only half the answer	Even with a more enhanced version model, the DfT model is insufficiently robust to be a sole basis for policy decision. Wider evidence, patterns and aspirations of growth will always be vital to any decision about additional airport capacity.
A step change is needed	There needs to be a step change in the way the forecasts are used and considered, and a greater role for probability analysis by peer reviewers who sense check the findings.

2 What the DfT forecasts tell us about aviation policy

2.1 Do the DfT forecasts support or challenge the argument that additional capacity is needed?

The DfT's UK Aviation Forecasts 2013 argue that passenger demand will slow from the historic growth rate of 5% per annum to a new, lower trend rate of between 1-3% per annum by 2050. This decline has significant implications for policy-makers and businesses within the aviation sector. It represents a drop in the central forecast of 25 million passengers per year (mppa) by 2050 from the DfT's 2011 forecasts, and the allocation of traffic between the UK's largest ten airports has markedly changed, summarised in Table 2.

Table 2 – Percentage change in demand for seven major airports to 2020

Airport	Manchester	Gatwick	Heathrow	Luton	Stansted	East Midlands	Birmingham
Change 2011 vs 2013	-12%	+6%	-6%	+17%	0%	+33%	-40%

Under the current 2013 scenario, all airports in the South East will be full by 2030. And if no new capacity is added, Birmingham, Bristol, East Midlands, Manchester, Liverpool and Edinburgh airports will also be full by 2050. If, as seems unlikely, the results are accepted as robust and the Airports Commission uses them to answer a static demand-management question, there are two overriding long-term policy options.

i. Do nothing – by making the best use of existing capacity and pursuing policies that aggressively redistribute traffic, the UK would not need additional capacity to cater for forecasted demand growth for at least thirty years.
ii. Build a third runway now – Heathrow is effectively at capacity and the Commission should expand it now, ignoring all other considerations and make allowances for the possible need for further expansion by the mid-2030s.

The problem with both of these proposals is that they assume that policy-makers have a carte blanche to build runways or allocate traffic wherever they want. Equally, the demand figures are representative of a historic data about the past UK economy projected forward. Both the DfT and the Airports Commission recognise these failings, stating that 'in making any predictions about the future there is inherent uncertainty'. More importantly for policy-makers, 'the uncertainty reflected by the range of the national level is compounded at the level of the individual airport'. In other words, the DfT acknowledges that its model is poor at forecasting future growth at individual airports. It follows that the model is not capable of reconciling the demands of future regional GDP growth (and thus the true demand for connectivity) against available regional capacity.

Birmingham Airport believes that policy-makers would be taking a considerable risk if they decided to base any long-term aviation strategy on the DfT's 2013 growth forecasts alone. The risk of error would be further compounded if the forecasts were taken to mean that a long-term policy solution should solely rely on expansion at one airport. Birmingham Airport considers that one of the primary benefits of the Airports Commission is to challenge the status quo and consider other solutions to the aviation capacity issues. The success of the major UK regional economies – London, Birmingham, Manchester, Wales/ Bristol, Scotland, Newcastle – will depend on their ability to have a major airport providing the specialist international connectivity its local economy needs to prosper in the future. Considerations about future capacity should help shape the future, not just reflect the past.

3 A detailed critique of the DfT's forecasting model

The 2013 DfT forecasts attributed the downward revision of growth to changes in two main factors, the validity of which are scrutinised below:-

i. A series of judgement-based assumptions to reflect different levels of market maturity, which reduce the forecast demand by around 7% in 2030 and 21% in 2050.
ii. A predicted slowing down of the long-term decline in average air fares caused by the inability of airlines to cut operating costs in line with rising fuel and CO_2 emission costs.

Birmingham Airport believes that these findings are debatable because there are underlying weaknesses in the assumptions, inputs, rationality and outputs underpinning the DfT models. These are summarised in the sections below, but for a more detailed evaluation of the DfT's demand forecast see York Aviation's critique - Appendix 1.

3.1 The key inputs and assumptions

The task of creating a model that forecasts 31 airports in the UK at the same time, across a large range of passenger segments and destinations is undoubtedly difficult. There are numerous problems behind the key inputs, summarised below.

3.1.1 The CAA Passenger Survey data

- Data on origins and destinations is coarse.
- There is no information on actual journey times.
- The fare data collected is partial and unreliable as it relies on passenger memory.
- Many surveys are out of date.

The absence of an universal data set for all 31 airports collected at a single point in time is a major barrier to generating accurate passenger metrics.

3.1.2 Government data sources

The assumption inputs used by the DfT model are from official government sources or internationally recognised institutions, however:-

- The Office of Budgetary Responsibility has a tendency to be over optimistic.
- Macroeconomic forecasting is notoriously difficult.

While the use of inputs matches the Government's aim to ensure consistency across departmental policy analysis, in some cases the sources used are not necessarily the best ones in terms of their track-record for accuracy.

3.1.3 Customer choice

- The airline market is a characterised by monopolistic competition, but the DfT model ignores the impact of frequent flier schemes and airline preferences on consumer choice.
- Many destinations, particularly in the short-haul leisure market, are substitutable. This consideration is not included alongside airport substitutability.
- Customers are treated as though they have perfect information about the options available to them in terms of surface access and air travel. This is not the case, and undermines the DfT's assumptions that current travel patterns reflect the true demand in the market. In the event that passengers are given greater consumer information, there would be more choice and more competition in the market. This is not reflected in the current DfT passenger demand outputs.

As a result, assumptions about passenger choice may be unreliable and based on historic patterns of air services and competitive interactions.

3.1.4 Exaggerating the past

The DfT's model is calibrated on past behaviour and is not necessarily reflective of future behaviour. This leads to forecasts that represent an exaggerated version of the status quo. The model makes no assumptions about the underlying potential of the catchment area. This becomes a particular problem where historic data may no longer reflect current market interactions or indeed future economic growth plans.

3.1.5 Reliance on the 'spill' effect from Heathrow

The defining feature of the DfT's model is the allocation of traffic spilling over from capacity constraints at Heathrow. The model uses a mechanism which either pushes traffic to the nearest available airport or prices passengers out of flying. The concept is sound but its outputs are highly contestable because:-

- This is a theoretical construct and the way in which passengers behave in the face of constraints cannot be directly calibrated.
- The magnitude of effects, particularly the balance between pricing off and reallocation, is supposition but plays a significant role within the NAPAM outputs.
- The use of price elasticities to decide the pricing off substitution effect ignores future changes to surface access connectivity. For instance, High-Speed 2 (HS2) would change the cost, frequency and availability of passenger substitution across the UK airport network, notable at London, Birmingham and Manchester, yet no indication is given that this is accounted for within the DfT model. This is worrying given the likely transformative effect of HS2 on air-rail connectivity.
- Applying an air fare elasticity to the generalised cost facing passengers, which is a function of the monetarised cost of accessing the airport and waiting for flights, may be erroneous and produce skewed results.

3.1.6 Peer reviews should sense check results

Peer review experts have critiqued the DfT model, and generally supported its main workings. However, the reviews have been from a technical execution basis rather than considerations of rationality, i.e. whether the results make sense in the 'real' world:-

- The Airports Commission's discussion paper refers to the absence of 'probability analysis' within the DfT's approach. This should be rectified.
- Regional airports have particular concerns about the validity of the DfT model because it is often at odds with their own market analysis. This makes it difficult for regional airports to use national statistics during their route development activities.

3.2 Weaknesses in the rationality of the results

3.2.1 Poor granularity

One of the biggest issues for Birmingham Airport is that the economic growth assumptions are insufficiently granular. They do not adequately differentiate between regional economies, or major economies within the main geographical international markets:

- There is no segmentation of economic inputs in terms of the geography of the UK - they are all national. The ability of the DfT model to make suitable and accurate adjustments based on the population and income of the 455 origin/ destination zones is unclear.
- The use of four geographical zones as the basis for the definition of the international market within the NAPDM is questionable. The current DfT groupings do not now offer a sufficient level of granularity to enable a real picture of changing spatial dynamics to emerge.

3.2.2 Poor treatment of market maturity

The DfT's analysis of the impacts of market maturity is at the heart of its downward revision of growth trends from 2011 compared to 2013, however:-

- Assuming that the EU Emissions Trading Scheme will lead to an increasing level of carbon costs is highly contestable. The Department of Energy and Climate Change (DECC) forecasts have changed significantly over recent years, reflecting an immature marketplace. The DfT model does not account for the possibility that estimates of carbon prices will continue to move substantially with knock-on effects to air traffic forecasting.
- The assumption that the maturity of the UK market will reduce total demand by around 21% is the central case and is based on analysis by Dr Anne Graham dating from 2000. Whilst notable, this methodology is out of date and shows that there is considerable uncertainty behind a major input for the DfT's future modelling of maturity effects. A small error in this judgement could result in a very significant change in the market.
- Aviation markets in the UK are not uniform in their maturity, their relative sensitivity to economic growth or changes in prices. Higher price elasticities in relation to economic growth and potentially air fares might be experienced in relation to demand in regions away from London and the South East.

3.2.3 Treatment of uncertainty

The Airports Commission's discussion document acknowledges that the treatment of uncertainty in inputs by the DfT relies on sensitivity and scenario testing. This treatment is relatively weak because:-

- The scenario testing appears relatively unsophisticated and the scenario testing appears to be relatively arbitrary in the inputs it varies and therefore the high and lows it actually defines.
- The use of probability techniques, such as Monte Carlo analysis (see York Aviation – Appendix 1), would add to the accuracy of the DfT forecast outputs.

3.3 The limitations of the NAPAM

3.3.1 The missing parts are ignored

The NAPAM is sensible, but the failure to recognise the limitations of the model is the greatest failure of the DfT's approach to forecasting. It is based on a monetarised generalised cost function consisting of three elements: access cost, frequency cost and capacity constraint costs. There are three missing components:

- Air fares – these were taken out of the model because the data was criticised for being unreliable, but the implications of exclusion need consideration. The DfT's assumption that average fares remain constant ignores the fact that routes are not exogenous entities, e.g. capacity constraints and their interaction across different routes. This results in problems in relation to how and which passengers are priced off or reallocated and from which markets.

- Flight times – the failure to include a flight time cost reduces the overall absolute starting point generalised cost, potentially upsetting the balance between a constraint cost and other elements of the generalised cost equation.
- Access journey times – there is no indication that the current DfT model investigates the likely impact of air-rail improvements. How these will affect door-to-door journey times, surface access connectivity or lead to a reduction in the cost of CO_2 emissions resulting from the shift away from cars as the dominant modal share for access airports in the UK.

3.3.2 The frequency term and mechanism are misleading

The frequency term used by the NAPAM articulates the supply side reaction to the growth in demand in the market. It is a key driver in market share over time. But there are issues when frequency is not specified correctly and issues around how this is then calibrated:-

- The 'ballooning effect' – the calibration of the generalised cost facing passengers is weighed more on frequency that access times, drowning out passengers' considerations around distance. It ignores how airports will grow in line with the potential of their own catchments, so rival airports with a frequency advantage in 2013 quickly develop more traffic over time.
- Reinforcing the past - future flight allocation reinforces historical trends because the difficulty with a limited generalised cost function is that if an airport has no history in a given market its attractiveness in terms of its constant will be low. This is a major weakness, especially for Birmingham Airport because it ignores our route development activities.
- This is particularly important for manufacturing centres. Multinationals disperse their supply-chain across the global market, but seek to maintain just-in-time production methods. This places a premium on direct aviation links and is a contributing factor to the emergence of new city pairs, and the resulting new aviation links between them. Failing to make any assumptions or predictions about the future interactions between regional UK economies and international trading partners means that demand forecasts are too static. The Airports Commission needs to use a demand model that incorporates dynamic analyses and engages with future economic scenarios.

3.4 The weakness of the treatment of the international passengers market

The NAPDM has three major shortcomings in its analysis of international passengers:

- It does not consider the potential size of the future international market or indeed how UK airports compete with overseas airports for this market. This is a crucial consideration for hub capacity considerations, if that is the emphasis of the Airports Commission's analysis.
- The model assumes Heathrow has a 'right' to the slice of the international to international interline market. If Heathrow continues to lose destinations and become a less attractive hub there is the potential for others to pick up substantially more traffic from Heathrow than is implied by the DfT. The real determinant of hubbing is the extent to which airlines wish to attract transfer traffic either to sustain their network or to maximise profitability by creating externalities for the customer (in extended surface access, Hotels, etc).
- The DfT model needs to effectively consider the competitive dynamic with other hub airports, especially the growth of Istanbul and Turkish Airlines.

4 Irregular results from the DfT's forecasting model

4.1 Case studies of irregularities created by DfT forecast

Analysis by Birmingham Airport and York Aviation reveals that the findings made by the DfT's aviation forecasting model are highly contestable. Table 3 shows gives examples of the irregularities caused by the inconsistencies within the model:

Table 3 – evidence that the DfT forecasts generate huge irregularities

Example	Explanation
Long-haul traffic at Stansted Airport collapses in favour of Birmingham Airport	Increases by 2.5 mppa by 2030 but all lost by 2050, because sucked in by Birmingham Airport, which goes from +2.6 mppa long-haul passengers in 2030 to c. +15mppa by 2050. This is caused by the 'ballooning' effect that underpins the DfT's frequency term.
Medium term traffic collapses at Birmingham	To generate the 2013 model, the DfT discussed plans with certain airports – such as Gatwick's 55m strategy and Luton's terminal expansion – and included them within its model. This has generated a 40% collapse in Birmingham Airport's medium-term traffic forecasts, which as the example above shows, miraculously returns in the very-long term. This level of traffic fluctuation is not reflective of the true state of the airport market.
Growth at East Midlands	The 'spill over' effect takes growth in passengers from +100,000 by 2030 to +14.1 mppa by 2050. A jump of 14 million passengers is improbable.
Long-haul traffic in England	The 'spill over' effect pushes demand from London up until it hits an airport with capacity and a track record in long-haul traffic, wherein a ballooning effect occurs until that airport reaches capacity and then the process rolls onto the next available airport. This is not how air services develop in practice.

5 How to improve the DfT forecasting model

5.1 Recommendations on how to improve the DfT forecasting approach

1. There needs to be a step-change in the way the forecasts are used and considered. The model is just that. It is a piece of evidence that needs to be viewed in the round. It is not sufficiently robust or accurate to be a sole basis for policy decisions. An essential function of the Airports Commission should be to recommend alternative solutions which could radically change the aviation map, and aviation forecast splits, of the UK in the future. Wider evidence on growth and patterns of growth will always be vital. Even an enhanced version of the model should still only serve to inform policy, not dictate it.
2. The approach to analysing uncertainty in the inputs to the NAPDM needs overhauling, and should include probability analysis and an appropriate approach to uncertainty.
3. The frequency term within the model needs re-examining to consider the 'ballooning' effect it creates when traffic is spilled out from constrained airports. The model needs to place how the airports will be developing their own catchment areas at the centre of its model.
4. The treatment of international to international hub traffic needs to be substantially overhauled, particularly in the context of the terms of reference for the Davies Commission.
5. There needs to be substantially more granularity within the model as it is inaccurate to paint the UK economy and aviation market as homogenous.
6. The treatment of price elasticities and future changes to surface access connectivity within the model needs to be re-examined, particularly given the likely transformative effect of HS2.
7. The Airports Commission needs to revisit Dr Anne Graham's work in 2000 on market maturity and re-examine DECC's CO_2 emissions forecasts. Small errors in these considerations generate significant variations in the model's output and are at the heart of the downward revision in traffic between the DfT's 2011 to 2013 forecasts.
8. Peer review needs to include industry forecasting expertise to sense check the results to how airlines are actually likely to respond to growth and constraint, as well as how new entrants might be attracted to the market.

York Aviation

Birmingham Airport

Overview of and Commentary on DfT's Approach to Demand Forecasting

1 BACKGROUND

1.1 In February 2013 the Davies Commission issued the first in a series of discussion papers designed to provoke debate and ultimately inform its thinking with regard to its consideration of the requirement for and potential location of new airport capacity in London and the South East. This first discussion paper considers demand forecasting and specifically it asks for comments and suggestions with regard to how the current DfT approach to demand forecasting could be improved.

1.2 Birmingham Airport has asked York Aviation to provide an overview of the DfT model, its operation, the key inputs to the model and an assessment of its strengths and weaknesses to assist it in preparing its response to the Davies Commission. Our paper is based on:

- the discussion paper published by the Davies Commission;
- the UK Air Traffic Forecasts 2013 published by DfT in January 2013;
- the UK Air Traffic Forecasts 2011 and the accompanying peer review documentation;
- our own knowledge and experience in relation to the DfT model developed over the last 15 years.

1.3 The paper has been structured as follows:

- initially we set out an overview of the structure and operation of the DfT model;
- then we set out a more detailed schematic of the DfT model, including specifying the key inputs and assumptions within the process to the extent that they are known and published;

- we then move on to consider the strengths and potential weaknesses relating to the inputs to the model and consider in overview how weaknesses might be addressed;

- the strengths and weaknesses of DfT's approach in terms of the methodologies and specification of the model;

- we seek to highlight possible evidence within the 2011 and 2013 forecasts of the problems with the model.

2 OVERVIEW OF THE DFT MODEL

2.1 The DfT forecasts are developed using two separate models[1] that interact with each other. These models essentially follow a four stage process to get to a passenger demand forecast:

- the **National Air Passenger Demand Model** initially forecasts the total quantum of passenger demand for the UK as a whole split in to a number of passenger segments. This model is based around time series regression analysis of the link between historic air passenger demand and the key drivers of demand, identified as economic growth and the cost of travel. The passenger segments are based around UK and foreign passengers, those travelling for business or leisure, and five separate world zones (domestic, Europe, OECD countries, Newly Industrialised Countries and Less Development Countries). This model also estimates the number of international to international interline passengers expected to use UK airports. There are a total of 19 different passenger segments modelled;

- the **National Air Passenger Allocation Model** splits the level of demand nationally in any given market segment in any given year in to 455 origin/destination zones around the UK based on CAA Passenger Survey data adjusted to reflect changes in demographics drawn from DfT's TEMPRO planning model, which includes estimates of changes in population in each zone;

- the **National Air Passenger Allocation Model** is then used to allocate this demand to one of the 31 UK airports covered by the model. This is based on a statistical analysis of historic patterns of passenger choice based around access time to different airports and the frequency offered for the relevant route at different airports. This is initially done on an unconstrained basis where no capacity constraints are assumed at the airports in the model. It is at this point that a more detailed breakdown of destinations used by passengers is introduced. The model considers passenger demand in terms of 48 routes or groups of routes. These are

[1] The DfT model also feeds in to models that consider fleet mix and ultimately CO2 emissions. These have not been considered here.

generally either large individual destinations or, particularly for longer haul destinations, world areas;

- following the initial unconstrained allocation of passengers, the model moves on to apply capacity constraints to produce a constrained forecast for the 31 airports. The model uses a constraint or shadow cost to make individual airports less attractive to passengers at airports where the capacity of the airport is breached. This has two effects. Firstly it simply prices some passengers out of the market entirely based on a price elasticity (i.e. they are assumed to no longer travel) and secondly it results in some passengers being reallocated to other airports.

2.2 The model provides demand forecasts in a considerable level of detail. It potentially enables individual routes (or route groups) to be forecast, including business/leisure splits and inbound and outbound proportions. However, DfT recognise that the level of uncertainty around the forecasts increases as the level of disaggregation increases and hence limited detail is usually published. However, recent statements, notably in the HMRC report on the impact of price differentials, suggest that the level of uncertainty can be high even at an individual airport level. Indeed, the 2011 forecasts rounded the results for larger airports in the longer term to the nearest 5 million passengers in any given year. Furthermore, at para 1.4 of the 2013 forecasts, DfT states that *"the uncertainty reflected in the range at the national level is compounded at the level of the individual airport. At the airport level, the DfT forecasts may differ from local airport forecasts. The latter may be produced for different purposes and may be informed by specific commercial and local considerations."* This does not give confidence as to the effectiveness of the model for forecasting individual airports. Clearly this has implications in relation to the forecasting of hub capacity in London and the South East.

3 KEY INPUTS AND ASSUMPTIONS

3.1 In **Figure 1** below we have set out a schematic diagram that shows how the process described above works and the key inputs and assumptions that feed in to the model at each stage. This diagram essentially seeks to give a simplified view of the flow diagram set out by DfT in its latest forecasting document[2].

3.2 It should be noted that the individual passenger segments or markets may use different combinations of either the key drivers of overall demand that feed in to the National Air Passenger Demand Model or passenger decision criteria that feed in to the generalised cost function within the National Air Passenger Allocation Model. However, the broad principles remain constant across all the markets.

[2] UK Aviation Forecasts – Department for Transport (January 2013), page 14.

3.3 In general, the assumption inputs to the DfT model are, wherever possible drawn from official government sources. For instance, UK economic growth metrics and exchange rates use Office for Budgetary Responsibility forecasts and oil prices and carbon allowance costs are based on forecasts from DECC. Other inputs, notably foreign GDP, are drawn from sources such as the IMF and Enerdata. Trends in key aviation industry variables such as fuel efficiency or airline productivity are based on DfT analysis of industry evidence. In general, the aim is to ensure consistency across Government in terms of the inputs used for policy analysis. It may mean that in some cases the source used is not necessarily the best one in terms of its track record of accuracy.

3.4 Despite the level of detail available regarding the model and its inputs and assumptions, there are a number of areas where there is uncertainty regarding its operation, specification and inputs. These include:

- ✈ the basis on which the allocation of total demand to the origin and destination zones around the UK is changed over time is not entirely clear. The 2013 Aviation Forecasts appears to suggest that this based on demographic trends (see page 14) but the Airports Commission Discussion Paper suggests that there may be an income element to this as well (see Page 16);

- ✈ whether flight times are included within the generalised cost function within the Allocation Model is also unclear. The recent HMRC report on the impact of price differentials at UK airports states clearly that only surface access costs are considered. However, the 2013 Aviation Forecasts report (page 24) suggests that flight duration is a determinant that may be important in some markets;

- ✈ the precise workings of the element of the model that estimates frequencies at each airport in different markets in the future is not clear. Broadly, the model increases frequency over time to meet the demand allocated to each airport where existing services are in place. However, it also tests for the viability of new routes as part of this process. It is this element of the process that is unclear. We assume that a notional new service is added to an airport to test for viability but this service interacts with other potential new services. Are new services tested independently or in groups or all at the same time. Clearly, the extent of competition will be a key determinant as to whether the route is viable or not. Currently, we do not have any visibility as to this process and hence cannot comment on its appropriateness.

Figure 1: DfT Aviation Forecasting Framework

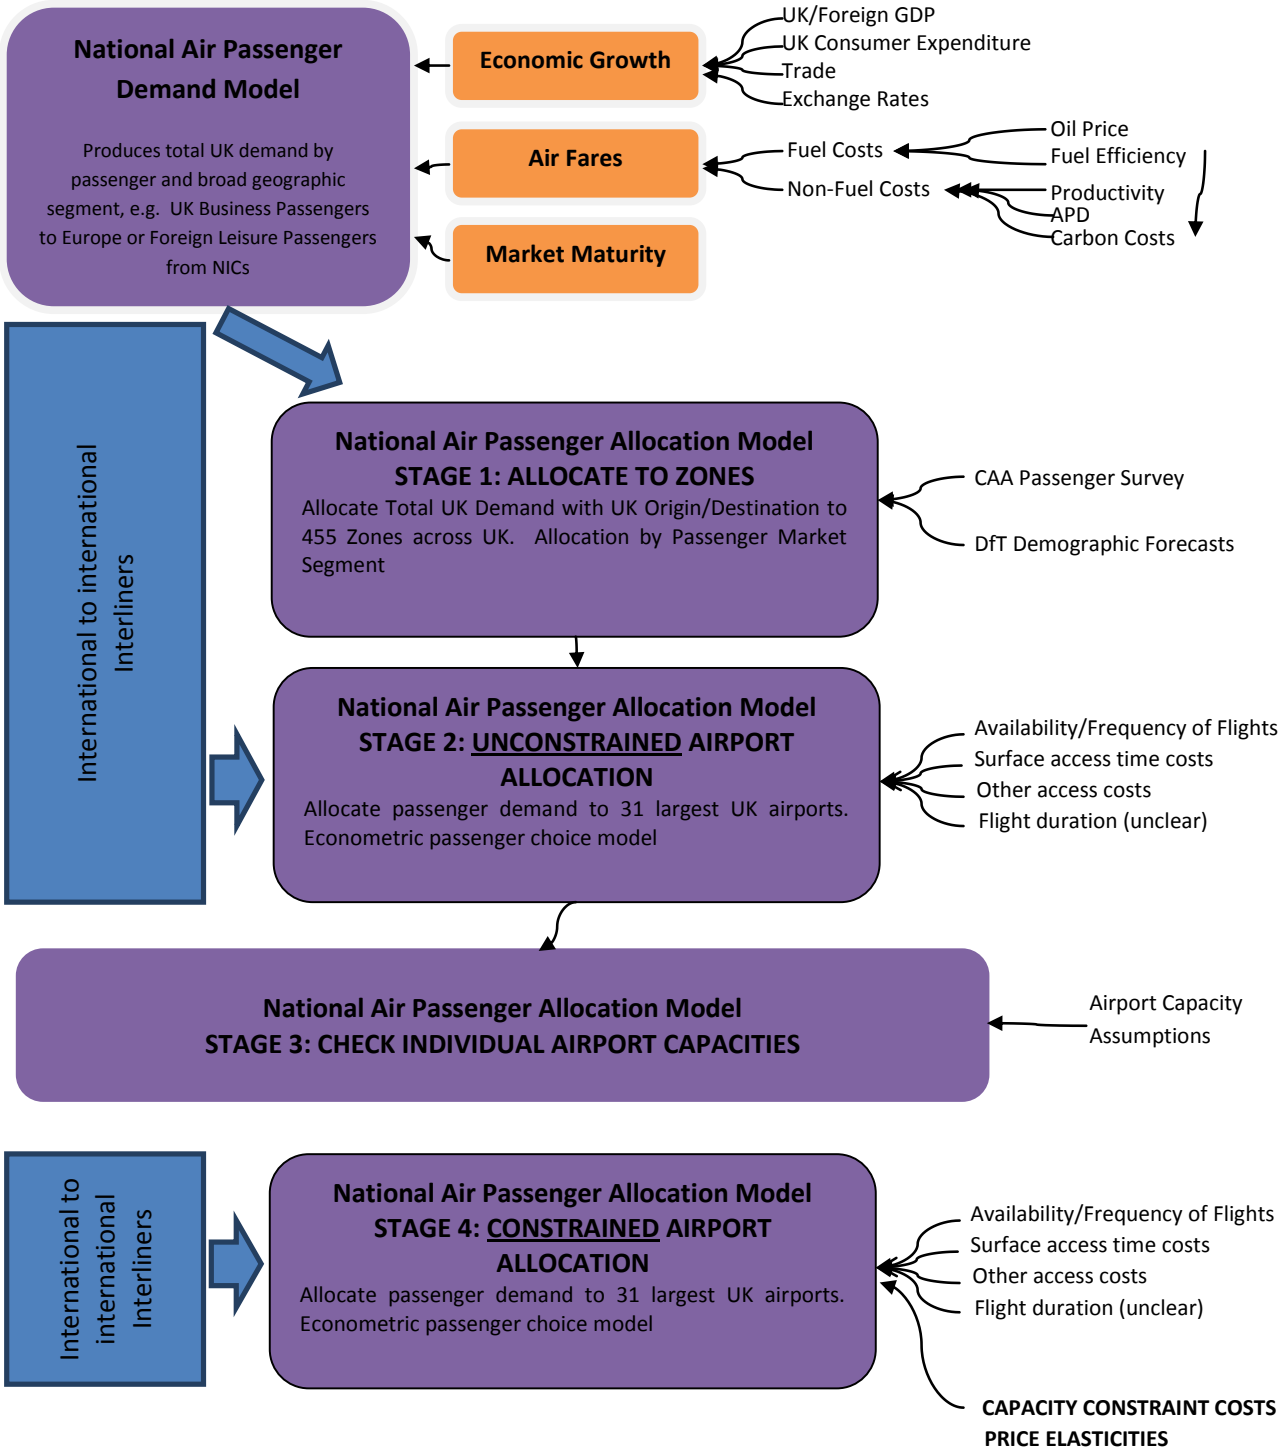

3.5 Before moving on to comment on the strengths and weaknesses of various elements within the model, it is worth commenting on the DfT model and approach in overview. At the outset, it should be said that the task attempted is ambitious. Creating a model that forecasts 31 airports in the UK at the same time, across a large range of passenger segments and destinations is unquestionably difficult. The predominant problem is the availability and reliability of data. The CAA Passenger Survey is an invaluable source of information but it is far from perfect. Origins and destination areas are relatively coarse, there is no information on actual journey times within it, the fare data collected is partial and notoriously unreliable and hence unusable in this circumstance (as it relies on passenger memory and response), there is no universal dataset for all 31 airports collected at a single point in time and, indeed, some of the airports have not had surveys done at them for many years and hence passenger choice assumptions may be unreliable and based on historic patterns of air services and competitive interactions.

3.6 There are also significant elements in customer choice that are not even considered, such as the influence of frequent flyer schemes, preference (or otherwise) for particular airlines or the fact that there is in fact substitutability not just between airports but also destinations (one short haul sun destination is potentially very similar to another).

3.7 It should also be remembered that the model is, by necessity, calibrated on past behaviour. This is not necessarily reflective of future behaviour. This has a tendency to lead to an exaggerated version of the status quo as demand grows in to the future. Airports develop along the lines of what they do now rather than in line with the actual underlying potential of their catchment areas because that is what the statistical analysis of past behaviour identifies. This can be a particular problem where historic data may no longer reflect current market interactions.

3.8 It is also worth keeping in mind the impact on the model of perhaps the key defining feature of the UK aviation market currently, the constraint in the London system centring on Heathrow. The model produces constrained airport forecasts using a mechanism which prices passengers away from constrained airports within the allocation model. This produces what is known commonly as ‘spill’ traffic, which is then allocated to other airports within the model. However, it should be recognised that this is a theoretical construct and the way in which passengers will behave in the face of constraint is assumed but cannot be directly calibrated. The broad concept is undoubtedly sound but the magnitudes of effects, particularly the balance between pricing off and reallocation, is largely a matter of supposition.

3.9 The DfT model has been the subject of significant investment and development time over the last 10 to 15 years. It has also been subject to Peer Review by a range of experts. These reviews have identified what might be considered relatively minor problems with the models but in the main have been positive. However, it should be recognised that these reviews have largely come at the model from a technical execution

basis rather than a consideration of rationality and whether its results do in fact 'make sense' in the real world. It is here that there has been considerably greater criticism, particularly from the airports industry. It is notable that there has, in recent years, been no sense check by industry experts as was normal practice in the 1990s. Regional airports in particular have concerns about the validity of the model, often relating to the fact that their own market analysis and forecasting is often significantly at odds with the model results. Below we have sought to comment from both perspectives, identifying some technical weaknesses and highlighting concerns regarding the rationality of results.

4 KEY INPUTS: STRENGTHS AND WEAKNESSES

4.1 Below we have considered specifically some of the key inputs to and assumptions in the DfT model and their strengths and weaknesses. This is as opposed to comments around methodology adopted and the specification of the models.

Economic Growth Assumptions

4.2 The economic growth forecasts feed in to the National Air Passenger Demand Model that defines the total demand in the UK in any given year. Currently, the DfT model uses UK economic growth forecasts that are developed by the Office for Budgetary Responsibility. This has the advantage that forecasts are consistent with other UK government policy analysis and that, in theory, the economic forecasts should be independent and not subject to particular political or commercial perspectives. The concern must be, however, that the OBR has not proved to be a particularly accurate forecaster. By its own recognition it has had a tendency to be over optimistic in relation to GDP growth[3]. This is, however, primarily a short term issue and it should be recognised that forecasting of this type in the current economic climate has been very difficult.

4.3 In relation to economic growth inputs for the UK, perhaps the greater issue in terms of the inputs to the DfT modelling is the fact that there is no segmentation in terms of the geography of the UK. Everywhere is assumed to experience the same rate of economic growth and, consequently, the same growth in air transport demand relating to economic growth. This is patently not true. There is substantially variation in the rates of economic growth across the regions and nations of the UK, which could clearly influence the speed at which individual airports grow. The additional difficulty here is understanding the extent to which these spatial differences may in fact be reflected when total demand is allocated to the 455 origin/destination zones around the UK. The adjustments made for population and possibly income growth in

[3] Forecast Evaluation Report – Office for Budgetary Responsibility (October 2012).

these zones may go some way towards addressing this issue but at present it is not possible to know.

Definition of International Markets within the National Air Passenger Demand Model

4.4 There must be some concerns as regards the relatively broad geographic groupings in to which international traffic is defined. Excluding domestic traffic, there are four geographic zones used by the DfT model to determine overall growth rates, Europe, OECD, NIC and LDCs. Separate growth rates are estimated using time series regression techniques for each of these groups. However, there must be some concerns as regards the homogeneity within some of these groups. In other words, is it sensible for instance to apply the same growth rates to US traffic as to Australian traffic. With increasing globalisation and the centres of economic growth in the world shifting and polarising, it must be questionable as to whether the current DfT groupings now offer a sufficient level of granularity to enable a real picture of changing spatial dynamics to emerge.

Immaturity of Carbon Markets and Impact on Forecasts

4.5 Moving forward one of the key drivers of the lower overall growth identified by DfT in its 2013 forecasts is the end to what has been a sustained period of falling air fares. One of the main causes of this is the increasing level of carbon costs associated with the EU Emissions Trading Scheme (or any similar future scheme). The difficulty here is the certainty that can be placed on the current forecasts of carbon allowance prices. The DfT model uses forecasts from DECC as its input to the model. These forecasts have changed significantly over recent years as approaches to carbon valuation have changed and better understanding has been developed of the operation of the market. However, the market remains relatively immature and as a consequence there must be the possibility that estimates of carbon prices will continue to move substantially with knock on effects to air traffic forecasting.

Market Maturity

4.6 The DfT model includes an assumption that over time air transport markets will mature and become less responsive to economic growth. This in itself is not an unreasonable assumption. Most products have a life cycle of this nature, whereby in the early stages take up is rapid compared to growth drivers and ultimately this slows. This inclusion of considerations around market maturity should be considered a strength. However, equally, the actual assumptions and methodology could be considered a weakness. The DfT assumptions are based around work undertaken by Dr Anne Graham of the University of Westminster in 2000. They essentially assume that different markets are currently at different levels of maturity and that therefore maturity assumptions should start at different points in time and that the effect will result in a gradual fall and convergence in economic growth elasticities over the long term. Whilst the general presumption may be valid, the work upon which the assumptions used are based is somewhat out of date and may not give a realistic indication of the current state of maturity of markets. Whilst, the DfT is quite clear that ultimately these are judgemental assumptions, this may be the only possible approach given there is very limited evidence of this effect to date against which some form of quantitative analysis could be performed. However, this means that there is considerable uncertainty in the future modelling of this effect. This would matter less if the impact was marginal but it is not and, as a consequence, a significant driver of future market growth rates is in fact judgemental. By 2050, the impact of market maturity is to reduce total demand by around 21% in the Central Case. A small error in this judgement could result in a very significant change in the market.

4.7 A further issue that has not been considered within the DfT forecasts is the extent to which aviation markets in different parts of the UK may in fact be more or less mature and hence more or less sensitive to economic growth or indeed changes in price. It would seem reasonable to assume that the London market is a relatively mature air transport market. However, it would also be reasonable to suggest that many UK regional airports are some way behind London in terms of the product life cycle. Hence, it would be reasonable to suggest that higher elasticities in relation to economic growth and potentially air fares might be experienced in relation to demand in regions away from London and the South East. This is something that has been suggested for some time by regional airports across the UK and has, to some extent, been recognised by DfT in the past. Earlier versions of the UK aviation forecasts produced around the time of the Future of Air Transport White Paper in 2003 included different growth rates for London and the regions to reflect this point.

General Treatment of Uncertainty in Inputs

4.8 DfT clearly recognises the importance of considering the impact of uncertainty in relation to the inputs in to its modelling process. It publishes sensitivity tests that examine the impact of varying individual inputs and also scenarios which alter a number of inputs at the same time to examine high and low case forecasts. However, although there are a considerable number of these tests they are relatively unsophisticated and the scenario testing appears to be relatively arbitrary in the inputs it varies and therefore the highs and lows it actually defines. It is perhaps therefore reasonable to suggest that currently the treatment of uncertainty in the DfT model is relatively weak. The use of probability techniques, such as Monte Carlo analysis, to analyse uncertainty would have considerable merit and is something that could be built in to the DfT framework. This is a key theme in the Davies Commission Discussion Document.

Pricing Off Elasticities

4.9 The recent HMRC report on the impact of price differentials at UK airports identifies that the mechanism by which passengers are priced out of the market within the model in the event of an airport being constrained uses the air fare elasticities identified within the National Air Passenger Demand Model. This may be a misapplication of these elasticities as they are being applied to generalised cost rather than an air fare. The generalised cost facing passengers is a function of the monetised costs of accessing the airport and waiting for flights. It does not include the fare as this variable has been dropped. Therefore, applying an air fare elasticity to this generalised cost may be erroneous and produce skewed results.

5 METHODOLOGY AND SPECIFICATION: STRENGTHS AND WEAKNESSES

5.1 Above we have outlined some strengths and weaknesses in relation to the key inputs and assumptions in the DfT model. However, this only presents some of the picture and in some cases, while the DfT model has flaws, these are potentially universal flaws that will face any model. Below we have set out some key issues in relation to the methodology and specification of the model, essentially the way in which the model actually works. At the outset, it is again helpful to make a couple of overarching points:

- ✈ our comments relate almost exclusively to the National Air Passenger Allocation Model. This is by some margin the more complex of the two models and in our view is the one that gives rise to most concerns. The regression techniques and specification of the National Air Passenger Demand Model are well established,

less complex and generally appear robust. While it is possible to question inputs to this process, the model itself appears robust;

- while we have significant concerns regarding the National Air Passenger Allocation Model and the results it produces, the type of model and the overall approach is sensible. The model is a multinomial LOGIT model, a type of model that is commonly used for transport forecasting in competitive markets. This type of model's ability to statistically analyse passenger choice in competitive environments is potentially powerful. However, these models are just tools. Their strengths and weaknesses need to be recognised, their results interpreted along with other evidence and, ultimately, the outputs considered for rationality. As much as anything it is the current failure to recognise the limitations of such a tool that is the greatest failure of the DfT's approach to forecasting.

Composition of the Generalised Cost Function

5.2 At the centre of the allocation model is a monetised generalised cost function. This identifies the cost to passengers of using different airport options and, based on past patterns, estimates the market share of each airport option on this basis. In broad terms, there are three elements to this generalised cost within the model:

- Access costs – the time and costs associated with getting to a particular option;

- Frequency costs – the cost associated with waiting for a flight at the relevant option. The greater the frequency, the lower the wait time and the lower the cost;

- Capacity constraint costs – if an airport is capacity constrained, the model applies an additional cost to this option to make it less attractive to passengers.

5.3 Clearly this is a simplified view of the world as are all models and DfT has stated that there are of course other components to a passenger's decision that are not reflected in this generalised cost. However, there are two that are potentially particularly important:

- Air fares – the level of air fare on offer is clearly an important consideration in passenger choice. If one airport offers consistently lower fares, it is likely to be more attractive. The fare term was excluded following comments in the 2011 Peer Review of the model on the grounds that the data available was unreliable. We would wholly support that conclusion but the implications of excluding the fare do need to be considered. DfT suggests that over time average fares to a single destination from different airports are probably similar and hence it is not likely to be a significant determinant in airport choice. This may to some extent

be true (although we would question if this is really the case in situations where LFAs at one airport are competing with full service airlines at another) but it is not just the relationship within the route market that is important within the model. Routes are not exogenous entities. As soon as capacity constraints enter in to the equation, then there is interaction across different routes. The overall size of the generalised cost then becomes important as constraint costs will be a larger proportion of overall cost on some routes than others and if a component as important as the fare is excluded from the original absolute generalised cost the impact of the constraint cost within this process could be skewed. This would result in problems in relation to how and which passengers are priced off or reallocated and from which markets;

- ✈ flight times – as described above, there is some confusion as to whether flight times are included within the generalised cost function or not. However, if they are not there is again the potential for difficulties once constraint costs start to be applied within the system. Again, the failure to include a flight time cost will reduce the overall absolute starting point generalised cost, potentially upsetting the balance between a constraint cost and other elements of the generalised cost equation. The effects are potentially similar to that of not including the fare.

Frequency Term and Mechanism

5.4 The specification of the frequency term within the model and the way in which it operates is central to many of our concerns regarding the DfT model. The frequency term is very important as it essentially articulates the supply side reaction to the growth in demand in the market. It is a key driver in market share over time. If frequency grows at one airport and not another, the former will become stronger and stronger over time, gaining market share. This is in essence correct. However, the difficulty comes if this frequency term is not specified correctly or there are issues around calibration. In our view both are an issue here.

5.5 We are concerned that over time frequency simply becomes too important a part of the generalised cost facing a passenger. The model is calibrated on a balance between access time and frequency that is correct currently. However, access times do not change significantly in most cases over time while clearly the number of frequencies does. The result is that the balance in the generalised cost function may be becoming upset, with frequency becoming more and more important, drowning out passengers' considerations around distance. This has the effect of making it hard for the model to recognise airports' abilities to grow in line with the potential of their own catchment areas as they begin to compete more and more on frequency rather than access time. The result is that airports that gain a slight advantage over others in one market or another rapidly develop frequency that sucks in more and more traffic resulting in a 'ballooning' effect in terms of the demand they

attract. We provide some examples of what believe to be this effect below.

Calibration is Beholden to History

5.6 As stated above, the allocation model is calibrated on historic patterns of behaviour by passengers. The result is that in some markets if an airport has no track record of performance it is unlikely to attract any traffic. The model estimates regression coefficients for the key choice determinants described above. It also estimates what are termed airport constants. These could be seen as reflecting the innate attractiveness of the individual airports to passengers within the given market stemming from the elements of passenger choice that are not modelled. The difficulty is that with a limited generalised cost function, if an airport has no history in a given market its attractiveness in terms of its constant will be low. As a result, even if there is demand in the future for a service it has not served before, it is very hard within the model for it to attract the traffic and sustain the service in the face of competition from incumbent airports in the market. This can lead to perverse patterns of development; particularly again when constraint effects are applied and significant volumes of demand start to spill out of London and the South East.

Treatment of International to International Interline Passengers

5.7 The way in which international to international interline passengers are forecast in the model is another area of particular concern and is particularly pertinent to the current debate around new capacity in London and the South East and the deliberations of the Davies Commission. The issue is in fact highlighted specifically within the Discussion Document.

5.8 This market is essential to the operation of a 'hub' airport. The function of a hub is concentrate demand from a wide geographic on a single point so an airline or airlines can consolidate demand to enable a broader range of destinations to be served at a higher frequency than would be possible based on airport's local catchment area.

5.9 The competitive market for these passengers is quite different to that for surface origin passengers in the UK. Competition is primarily beyond UK boundaries, Heathrow (the UK's only true hub) competes not with other UK airports but with overseas airports such as Amsterdam, Paris CDG, Frankfurt, Munich, Madrid, Dubai, Doha or Abu Dhabi.

5.10 Currently, the DfT model only considers a small slice of this market and only in a very simplistic way. The National Air Passenger Demand Model estimates the number of international to international interline passengers using UK airports as hubs using a relationship to overseas economic growth. The numbers at each airport simply grow in line with this forecast. When there is a constraint cost a proportion are simply priced out of the market. There is no consideration of the overall size of

the potential market or indeed how UK airports compete with overseas airports for this market. Given that the current discussions around additional airport capacity in London are focussing around maintaining hub status this is a crucial problem.

5.11 At present, the DfT model implicitly assumes that the UK (essentially Heathrow) has a 'right' to a slice of this market. This is not the case. Interline passengers are by nature footloose, they can shift routings very easily, much more easily than a surface passenger in many cases. This means that if Heathrow continues to lose destinations and become a less attractive hub there is the potential for others to pick up substantially more traffic from Heathrow than is implied by the DfT approach. Equally, the DfT approach cannot hope to replicate the potential of a rejuvenated London 'hub' in taking traffic from other hub airports. Ultimately, the decision to route a passenger through one hub or another is controlled by the airline or alliance in terms of how it prices tickets for transfer passengers in a competitive market. Hence, the real determinant of hubbing is the extent to which airlines need to attract transfer traffic to sustain their network.

5.12 The DfT model needs improvements in this area. It needs to make estimates of the total relevant market flows, perhaps using MIDT or similar data, and it needs to effectively consider the competitive dynamic with other hub airports, using a LOGIT or Quality of Service Index (QSI) type approach. This is not necessarily a simple task but it needs to be considered in more depth.

6 EVIDENCE OF ISSUES IN THE FORECASTS

6.1 We have described above a number of issues and potential problems with the DfT forecasting model. However, it is also helpful to consider what these issues mean in practical terms in relation to the outputs from the process. In the latest version of the forecasts, with total demand growing relatively slowly, up until 2030 the forecasts look relatively sensible. The 'spill' effects that highlight a lot of the problems we have described are not yet strong enough to cause major problems. However, by 2050, the irrationalities in the model are beginning to appear, suggesting that they are there earlier on but are simply are harder to spot in the published data available.

6.2 Below, we have set out a small number of examples that help to illustrate and provide evidence in relation to some of our concerns:

- ✈ Long Haul Traffic at Stansted Airport – the latest 2013 forecasts show a strange pattern in terms of long haul demand at Stansted. Unlike in previous versions of the forecasts Stansted does secure some long haul traffic by 2030 (around 2.5 mppa). This is rational. The Airport is close to London, has a strong natural catchment area (including strong long haul markets) and runway capacity and capability. However, by 2050 all this traffic has gone again. We

suspect that it has been sucked in to Birmingham instead, which has gone from having 2.6 million long haul passengers in 2030 to nearly 15 million by 2050. This does not seem rational. With a significant presence at Stansted established, presumably serving the natural catchment area, why would airlines then leave that base? The answer in the model probably relates to the weakness of Stansted's previous track record as a long haul airport and a frequency 'ballooning' effect at Birmingham;

- Growth at East Midlands – this is classic example of the model not allowing an airport to develop in its own catchment. Between 2010 and 2030, East Midlands grows by only 100,000 passengers, reaching 4.4 million in 2030. However, as 'spill' starts to roll out of London in earnest, by 2050 it has reached 14.1 million passengers. Again, however, there is a lack of rationality in the pattern. Despite having 14.1 million passengers, there is no long haul traffic. The Airport has no track record in these markets and hence it cannot sustain any services;

- Long Haul in the East of the UK – despite constraint in London and despite a number of regional airports developing significant long haul presences, notably Birmingham and Manchester, no airport in the Eastern half of the country, including Luton and Stansted (at 2050), develops any long haul traffic until as far North as Newcastle. Even then the pattern is strange. Between 2010 and 2030, Newcastle loses 100,000 long haul passengers, but by 2050 it has gained 600,000. This suggests that it is in fact a beneficiary from spill, probably from Manchester as it has filled up on the back of demand spilling from further South. Essentially, the model is just pushing demand from London until it hits an airport with capacity and a track record in long haul traffic. That airport then experiences a ballooning of frequency, which stops others from developing these services, until it itself is full and demand flows to the next airport in the chain. This is not how air services develop in practice.

6.3 Whilst it is easier to identify the patterns in the long haul category as it is separately identified, we suspect that the same errors infect individual short haul markets as well rendering detailed results spurious as congestion bites at Heathrow.

7 CONCLUSIONS

7.1 The DfT modelling is highly complex and has been developed over many years. However, there are significant issues, particularly around the functioning of the allocation model. In our view there are a number of priority areas for action:

- The approach to analysing uncertainty in the inputs to the National Air Passenger Demand Model needs overhauling, probably using a

probability approach. There will always be difficulties around input assumptions and forecasts but the effects of these can be minimised by appropriate approaches to uncertainty;

- The frequency term within the model needs re-examining to consider the 'ballooning' impact it seems to create when traffic is spilled. Fundamentally the model needs to allow airports to develop their own catchment areas appropriately;

- The treatment of international to international hub traffic needs to be substantially overhauled, particularly in the context of the terms of reference for the Davies Commission;

- There needs to be a step change in the way the forecasts are used and considered. The model is just that. It is a piece of evidence that needs to be viewed in the round. It is not sufficiently robust or accurate to be a sole basis for policy decisions. Wider evidence on growth and patterns of growth will always be vital. Even an enhanced version of the model should be viewed in this light.

- Peer review needs to include industry forecasting expertise to sense check the results to how airlines are actually likely to respond to growth and constraint.