INDEPENDENT PEER REVIEW OF ECONOMIC IMPACT OF NIGHT FLIGHTS RESEARCH STUDY

Richard Batley & Peter Mackie, ITS Leeds

15th August 2017

1. INTRODUCTION

This document reports a peer review commissioned by the Department for Transport (DfT) to assess:

- the quality of the methodology for estimating the economic impacts of potential changes to the current night flight regime on airports, airlines, passengers and the public accounts;
- whether the methodology fits with accepted best practice; and
- whether the methodology could be improved.

The review was based primarily upon the final report of the study, namely:

- Economic impacts of night flights research study, SYSTRA, 2017.

But also drew upon a number of other related documents, namely:

- National Air Passenger to Airport Allocation Model (NAPAM): an overview of the multinomial logit re-estimation, DfT internal document, 2011.
- Peer review of NAPAM, John Bates Services for DfT, 2010.
- The nested incremental logit model: theory and application to mode choice, John Bates, David Ashley & Geoff Hyman, PTRC paper, 1987.

Also, there has been exchange of e-mails and three telecons with the DfT's project officer for this peer review.

The review was conducted in two stages. First, we were issued with a draft final report on 8th April 2017; this encompassed Chapters 1 to 4, detailing the study objectives and scope, stakeholder consultation, model development and modelling approach. We spent 3.5 days peer reviewing the draft final report, and submitted our interim findings to DfT on 25th April. Second, we were issued with the finalised report on 4th July 2017; this version added Chapters 5 and 6, detailing the illustrative model results and study conclusions. We spent a further 1.5 days completing our peer review, and submitted our findings to DfT on 7th July. DfT conducted an in-house review of our peer review, and we made some minor adjustments, before re-submitting on 15th August. The layout of our peer review report is as follows:

- Section 2 comments on the methodological perspective that has been employed, focussing upon two areas as directed by DfT: first, the high-level market response to the night flights regime; and second, more specific considerations around the airport choice model and shadow fare.
- Section 3 comments on the empirical results, again focussing upon two areas: first, the general plausibility of the reported findings; and second, specific considerations around sensitivity testing of the shadow fare.
- Section 4 summarises our responses to the three key peer review questions detailed above.

2. COMMENTS ON METHODOLOGICAL APPROACH

2.1 Background comments on the methodology

The background to SYSTRA's study is clearly explained in their report, namely that the current night flight regime (NFR) – which limits number/type of aircraft permitted to fly and amount of noise that can be produced at the three principal London airports (LHR, LGW, STN) between 23:00 and 7:00 – will end in October 2017. The study was therefore undertaken *'with the purpose of improving the Government's evidence base on this issue and therefore to inform its assessment of the impacts of the night flights regime that should apply from October 2017 onwards'*.

SYSTRA's report notes that key outputs of the study were to be: *'flexible models for each airport which can be used by the Department to estimate the economic impacts on airports, airlines, passengers and the public accounts from a range of potential changes to the current night flight regime'*. It is also noted that, whilst of significant interest to the Department, *'the impacts associated with changes in noise pollution from aircraft'* fall outside the scope of the study.

SYSTRA's report further sets out the rationale for, and key principles of, the Government's approach to the NFR, as detailed in the Aviation Policy Framework (APF) 2013. The key principles encompass:

- limits on the number of flights during the night quota period (NQP), from 23:30-06:00
- restrictions on the noisiest aircraft types, from 23:00-07:00; and
- limits on the total amount of noise during the NQP

The study focusses on three specific time periods, namely:

- Entire night period; 23:00 to 07:00
- Night quota period (NQP); 23:30 – 06:00
- Shoulder periods; 23:00 to 23:30 and 06:00 to 07:00

The economic impacts considered to be within the scope of the study are first order effects of NFR experienced by:

- Airports
- Passenger airlines
- Cargo airlines
- Passengers
- Public accounts

The study also seeks to capture any 'behavioural change' resulting from the aforementioned changes, as well as 'knock-on' effects between the five groups – these might be referred to as second order effects.

It is further noted that GDP and employment impacts (i.e. wider economic benefits) fall outside the scope.

In seeking to deliver the key outputs of the study (as noted above), SYSTRA develop a comprehensive modelling framework which is illustrated in Figure 2 of their study report, and is reproduced below.

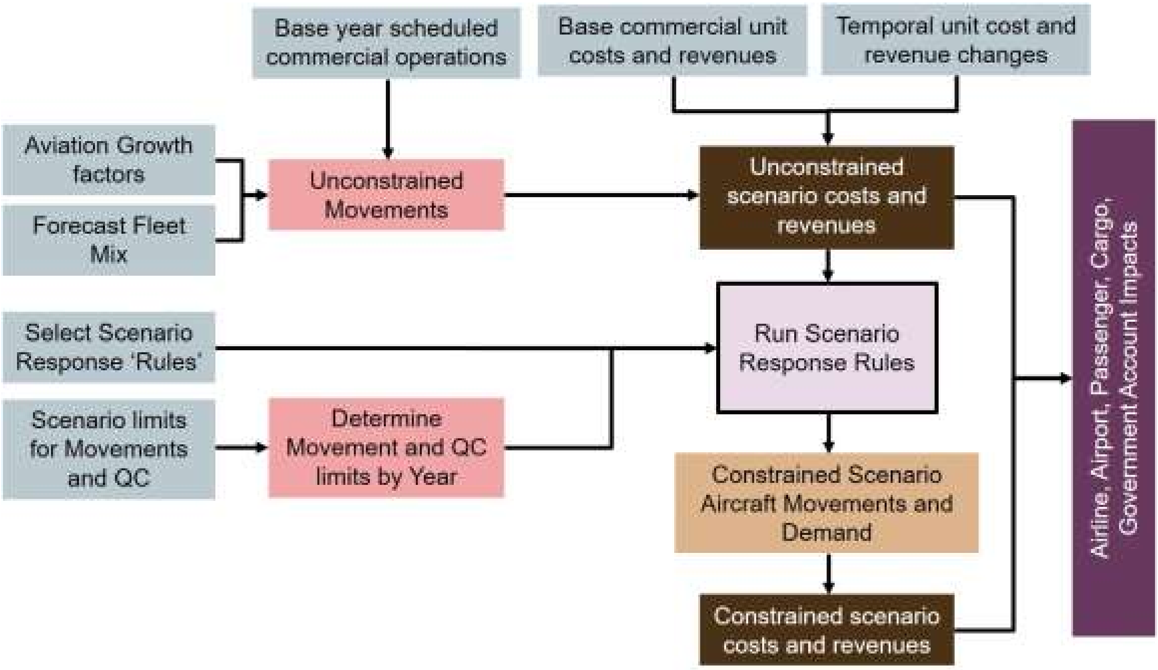

The key features of the modelling framework include the following:

- The foundation of the models is the actual base year (2014/15) commercial aircraft operations to which base year commercial unit revenues and costs are added (this covers airline costs and revenues, airport costs and revenues, and APD and VAT).
- A set of unconstrained[1] aircraft movements in future years are calculated by applying the aviation growth factors and the assumed future fleet mix to the base year operations.
- Unconstrained scenario revenues and costs in future years are calculated as a reference point for the impact analysis and a basis for computing the revenues and costs for the night flights policy scenario under consideration.

[1] By the way, the definition of the unconstrained scenario is not transparent in the text and requires careful reading; at first, we could not understand the results for Heathrow. It should be made clear that it is not truly unconstrained, which would be where all airports had 24 hour operation and passengers and airlines made their choices within that. There would be no maximum flight numbers (other than capacity restrictions) or any other quantitative restriction. In that scenario, Heathrow would have a lot more traffic than it currently does, some diverted, some generated, and the time of day pattern would be different. If that was the unconstrained case there would be very large numbers in the tables for the constrained baseline. The fact that the numbers for Heathrow are so small suggests that this cannot be what DfT and SYSTRA are viewing as the 'unconstrained scenario'.

- The commercial response rules are then 'run' where there are constraints on the total operations and/or QC under the night flights regime policy scenario under consideration compared to the unconstrained scenario.
- The models then generate a constrained scenario for aircraft operations, demand, costs and revenues. That is, the total operations, demand, costs and revenues under the night flights regime policy scenario being considered.
- The impact of each night flights regime policy scenario (including the baseline) is computed relative to the unconstrained scenario.
- For future years, the models estimate the total revenues and costs for each flight grouping, covering airline costs and revenues, airport costs and revenues, and APD and VAT, based on the number of flights estimated in the flight groupings, and the average revenues and costs per flight assumed for the flight groupings.

2.2 Observations on the high-level market response to the NFR

2.2.1 Overarching comments

Our understanding is that the principal application of this work is to flights of 2-4 hours duration, predominantly leisure in nature. However there are some long distance overnight flights in the mix. If that is correct, we would suggest that the Department uses some simple back-of-envelope sense-checks on the orders of magnitude of the results.

Consider a very simple pattern in which departures and arrivals occur in a rectangular distribution over time from 0600 to 2200 in the reference case and the policy test is to change the hours of operation to 0700 to 2200. This will lead to a complete rescheduling of all the rotations but in the case cited above the productivity of the system will fall by one-sixteenth. Applying that to the prices, quantities and marginal costs in the system would provide a back-of-envelope calculation of the limiting case, excluding consumer surplus. Variations on that theme could be used if the departure/arrival rate in the 0600-0700 time period were less than one-sixteenth. We recommend that the Department establishes some simple benchmarks against which to judge the model outputs. This approach is only reasonable for the short haul market where several rotations per day are being achieved and the schedule can be re-optimised in response to constraints. This approach has both the virtue and limitation of simplicity. Obviously it is not a satisfactory approach for long haul.

2.2.2 Assessing the SYSTRA work

Turning to the SYSTRA work itself, there are essentially three elements by which one would want to judge:

- the inputs, including the scenarios to be tested;
- the model properties and consistency with NAPAM;
- the model outputs and their credibility.

Given that several scenarios are tested, an important thing to do is to establish the credibility of the pattern of the results. As the John Bates (2010) peer review of NAPAM pointed out, the understanding of the choice 'travel versus not travel' is weak relative to other choices, so the handling and results for that response are of particular interest relative to retiming and reassignment.

2.3 Observations on the airport choice model and shadow fare

2.3.1 Overarching comments

Within SYSTRA's modelling framework, there would seem to be a series of sub-models, one of which is a demand/passenger response model. SYSTRA note under point 4.3.3 of their report that: *'The implementation approach for this draws on the DfT National Passenger Allocation Model (NAPAM) coefficients to provide a model and coefficients to allow the representation of the impacts of commercial responses. NAPAM is an aggregate multinomial logit model which distributes passengers across a set of available route options based on their relative costs. Within this research study the model has been applied in incremental form – that is to adjust from the unconstrained situation'*.

Turning to NAPAM itself, the function of this model is to input demand forecasts for air passenger travel between zones, and allocate this demand to UK airports, subject to capacity constraints. The model comprises three components:

- An airport choice model which allocates passenger demand to airports on the basis of access costs and service levels.
- An air traffic movement (ATM) model which converts passenger demand into ATMs (i.e. routes) on the basis of aircraft mix.
- A capacity restraint which generates a 'shadow cost' for any given route on the basis of runway and terminal capacity; these costs are fed back to the airport choice model and, in due course, to the forecasts of air passenger travel between zones.

Within the time constraints of our peer review of SYSTRA's final report, we have undertaken an expedient but considered assessment of the airport choice model and the method used to calculate the shadow fare, and this reveals three key points. Each of these points is considered in the sub-sections that follow, but can be summarised as follows.

- SYSTRA adopted an 'incremental' version of the airport choice model, which was advocated by Bates (2010), but entails a material departure from the 'absolute' version of the airport choice model estimated in the context of NAPAM. The idea of using an incremental model has its attractions, but in adopting this model, SYSTRA's practice of 'recycling' parameter estimates from an absolute model could be problematic. In other words, this practice will effectively introduce bias into the parameter estimates – it is difficult for us to advise on the scale of the bias, since this will be an empirical issue.
- Subject to minor caveat, we have confirmed that SYSTRA's formula for calculating the shadow fare (noted under point 4.3.12 of their report) is correct.
- In applying the formula for shadow fare, there is a need to convert from utility to money units, but the NAPAM parameters offer a range of potential numeraire, and different numeraire will give different answers. For reasons of robustness, there is a particular question mark over whether the fare attribute should be included in the list of candidates.

2.3.2 Interchangeability between absolute and incremental versions of the NAPAM discrete choice model

At the heart of NAPAM, there is a discrete choice model with utilities specified in 'absolute form'. However, a well-known property of discrete choice models is that of 'translational invariance' (e.g. Train, 2009, p19[2]), meaning that a common constant can be added to the utilities of all alternatives without empirically changing the model. This property means that any specification in absolute form can be readily translated into a corresponding specification in 'difference form' without empirically changing the model. In other words, these two specifications are not identifiably different, all else equal.

To illustrate, using a simple two-alternative example for the context of a value of travel time (VoT) problem, define[3]:

$$U_{1,n} = \lambda \cdot \left(\alpha \cdot T_1 + \beta \cdot C_1 + \varepsilon_{1,n} \right) \quad (1)$$

$$U_{2,n} = \lambda \cdot \left(\alpha \cdot T_2 + \beta \cdot C_2 + \varepsilon_{2,n} \right) \quad (2)$$

Where:

$U_{i,n}$ is the utility of alternative $i = 1,2$ for individual $n = 1,...,N$;

T_i is the travel time of alternative $i = 1,2$;

C_i is the travel cost of alternative $i = 1,2$;

α and β are parameters representing the marginal utility of travel time and cost respectively, where it would be expected that $\alpha, \beta < 0$, since both travel time and cost are generally considered to be 'bads';

λ is a parameter which sets the scale of utility, where $\lambda > 0$ by definition;

$\varepsilon_{i,n}$ is an independent and identically-distributed random term, associated with the utility of alternative $i = 1,2$ for individual $n = 1,...,N$.

Assuming that the random term were distributed IID Gumbel, the logit probability of choosing alternative 1 would be given by:

$$P_{1,n} = \frac{\exp\left(U_{1,n}\right)}{\exp\left(U_{1,n}\right) + \exp\left(U_{2,n}\right)} \quad (3)$$

Now, on the basis that *'only differences in utility matter'* (e.g. Train, 2009, p19), $U_{1,n}$ can be subtracted from (1) and (2) to give:

---

[2] Train, K. (2009) Discrete Choice Methods with Simulation. Cambridge University Press, Cambridge.

[3] Of course, the NAPAM model considers a choice set comprising multiple alternatives, but the two-alternative example will be adequate to illustrate the conceptual issues of interest.

$$\Delta U_{1,n} = 0 \tag{4}$$

$$\Delta U_{2,n} = \lambda \cdot \left(\alpha \cdot \Delta T + \beta \cdot \Delta C + \Delta \varepsilon_n \right) \tag{5}$$

Where:

$$\Delta U_{1,n} = U_{1,n} - U_{1,n}$$

$$\Delta U_{2,n} = U_{2,n} - U_{1,n}$$

$$\Delta T = T_2 - T_1$$

$$\Delta C = C_2 - C_1$$

$$\Delta \varepsilon_n = \varepsilon_{2,n} - \varepsilon_{1,n}$$

In this way, we translate the 'absolute form' for (1) and (2) into their respective 'difference form' equivalents (4) and (5).

Assuming again that the random term were distributed IID Gumbel, the logit probability of choosing alternative 1 would now be given by:

$$P_{1,n} = \frac{\exp(0)}{\exp(0) + \exp\left(\Delta U_{2,n}\right)} = \frac{1}{1 + \exp\left(\Delta U_{2,n}\right)} \tag{6}$$

Despite this reformulation, the probabilities (3) and (6) will be empirically equivalent.

In passing, note that if alternative 2 were slower and cheaper than alternative 1, it must be the case that $\Delta T > 0$ and $\Delta C < 0$. With reference to the random term in (5), the following identity holds by definition:

$$\text{var}\left[\lambda \cdot \Delta \varepsilon_n\right] = \lambda^2 \cdot \text{var}\left(\Delta \varepsilon_n\right) \tag{7}$$

Taking the square root:

$$\text{s.d.}\left[\lambda \cdot \Delta \varepsilon_n\right] = \lambda \cdot \text{s.d.}\left(\Delta \varepsilon_n\right) \tag{8}$$

Rearranging:

$$\lambda = \frac{\text{s.d.}[\cdot]}{\text{s.d.}(\cdot)} \tag{9}$$

In combination with (5), (9) shows that the λ parameter is associated with the distribution of the random term.

Perhaps with the intention of following Bates' advice (2010, p27), SYSTRA has seemingly sought to exploit the aforementioned 'translational invariance' property, by re-specifying the discrete choice model in 'incremental form', meaning that utility differences are specified in terms of differences between Do Something (DS) and Do Nothing (DN) scenarios (as opposed to differences between alternatives under any one scenario, as in (4) and (5) above).

However, a defining feature of incremental logit is that the probability statement re-weights the (exponents of the) utilities in the following manner:

$$P_{1,n} = \frac{D_1 \cdot \exp(0)}{D_1 \cdot \exp(0) + D_2 \cdot \exp(\lambda \cdot \Delta U_{2,n})} = \frac{D_1}{D_1 + D_2 \cdot \exp(\lambda \cdot \Delta U_{2,n})} \qquad (10)$$

Where D_i is the absolute demand for alternative $i = 1,2$ under the DN. The implication of (10) is that, unless $D_1 = D_2$, it will no longer be equivalent to (6).

Therefore, contrary to the practice which seems to have been employed by SYSTRA, it will not in general be valid to interchange parameters between 'absolute' (as used in the NAPAM re-estimation report) and 'incremental' (as used in the SYSTRA report) versions of the discrete choice model.

In other words, given the incremental formulation of SYSTRA's model, it is important that the model parameters (i.e. the β) are re-calibrated to the absolute demands (i.e. the D_i).

2.3.3 Deriving the shadow fare from an incremental version of the NAPAM discrete choice model

Corresponding to (10) from incremental logit, the demand (as opposed to choice probability) for alternative 1 under the DS is given by:

$$D_1' = D_1 \cdot \frac{D_1}{D_1 + D_2 \cdot \exp(\lambda \cdot \Delta U_2)} \qquad (11)$$

Where the DS in this case entails some change for alternative 1 but no change for alternative 2 (and we now simplify notation by dispensing with the *n* index).

Again perhaps following Bates' advice (2010, p36) concerning the arbitrariness of the DN for alternative 1, SYSTRA appear to have deferred to the approximation[4]:

$$D_i' = D_i \cdot \exp(\lambda \cdot \Delta U_i) \text{ for } i = 1,2 \qquad (12)$$

Where:

ΔU_i is the difference in utilities between the DN (i.e. unconstrained) and DS (i.e. constrained) scenarios

λ is the scale factor, which for consistency with logit must now fall within the range $0 < \lambda \leq 1$.

Rearranging:

$$\frac{D_i'}{D_i} = \exp(\lambda \cdot \Delta U_i) \qquad (13)$$

[4] Although in the present case, there is in practice only a utility change for alternative 2.

Taking logs:

$$\ln\left(\frac{D_i'}{D_i}\right) = \lambda \cdot \Delta U_i \tag{14}$$

Within the SYSTRA report, ΔU_i is specified using a so-called 'individual' formulation, whereby a separate parameter is specified for each of several components of generalised cost, as follows:

$$\Delta U_i = \beta_{access} \cdot \Delta U_{access,i} + \beta_{fare} \cdot \Delta U_{fare,i} + \delta_{retime}\beta_{wait} \cdot \Delta U_{wait,i} + \beta_{transfer} \cdot \Delta U_{transfer,i} \tag{15}$$

Where:

$\Delta U_{access,i}$ is the difference in access costs between the DN and DS for alternative $i = 1,2$, and so on for the other attributes

β_{access} is the marginal *utility* of the difference in access costs, and so on for the other attributes

δ_{retime} is the retiming adjustment (applied to wait time in the event of schedule displacement)

Given the formulation (15), it is important to make the distinction that[5]:

- $\Delta U_{access,i}$ is defined in money units, and so on for the other attributes, whilst
- ΔU_i, having applied the relevant marginal utility for each attribute, is defined in utility units (or 'utils').

Now substituting for ΔU_i in (14) using (15):

$$\ln\left(\frac{D_i'}{D_i}\right) = \lambda \cdot \left[\beta_{access} \cdot \Delta U_{access,i} + \beta_{fare} \cdot \Delta U_{fare,i} + \delta_{retime}\beta_{wait} \cdot \Delta U_{wait,i} + \beta_{transfer} \cdot \Delta U_{transfer,i}\right] \tag{16}$$

In an *estimation* context (such as the NAPAM re-estimation study[6]), the scale factor λ will not however be identifiable, since it will be confounded with the various β parameters. That is to say, the following model will in practice be identified:

$$\ln\left(\frac{D_i'}{D_i}\right) = \hat{\beta}_{access} \cdot \Delta U_{access,i} + \hat{\beta}_{fare} \cdot \Delta U_{fare,i} + \delta_{retime}\hat{\beta}_{wait} \cdot \Delta U_{wait,i} + \hat{\beta}_{transfer} \cdot \Delta U_{transfer,i} \tag{17}$$

Where $\hat{\beta}_{access} = \lambda \cdot \beta_{access}$, and so on for the other attributes

If however – and without materially changing (17) – we divide both sides by the marginal utility of fare (or some other cost-based attribute – see discussion to follow), we get:

$$\frac{\ln\left(D_i'/D_i\right)}{\hat{\beta}_{fare}} = \frac{1}{\hat{\beta}_{fare}} \cdot \left[\hat{\beta}_{access} \cdot \Delta U_{access,i} + \hat{\beta}_{fare} \cdot \Delta U_{fare,i} + \delta_{retime}\hat{\beta}_{wait} \cdot \Delta U_{wait,i} + \hat{\beta}_{transfer} \cdot \Delta U_{transfer,i}\right] \tag{18}$$

[5] There is a similar discussion on page 5 of Bates (2010).

[6] Whilst the NAPAM re-estimation study employed an absolute version of the discrete choice model, essentially the same problem will apply.

Which simplifies to:

$$\frac{\ln\left(D_i'/D_i\right)}{\hat{\beta}_{fare}} = \frac{1}{\beta_{fare}} \cdot \left[\beta_{access} \cdot \Delta U_{access,i} + \beta_{fare} \cdot \Delta U_{fare,i} + \delta_{retime}\beta_{wait} \cdot \Delta U_{wait,i} + \beta_{transfer} \cdot \Delta U_{transfer,i}\right] \qquad (19)$$

Simplifying further, (20) gives a measure of the shadow fare – in money units – free of the λ scale[7]:

$$\frac{\ln\left(D_i'/D_i\right)}{\hat{\beta}_{fare}} = \frac{\Delta U_i}{\beta_{fare}} \qquad (20)$$

This derivation confirms that, subject to the caveat that the demand statement (12) is an approximation to (11), SYSTRA's formula for calculating the shadow fare (noted under point 4.3.12 of their report) is correct. However, following section 1, in order for (20) to be empirically meaningful, it is important that the model parameters (i.e. the β) are calibrated for the absolute demands (i.e. the D_i).

2.3.4 Choice of numeraire

The preceding discussion has made the important distinction between measurements in money or utility units, and thus highlighted the crucial role played by the marginal utility of (the difference in) money/cost in translating from one dimension to the other.

However, a further point made by Bates (2010) is that the 'individual' formulation (15) potentially yields four variations on a theme – in terms of different notions of the marginal utility of money or cost. In fact, Bates expressed concerns regarding the robustness of the fare parameter, recommending that this should be removed from the econometric models (p26 of his review). For purposes of estimating the shadow fare, Bates instead recommended the surface access costs parameter (p42), since this is naturally quantified in money units (as opposed to a conversion from time units via the VoT) and was estimated robustly in the course of the NAPAM re-estimation.

To illustrate the potential impacts of this change using SYSTRA's assumed parameters (point 4.3.5 of their report)[8], the tables below compare the range of shadow fares using fare as the numeraire vs. using surface access costs, for ratios of demand (i.e. DS demand/DN demand) ranging from 0.95 to 0.50.

In the first table we use fare as the numeraire (i.e following (20)), and elicit shadow fares in the range £3.40-£45.90 for leisure and £114.24-£1543.76 for business.

---

[7] Although not used by SYSTRA, it is worth noting the corresponding measure of the shadow fare deriving from the full demand specification (11), which in money units would be given by:

$$\ln\left(\frac{D_1' \cdot D_1 + D_1^2}{D_1' \cdot D_2}\right) \Big/ \beta_{fare} = \Delta U_2$$

[8] In passing, we can confirm that the model parameters tabulated under point 4.3.5 of the SYSTRA report are correctly transposed from the recommended models on p17 and p22 of Bates' review (2010) of NAPAM.

In the second table we instead use surface access costs as the numeraire, and elicit shadow fares in the range £0.38-£5.17 for leisure and £1.69-£22.88 for business.

Thus the shadow fares are markedly different, depending on which numeraire is adopted.

D'/D	betafare_leis	betafare_biz	shadow_leis	shadow_biz
0.95	-0.0151	-0.00045	3.396907	114.239
0.9	-0.0151	-0.00045	6.977518	234.6559
0.85	-0.0151	-0.00045	10.76284	361.9575
0.8	-0.0151	-0.00045	14.77772	496.979
0.75	-0.0151	-0.00045	19.05179	640.7173
0.7	-0.0151	-0.00045	23.62086	794.3763
0.65	-0.0151	-0.00045	28.52867	959.4274
0.6	-0.0151	-0.00045	33.82951	1137.696
0.55	-0.0151	-0.00045	39.59185	1331.486
0.5	-0.0151	-0.00045	45.90379	1543.758
D'/D	betasac_leis	betasac_biz	shadow_leis	shadow_biz
0.95	-0.134	-0.0303	0.382786	1.692848
0.9	-0.134	-0.0303	0.786273	3.477245
0.85	-0.134	-0.0303	1.212828	5.363661
0.8	-0.134	-0.0303	1.66525	7.364474
0.75	-0.134	-0.0303	2.146881	9.494458
0.7	-0.134	-0.0303	2.661753	11.77145
0.65	-0.134	-0.0303	3.214798	14.21726
0.6	-0.134	-0.0303	3.812132	16.85893
0.55	-0.134	-0.0303	4.46147	19.73059
0.5	-0.134	-0.0303	5.17274	22.87614

These results illustrate that there needs to be further consideration of which numeraire is the most defensible – but our inclination is to avoid using fare.

3. COMMENTS ON EMPIRICAL RESULTS

3.1 Some remarks on the plausibility of the results

SYSTRA's model results refer to the testing of several DS scenarios vis-à-vis the DN scenario of the existing NQP regime. The report openly concedes that the scenarios are designed simply to illustrate the workings of the model and have no grounding in policy options being actively considered by DfT. We will not therefore comment upon the design of the scenarios, and instead focus upon the properties of the model as revealed by this testing exercise.

First and foremost, we would remark that – from the information given in Chapter 5 – we found it difficult to assess the plausibility of the results. It would be better to introduce the big tables (i.e. Table 5.5 onwards) by showing a sample calculation for one flight grouping at one airport at the per passenger level. So we would see base average fare, base MC per passenger (airport and airline), and then the fare and any relocation costs per passenger and any variation in the MC. This would begin to make Chapter 5 more accessible to the reader. This issue is compounded by the sensitivity testing, which focuses upon a single crude example (see the additional comments in the following sub-section).

Referring to Table 5.5 for example, we think that at least some discussion of the patterns in the results is needed. Does SYSTRA have comments on plausibility? For example, the results from Flights Lost, Retiming with Displacement and Flights via non UK hub are close together. Is this plausible? Also, what exactly is the base number and the percentages in the table? Some description is needed. Also, retiming to shoulder periods produces a massive PS benefit, so would airlines be operating night flights if this scenario was really feasible?

Another point about these tables concerns the labelling. We believe that the main results concern the performance of the baseline case relative to the unconstrained case *if* the whole of the response is taken by not travel *or* retiming *or* reassignment etc. This needs to be made clear. Also, of course, there is nothing within this work which enables us to forecast which of these responses will actually apply, or more realistically which combination. That places further stress on the correctness of several responses having total economic value change effects which are similar in magnitude, which is not what we were expecting ex ante.

Turning to Table 5.6, our reading is that the increase in the shadow price is around £25 per passenger. Is that reasonable (for Stansted)? More generally, the write-up is not terribly generous to the reader, so for example what is the base fare against which that £25 should be compared? And how sensitive is the shadow fare to the choice of numeraire?

Finally, the consultants need to put in a statement that these tables refer to a 'CBA of all users and operators at UK airports impact' and not to a 'UK users and operators' CBA. We do not have a problem with this since it follows the Airport Commission's approach, but it needs saying.

3.2 Some reflections on the implications for model specification

Our reading of the results also gives us cause to reflect back upon various aspects of model assumptions and properties, as follows.

- With reference to point 4.3.9 in the SYSTRA report, '*A key assumption here is that passengers prefer flying in the NQP and this may not always be the case*'. We interpret this assumption as meaning that there is essentially a uniform fare within flight groupings with passengers allocating themselves to time of day according to preferences. So then under the scenario tests, passengers suffer an increase in the shadow fare *plus* any 'relocation' costs to outside the NQP. Our instinct is that, while this might be reasonable for long haul scheduled, for short haul trips of two or three hours duration another model is likely to be at work. This world is more likely to be one of non-uniform pricing (the fare from Stansted to Ibiza is on average substantially lower for the 2am flight than the 2pm flight). People are allocating themselves to flights on the basis of their generalised costs. So adding the shadow fare change and the relocation costs would be wrong if this is the correct model. This obviously significantly affects the results. A related point is that, as we understand it, the NAPAM elasticities are not time of day differentiated. That seems a weakness when it comes to analysing time of day specific policy scenarios.

- Again with reference to point 4.3.9, the report notes that: *'Standard values of travel time for business and leisure passengers from WebTAG are used in the models'*. Our appreciation is that there has never (or not for many years) been a study of VoT for air travel whether business or leisure, UK or non-UK, long haul or short haul. Therefore DfT is not in a good place concerning this value set.
- With reference to point 4.3.13, it is unclear why the fare-capping assumption has been made. Some explanation is desirable of what is going on and whether this is significant or just dealing with some outliers.

3.3 Sensitivity testing

Section 5.6 of SYSTRA's final report details a single sensitivity test whereby the default parameters from the NAPAM model are *'reduced in sensitivity by factoring by 0.8'*. More specifically, the 'default' and 'test' parameters are detailed in the following table, which is reproduced from the report.

PARAMETER	DEFINITION	BUSINESS VALUE	LEISURE VALUE
Surface Access Cost (β_{access})	VOT * Additional surface access travel time	-0.0303	-0.134
Air Fare (β_{fare})		-0.000449	-0.0151
Cost of (Direct) Wait Time (β_{wait})	VOT * Displacement of travel time to the shoulder periods	-0.09	-0.103
Cost of Transfer Time ($\beta_{transfer}$)	VOT * Additional transfer time	-0.0503	-0.187
A	Coefficient in the wait time formulation	0.2	0.4
Re-timing adjustment ($\delta_{re\text{-}time}$)	Derived parameter applied to wait time per hour of schedule displacement	0.0625	0.0625

Reasoning this test, SYSTRA comment that: *'This factoring implies that the assumptions regarding the time penalties for travelling at different times, to different UK airports or via non-UK hubs have a smaller impact, but that a larger shadow fare would be required to price off passengers in response to a constraint on operations'*. The commentary here is on the terse side and it took some effort to work through the implications. With reference to (20), if we apply SYSTRA's 'factoring' to our earlier calculation of the shadow fares – again considering fare and surface access costs as alternative numeraire – then we find that the effect of the factoring is simply to increase the shadow fare by 25% across the board (i.e. 1.25=1/0.8, irrespective of purpose or the proportionate demand reduction between the DS and DN).

D'/D	betafare_leis	betafare_biz	shadow_leis	shadow_biz	test/default_leis	test/default_biz
0.95	-0.01208	-0.0003592	4.24613364	142.798704	1.25	1.25
0.9	-0.01208	-0.0003592	8.72189699	293.319921	1.25	1.25
0.85	-0.01208	-0.0003592	13.4535538	452.446908	1.25	1.25
0.8	-0.01208	-0.0003592	18.4721483	621.223695	1.25	1.25
0.75	-0.01208	-0.0003592	23.8147411	800.896638	1.25	1.25
0.7	-0.01208	-0.0003592	29.5260715	992.970334	1.25	1.25
0.65	-0.01208	-0.0003592	35.6608374	1199.28429	1.25	1.25
0.6	-0.01208	-0.0003592	42.2868894	1422.12033	1.25	1.25
0.55	-0.01208	-0.0003592	49.4898179	1664.35691	1.25	1.25
0.5	-0.01208	-0.0003592	57.3797335	1929.69705	1.25	1.25
D'/D	betasac_leis	betasac_biz	shadow_leis	shadow_biz	test/default_leis	test/default_biz
0.95	-0.1072	-0.02424	0.47848222	2.11606	1.25	1.25
0.9	-0.1072	-0.02424	0.98284063	4.34655593	1.25	1.25
0.85	-0.1072	-0.02424	1.51603479	6.7045763	1.25	1.25
0.8	-0.1072	-0.02424	2.08156298	9.20559205	1.25	1.25
0.75	-0.1072	-0.02424	2.68360142	11.8680723	1.25	1.25
0.7	-0.1072	-0.02424	3.32719164	14.7143129	1.25	1.25
0.65	-0.1072	-0.02424	4.01849735	17.7715724	1.25	1.25
0.6	-0.1072	-0.02424	4.7651644	21.0736643	1.25	1.25
0.55	-0.1072	-0.02424	5.57683769	24.6632426	1.25	1.25
0.5	-0.1072	-0.02424	6.46592519	28.5951807	1.25	1.25

SYSTRA appear to draw some comfort from this test, especially since it provokes: *'Only small (only up to around 1%) differences in the total monetary impacts and on the numbers of passengers (and aircraft operations) overall'* vis-à-vis the default parameters. However, it is worth making a few remarks in response to this finding:

- This is only one, and somewhat crude, sensitivity test, which does not fully demonstrate the sensitivity of the model results to the shadow fare.
- Referring back to the earlier discussion concerning the choice of numeraire, we note that this sensitivity test serves to amplify the absolute discrepancy between the shadow fares that arise from using surface access cost rather than fare in the denominator.
- Whilst highlighting the aforementioned weaknesses, we support SYSTRA's focus on the shadow fare as potentially a key area of sensitivity in the model.

4. RESPONSES TO PEER REVIEW QUESTIONS

Finally, returning to the questions set out at the beginning of this note, we conclude by making the following responses.

- **The quality of the methodology for estimating the economic impacts of potential changes to the current night flight regime on airports, airlines, passengers and the public accounts**

Overall, we think SYSTRA have been asked a very difficult question which they have made a heroic attempt to answer. In some areas, particularly the supply side, we feel comfortable with the approach and the quality of the work done, although our expertise on the details is strictly limited. On the demand side, we have made a number of points in the text above and our main issue is that it is difficult to judge their significance for the results. There has been no interaction between ourselves and SYSTRA on the content above and it is possible that there are some misconceptions on our part. But we think DfT should consider our points. On the appraisal side, we think the report is consistent with the approach taken by the Airports Commission in its work on the airport and airline operators and air travellers. The tone of the report, for example the list of limitations in section 3.2, is appropriate.

- **Whether the methodology fits with accepted best practice**

It would be better to think of this report as operating close to the frontiers of knowledge and helping to form good practice. There are few benchmarks with which to compare it. For that reason alone we would urge caution in placing great weight on the results in policy formulation work.

- **Whether the methodology could be improved**

Given the current state of knowledge, and putting to one side particular points we have made, it would be difficult to improve the methodology in the short term. Probably at stake here is the suitability of NAPAM as a tool for appraising the demand effects of interventions in particular sub-markets such as night flights. We suspect a non-uniform demand model would be more appropriate. Also the relevance of standard WebTAG values to air travellers is a significant unknown. These are non-trivial issues which could not be resolved quickly. Having said that, the write up of Chapter 5 could be improved in particular respects to make it more user friendly and Chapter 6 (Summary and Conclusions) is weak for a piece of work of this length and potential significance.

Ends.

The Department for Transport has actively considered the needs of blind and partially sighted people in accessing this document. The text will be made available in full on the Department's website. The text may be freely downloaded and translated by individuals or organisations for conversion into other accessible formats. If you have other needs in this regard please contact the Department.

Department for Transport

Great Minster House

33 Horseferry Road

London SW1P 4DR

Telephone 0300 330 3000

General enquiries https://forms.dft.gov.uk

Website www.gov.uk/dft