



UNIVERSITY OF LEEDS

Independent Peer Review of Recent Research on the Existence of Scarcity Rents at Heathrow

FINAL REPORT

Richard Batley, Phill Wheat & Peter Mackie

22nd August 2019

Study Funded by the Civil Aviation Authority (CAA)

ITS

Independent Peer Review of Recent Research on the Existence of Scarcity Rents at Heathrow

FINAL REPORT

Contents

1	NON-TECHNICAL SUMMARY	5
1.1	What are scarcity rents?	5
1.2	Why do scarcity rents matter?.....	5
1.3	Is scarcity the only source of rent to the airport/airline system?	5
1.4	Are scarcity rents directly observable?.....	5
1.5	How might scarcity rents be observed?.....	6
1.6	Does economic theory suggest there could be a scarcity rent at Heathrow?.....	6
1.7	To whom does any such scarcity rent accrue?	7
1.8	Does the empirical evidence suggest there could be a scarcity rent? How reliable is this evidence?.....	8
1.9	What would happen to the rent if capacity were increased?	8
1.10	Recommendations	8
2	INTRODUCTION.....	10
3	GIVEN THE CURRENT REGULATORY REGIME IN THE UK, DOES THEORY SUGGEST THAT THERE COULD BE A SCARCITY RENT AT HEATHROW?.....	15
3.1	Comments on the theoretical framework presented by Frontier and FTI	15
3.2	Additional insights from the Cournot model	20
3.2.1	Theoretical rationale.....	21
3.2.2	Diagrammatic illustration	21
3.2.3	Conclusions from the Cournot model.....	23
3.3	Additional insights from oligopoly theory	23
3.3.1	Theoretical rationale.....	23
3.3.2	Empirical illustration	24
3.3.3	Conclusions from oligopoly theory	27

3.4	What would happen to rents if traffic growth were to continue but capacity not increased?	28
4	DOES THE EMPIRICAL EVIDENCE SUGGEST THAT THERE COULD BE A SCARCITY RENT? – INSIGHTS FROM FARES ANALYSIS	29
4.1	Introductory comments	29
4.2	Comparative assessment of the Frontier and SEO/Cranfield econometric analyses	30
4.3	The fundamental limitation of the ‘fares’ data – it is actually average revenue per route	31
4.4	Distinctions between modelling approaches.....	33
4.4.1	Approach to quantifying scarcity rent: dummy variable approach versus explicit measure of capacity	33
4.4.2	What does the measured ‘scarcity rent’ in each approach capture?	35
4.5	Cross-sectional data and panel data	38
5	TO WHOM DOES ANY SUCH RENT ACCRUE?	39
5.1	Background	39
5.2	Rent capture by the airline	40
5.3	Rent capture by the airport and government.....	40
5.4	Synthesis	41
6	WHAT CAN BE DONE TO STRENGTHEN ANALYSIS OF SCARCITY RENT AT HEATHROW?	42
6.1	Background	42
6.1.1	Analysis of secondary slot market	42
6.1.2	Analysis of D/S interaction.....	43
6.2	Strengthening the triangulation approach	43
6.2.1	Secondary slot market	43
6.2.2	D/S interaction	44
6.3	Strengthening the econometric approach.....	44
6.3.1	What should a capacity index capture?	45
6.3.2	What variables should be included in the model?	46
6.4	An alternative approach: analysing generalised cost instead of fares	47
7	SYNTHESIS AND RECOMMENDATIONS.....	48
7.1	Introduction	48
7.2	Theoretical insights	48
7.3	Empirical insights	49
7.4	Recommendations	50

ANNEX A: RESPONSES TO THE FRONTIER, FTI AND IAG CRITIQUE 51

 A1 Given the current regulatory regime in the UK, does theory suggest that there could be a scarcity rent at Heathrow? 51

 A2 Does the empirical evidence suggest that there could be a scarcity rent? – insights from fares analysis..... 56

 A3 To whom does any such rent accrue? 60

 A4 What would happen to the rent if capacity were increased? 62

ANNEX B: DERIVING THE LERNER INDEX AND SCARCITY RENT 65

1 NON-TECHNICAL SUMMARY

This report documents a review on the topic of scarcity rents at Heathrow. The background is that several reports and papers have been written by and for various parties in which there is not a meeting of minds on this subject, and the CAA has therefore sought an independent opinion of this body of work. The review is organised around a series of questions, and cites various academic papers and consultants reports relevant to the arguments.

1.1 What are scarcity rents?

Scarcity rents are excess profits caused by capacity limits, typically slot restrictions at airports. Consider a situation in which capacity utilisation at an airport serving a city region is close to its limits and demand is growing over time. Assume that there are no close substitutes for the airport. Ultimately, as capacity limits are approached, demand will need to be rationed to available capacity. This rationing could be by quantity, but in Western economies will be by price. The scarcity rent is the excess of the market clearing price over the uncongested price. Who receives this rent depends on the regulatory regime facing the various parties such as the airport, the airlines and the Government, including who has the property rights to the slots.

1.2 Why do scarcity rents matter?

They matter for several reasons. The regulator (CAA) is charged with protecting the interests of consumers, so if prices rise sharply in real terms because of scarcity, this affects their remit. Government, in considering the case for new capacity, needs to estimate the balance of effects on prices of the reduction or elimination of scarcity rents versus any increase in aero charges required in order to fund the capacity enhancement.

1.3 Is scarcity the only source of rent to the airport/airline system?

No. The fundamental source of rent is market power. There are economic and historical reasons why hubs are an efficient way of creating connectivity through interchange. This is likely to be associated with small group competition in particular markets possibly with entry barriers. Marketing phenomena such as frequent flyer points, travel agent commissions, ownership of slots and gates are supporting factors. But slot scarcity accentuates market power, operating in conjunction with the other market factors, often referred to as hub dominance. Given that different sub-markets have different elasticities, slot scarcity is likely to push the shape of what is provided at a hub in the direction of the least elastic traffics with the highest mark-ups. A corollary is that if new capacity is brought on-stream, the shape of the service provision at the hub will change.

1.4 Are scarcity rents directly observable?

No. It follows from the above that scarcity rents are an element in an overall economic system and are not directly observable. But that does not mean they do not exist; it means that they need to be observed indirectly, and that precision is problematic. But that is true of much modelling and forecasting within the transportation sector.

1.5 How might scarcity rents be observed?

There seem to be four possible avenues. The first is to observe what is happening in the secondary slot market and convert that into an equivalent price per seat. What seems clear is that there are very different values by time of day/week, so at the very least it is not credible to say that there is no scarcity value of a Mon-Fri morning arrival/departure slot pair at Heathrow. However, the secondary market is relatively thin and whether the reported trade values are representative of the value of all slots in a particular time period is open to question.

The second is to observe price differences for city-pairs between congested and uncongested airports. Again, this is a useful thing to try to do, but there are multivariate problems – ideally we want to control for differences in the traffic composition (journey purpose, class of travel, flexibility of ticket etc.), differences in the airlines' offers (seat pitch, meals, frequent flyer points etc.), differences in the access costs to the airports, differences in aero charges, congestion costs and turnaround times. We believe that price differences are the outcome of an amalgam of effects rather than a pure scarcity rent.

The third is to do some kind of accounting study of airline balance sheets. Airlines clearly 'own' valuable property at capacity constrained airports, and how they handle that in their accounts is of interest. Specifically, if airlines view scarce slots as tradeable assets, then this has implications for the calculation of equivalent price premia.

The fourth is to examine the dynamic interaction between demand and supply across relevant airport networks, and infer the equilibrium 'shadow price' arising at any given airport which needs to impose a price premium to match demand with available capacity.

1.6 Does economic theory suggest there could be a scarcity rent at Heathrow?

The papers we have read conjecture that the marginal cost curve for Air Traffic Movements (ATMs) is a reverse L shape, given fixed infrastructure capacity, permitted times of operation etc. However, we feel that there is a lack of clarity concerning short vs. long run properties of the cost function. In practical terms, we would conjecture that there are three relevant permutations:

- Plane size/route structure fixed and infrastructure fixed – which we would characterise as the very short run
- Plane size/route structure variable, infrastructure fixed – short run
- All factors variable – long run

Thus, in the short run, we might expect there to be more of a three-part curve – the uncongested section, an upward sloping section where more movements can be handled at the expense of increased queues, delays and turnaround times, and a vertical section given technology, departure intervals, approach controls etc. Airlines supply ATMs but sell seats. In the vicinity of capacity for ATMs, a number of behaviours make seats a bit more elastic than ATMs – larger planes replace smaller on given routes, frequencies on thick routes with wide bodies displace frequencies on thinner routes with narrow bodies, load factors increase as

travellers accept less preferred departure times. So until every departure out of Heathrow is a full A380, there can be some seat growth without ATM growth, assuming terminal facilities, gates and the availability of aircraft are not the ultimate constraining factors.

So we see the realistic short run marginal cost (SRMC) curve as strongly upward sloping as capacity is approached – but not literally vertical. This suggests that scarcity rents are likely to vary quite appreciably by time of day/week and by ticket type. Presenting a single average number (£x per seat) could be very misleading in terms of the pattern of impact, even if arithmetically correct. Moreover, most of the burden of adjustment to exogenous demand growth has to be taken on the price side, so even if scarcity rents are low today, they might grow quite quickly. In this context, the crucial moving parts are the shape of the SRMC per seat curve and the demand/supply interaction at Heathrow (which depends partly on what is offered at other London airports and the degree of substitutability in the system). But overall, we think theory supports the notion that in airports where utilisation is 99% of rated capacity, scarcity rents of some scale are likely to arise.

1.7 To whom does any such scarcity rent accrue?

In principle, the rents could accrue to some combination of Government, airport operator and airlines, depending on who ‘owns’ the slots and whether/how the airport is regulated. In London, Heathrow is a regulated entity so aero charges are set on the basis of a permitted revenue yield relative to the Regulated Asset Base. So the airport operator is not able to increase aero charges so as to cream off the scarcity rent as they would in an unregulated environment.

That does not mean that the airport operator may not enjoy rents of other kinds. For example, if the RAB is set too high, or can be manipulated, or the permitted yield per unit of RAB is set too high, or there is regulatory lag which takes time to adjust the formula as traffic grows, then the airport operator will receive rents or supernormal profits. Whether economic regulation of Heathrow and Gatwick is efficient is beyond our brief. But if there are rents to the airport operator, they are not scarcity rents.

We conclude therefore that any scarcity rents accrue to the airlines as ‘owners’ of the valuable property. This is why, albeit in a thin market, slots at London Heathrow (LHR) are willingly exchanged for money. So then the question must be confronted *‘If airlines benefit from these scarcity rents, why aren’t they highly profitable?’*. One answer could be that the scarcity rents might not be very large, no more than a few pounds per passenger on average at the present day. Another is that some of the rents might be accruing not to the airline firm/shareholders but for example to existing and legacy labour through terms and conditions and pension commitments. A further answer could be that above-average profits on some routes may be offset by less profitable routes elsewhere within a given airline network. Some of the answers might be found by studies of the cost-side rather than the demand-side. Network airlines offer a product which is differentiated from point-to-point airlines and incur costs in order to do so. In assessing evidence for rents, price differences (and changes in) should be adjusted for cost differences (and changes in).

1.8 Does the empirical evidence suggest there could be a scarcity rent? How reliable is this evidence?

Both Frontier (2019) and SEO/Cranfield (2017) are pieces of work which should be taken seriously, and should form part of the evidence base on which informed judgement should be made. However, these studies are not without their limitations, key among which are the following:

- The fares data is average route revenue, when ideally, finer resolution by time of day, business/economy etc. is desirable.
- The SEO/Cranfield approach of representing scarcity rents by a measured variable (capacity) is useful in separating out the influence of this factor from other unexplained factors. We prefer this approach to Frontier's use of airport specific dummy variables.
- Notwithstanding the previous point, the representation of capacity within both Frontier (2019) and SEO/Cranfield (2017) is overly simplistic, and we believe that there is scope to significantly strengthen this.
- There is a clear advantage in using panel data rather than cross-sectional; if further work is done, we would consider the use of a panel with airport-specific effects constant over time.
- The fares data is travel agency ticket sales data which only accounts for about half the market, although there is some adjustment of the raw data to reflect sales across all channels. It would be prudent to commission a detailed audit of this data to determine fitness for purpose.

We agree with the approach taken by Frontier (2019) in adopting a triangulation approach built around analysis of fares data, analysis of slot trading including their accounting treatment, and shadow cost approaches such as NAPAM. The value-added comes from interrogating each of the approaches in the light of the others.

1.9 What would happen to the rent if capacity were increased?

It is clear from theory that, if capacity is increased, rent will reduce, all else equal. However, in order to quantify this, the practical challenge is to forecast forward from the analysed position in the base year, say 2018, to the position in 2030 and 2040 during the life of the airport system both with and without enhanced capacity. Building a more secure description of the base year is one thing; using the model in predictive mode is another.

1.10 Recommendations

How to move forward from the position reached in summer 2019 is partially dependent on the timelines and decision points within which the scarcity rents issue is being addressed. But we think the following commissions could be considered by the regulator:

- Further work to enhance the existing regression analysis by Frontier (2019) by better treatment of congestion and capacity as outlined in section 6.3.
- A review and audit of the fares data, its representativeness of the market as a whole, and its fitness for purpose in this application, together with any recommendations for improved data, relevant disaggregation and validation.

- A small scoping study to assess the scope for re-couching the econometric model in terms of generalised cost rather than fare, especially with a view to achieving better representation of supply-side responses to congestion.
- Further work on the airline cost side to help determine whether price differences are partly a reflection of cost differences.

Alongside these suggestions for further commissioned work, the core of the way forward lies in policy work to develop the triangulation approach used by Frontier (2019) and described in Chapter 6. It would be useful if the fares approach and the analysis of the slot market were taken forward. More attention could be given to the insights from the NAPAM model since this incorporates demand/supply interaction which is helpful in representing what happens close to capacity. A crucial question for policy is how to forecast scarcity rents forward to 2030 and 2040. Sound judgement and interpretation of the different strands of evidence and their robustness will be essential.

2 INTRODUCTION

Our review is organised around the following questions, each of which accounts for one of the following chapters in the report:

- Chapter 3: Given the current regulatory regime in the UK, does theory suggest that there could be a scarcity rent at Heathrow?
- Chapter 4: Does the empirical evidence from fares analysis suggest that there could be a scarcity rent?
- Chapter 5: To whom does any such rent accrue?
- Chapter 6: What can be done to strengthen analysis of scarcity rent at Heathrow?
- Chapter 7: Combining theory and evidence, what do the findings tell us in terms of understanding the concept of scarcity rents and their existence at Heathrow?

Annex A extracts key arguments directly from the various consultants reports and offers a brief response to each such argument. Annex B is a technical annex relating to Chapter 3.

The context of this report is as follows. The Airports Commission¹ recommended to Government that a third runway and associated works should be constructed at Heathrow to expand capacity. The economic case consisted of the following components:

- The benefits to passengers in terms of lower air fares and improved accessibility to the air transport system relative to the Do-Minimum;
- The impacts on the wider economy;
- The environmental and local area impacts;
- The effects on airlines of reduced scarcity rents especially at Heathrow but also in the future at Gatwick;
- The capital and running costs of the capacity and the way this would be ultimately funded via aero charges.

The Civil Aviation Authority is the economic regulator for Heathrow and Gatwick. Its core remit is to consider the interests of air passengers in the regulatory decisions it makes, most obviously on the setting of aero charges now and over time and specifically the way in which the CAPEX for Runway 3 etc. is to be incorporated into the Regulated Asset Base (RAB).

In principle, there are two categories of passengers who will be affected by the new capacity at Heathrow. LHR is a constrained airport operating at over 99% of available runway capacity.

¹ The Airports Commission (2015) 'Final Report'. Available at:

https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/440316/airports-commission-final-report.pdf

Providing more capacity will enable increased frequencies and new routes. This will divert some passengers who prefer LHR but are squeezed out by capacity constraints and generate some new traffic which would not otherwise fly. The relevant elasticities and response properties within DfT's aviation modelling suite² are what drive the forecasts of this induced traffic. That is not the subject of this report, but is an important element in the overall cost-benefit analysis.

The other group affected by the new capacity is existing passengers using Heathrow. Potentially, there are two countervailing effects on them (leaving to one side any surface access improvements). Firstly, air passengers as a group will need to fund the capital expenditure over its life through aero charges. Secondly, there are argued to be scarcity rents at congested airports, of which LHR is one (see for example Starkie (2008)³). Supposing this to be the case, the additional capacity would increase competition and ticket prices would fall. The relative sizes of these two effects determine whether from a fares point of view passengers stand to gain or lose. There may also be non-fare effects on passengers and in a traffic growth environment the balance between the two effects could be expected to change over time.

In its Final Report, the Airports Commission (2015)⁴ judged that capacity expansion would not impact upon passenger fares, as the increase in airport charges resulting from expansion would not be passed through to passengers by airlines (p120). Over the period 2017-19, several key stakeholders have commissioned reports on this topic, summarised as follows:

- 2017: Heathrow Airport Limited (HAL) commissioned Frontier Economics to produce a report entitled '*Competition and Choice 2017 – A Report Prepared for Heathrow*'⁵. This report included an econometric analysis of airline fares and used the results of that analysis to conclude that fares would '*decrease by 23% relative to other London airports as a result of removing the capacity constraint*'.
- International Airlines Group (IAG) responded with a critique of the Frontier (2017) report and their own assessment of the impact of capacity constraint at Heathrow airport⁶.

² DfT (2017) UK Aviation Forecasts, available at: <https://www.gov.uk/government/publications/uk-aviation-forecasts-2017>

³ Starkie, D. (2008) Aviation Markets. IEA Studies in Competition and Regulatory Reform, Ashgate (especially Chapter 5 on 'Reforming Airport Regulation'.

⁴ Airports Commission (2015) Final Report, available at: <https://www.gov.uk/government/publications/airports-commission-final-report>

⁵ Frontier Economics (2017) '*Competition and Choice 2017 – A Report Prepared for Heathrow*'. Available at: https://www.caa.co.uk/uploadedFiles/CAA/Content/Accordion/Standard_Content/Commercial/Airports/HAL%20-%20Frontier%20Competition%20and%20Choice.pdf

⁶ Letter to Stephen Gifford, Consumer & Markets Group, Civil Aviation Authority, dated 6th July 2018, available at:

- 2018: CAA commissioned FTI to produce a report '*A Critique of Published Reports Regarding Scarcity Rents at Heathrow Airport*'⁷, which assessed the economic rationale and econometric analyses within the Frontier (2017) report and the IAG response. This report also identified areas where additional analysis could further illuminate these issues.
- 2018: IAG released a subsequent report entitled '*The Effect of Congestion at Heathrow Airport – Comments on Frontier and FTI Reports*'⁸ by RBB Economics. The report commented on the earlier Frontier (2017) and FTI (2018) reports, arguing that an increase in the cost of congestion was unlikely to cause ticket prices to increase. The report did not however present definitive evidence on the existence or scale of scarcity rents.
- 2019: HAL commissioned Frontier Economics to undertake a further study '*Estimating the Congestion Premium at Heathrow*'⁹. Instead of focussing upon econometric analysis as per their 2017 report, Frontier broadened their scope to incorporate two further analyses – namely: a) analysis of demand/supply interaction and the shadow price needed to 'choke off' excess demand and b) analysis of slot prices – and then sought to 'triangulate' between the three approaches.
- 2019: CAA commissioned FTI to undertake a further study '*Scarcity rents and slot prices at Heathrow Airport*', which is still in progress. This work employs data on the reported prices of slots traded in the secondary market at Heathrow airport as an 'indicator' of the level of scarcity rents.

Whilst the aforementioned portfolio of studies is the focus of our review, we also draw substantively on three further studies, namely:

- Airports Commission (2015) '*A Note from Expert Advisors, Prof. Peter Mackie and Mr Brian Pearce, on key issues considering the Airports Commission Economic Case*'¹⁰.

https://www.caa.co.uk/uploadedFiles/CAA/Content/Accordion/Standard_Content/Commercial/Airports/Files/IAG%20CAP1658%20response%20FINAL.pdf

⁷ FTI Consulting (2018) '*A Critique of Published Reports Regarding Scarcity Rents at Heathrow Airport*'. Report to CAA. Available at: <http://publicapps.caa.co.uk/docs/33/CAP1722b%20FTI%20Consulting%20-%20Report%20to%20CAA%20on%20Scarcity%20Rents%20-%20Final%20-%2030%20October%202018.pdf>

⁸ RBB Economics (2018) '*The effect of congestion at Heathrow Airport – Comments on Frontier and FTI reports*'. Report to IAG.

⁹ Frontier Economics (2019) '*Estimating the Congestion Premium at Heathrow*'. Available at: https://www.caa.co.uk/uploadedFiles/CAA/Content/Accordion/Standard_Content/Commercial/Airports/H7/Estimating%20the%20congestion%20premium%20at%20Heathrow.pdf

¹⁰ Airports Commission (2015) '*A Note from Expert Advisors, Prof. Peter Mackie and Mr Brian Pearce, on key issues considering the Airports Commission Economic Case*'. Available at:

This note summarised some of the key issues of relevance to the Economic Case for airport expansion. On scarcity rents, the note concluded: *‘Our view is that the assumptions concerning the working of the market for slots and the transfer of scarcity rents between producers and consumers produce CBA results which are a little conservative. It is vital to remember that the Do-Minimum reference case is a journey into the unknown and it is inherently more difficult to predict the behaviour of the actors in highly constrained conditions than in what, over the appraisal period, is broadly a facilitation strategy designed to accommodate demand growth’* (p3).

- Another report commissioned by the Airports Commission from SEO Economic Research (2015) *‘Expanding Airport Capacity: Competition, Connectivity and Welfare: Discussion of Options for Gatwick and Heathrow’*¹¹ also deserves attention. This report argued that airlines at Gatwick and in particular at Heathrow generate scarcity rents due to excess demand and price capping of charges. Furthermore, the report showed that these scarcity rents would increase in the future under a Do-Minimum scenario, because excess demand would increase further if no additional airport capacity were to become available. The study assumed that any increases in aero charges would be absorbed by the airlines through reduction in airline scarcity rents.
- Last but not least, SEO & Cranfield University (2017) *‘The impact of airport capacity constraints on air fares’*¹² conducted econometric analysis to investigate whether air fares were higher at airports subject to excess demand. They concluded that: *‘...higher levels of capacity utilisation are indeed associated with higher air fares, all other things being equal. We estimate this total additional fare premium at congested European airports at €2.1 billion today. Airport capacity shortages in Europe are becoming increasingly severe. Based on EUROCONTROL’s ‘Challenges of Growth’ forecasts, the total fare premium levied by airlines at congested airports is projected to reach € 6.3 billion by 2035’* (page i).

Across the 2017-19 reports, there is very little symmetry of views between the airport operator (HAL) and the main airline operator (IAG) at Heathrow. Consultants for HAL argue that there is evidence from ticket prices of appreciable scarcity rents enjoyed by airlines at LHR. Consultants for IAG argue that the analysis is flawed, and IAG argue that there are no scarcity rents or alternatively if there are, they accrue to the airport operator, not the airline. There is also disagreement over the size of the change in aero charges required to fund the scheme.

https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/438981/economy-expert-panelist-wider-economic-impacts-review.pdf

¹¹ SEO Economics (2015) Expanding Airport Capacity: Competition, Connectivity and Welfare – Discussion of Options for Gatwick and Heathrow. Report to the ITF of the OCED. Available at: https://www.itf-oecd.org/sites/default/files/docs/15cspa_airports-discussion-options.pdf

¹² SEO Amsterdam Economics & Cranfield Centre for Air Transport Management (2017) The impact of airport capacity constraints on air fares. Report for ACI EUROPE.

To help shed light on the relative merits of these competing arguments, the CAA has commissioned ITS Leeds to provide a commentary on the various reports, give a view, and suggest ways forward. To be clear about the scope of our work, we have been asked to consider the issues from two standpoints:

- 1) from the standpoint of economic theory;
- 2) from the standpoint of empirical analysis, focussing especially on econometric analysis of fares data.

It should be acknowledged that some of the aforementioned reports bring to bear other forms of empirical analysis, namely analysis of the secondary slot market, analysis of demand/supply interaction and the shadow price of capacity, and analysis of airline/airport balance sheet position. These alternative empirical analyses will be referred to at various junctures in our report, but detailed review falls outside our terms of reference.

3 GIVEN THE CURRENT REGULATORY REGIME IN THE UK, DOES THEORY SUGGEST THAT THERE COULD BE A SCARCITY RENT AT HEATHROW?

This chapter draws upon technical concepts from economic theory to consider whether there could be a scarcity rent at Heathrow. It should be emphasised that this chapter represents an ‘in principle’ discussion – even if the theoretical argument is made that a rent could arise, it does not necessarily follow that a rent will manifest ‘in practice’. The latter will be the concern of subsequent chapters of this report.

3.1 Comments on the theoretical framework presented by Frontier and FTI

The Frontier (2017) and FTI (2018) reports present a simplified and highly stylised representation of the market for seats on competing routes A and B operated by different airlines from a capacity-constrained airport such as Heathrow¹³, whereby:

- Demand is linear and downward sloping; it should be acknowledged that linearity is a common expositional simplification, but is unlikely to hold in practice.
- Short run marginal costs are constant (implying that $MC=AC$); again, this seems a reasonable expositional simplification, although in practice it is possible that the supply function could to some extent be upward sloping, and that MC could deviate from AC. In the context of LHR, MC will comprise operating costs directly incurred by airlines plus taxes and charges imposed by Government and the regulator respectively.
- At some quantity, supply is subject to a hard (vertical) constraint, which in the present context will be dictated by the availability of slots, planes and the size of those planes. Again this construct would seem reasonable in terms of expositional convenience, although in practice it is possible that the constraint could be more graduated, reflecting some degree of supply-side flexibility in managing capacity (we will return to this point below).
- Routes A and B represent separate sub-markets subject to different capacity constraints.

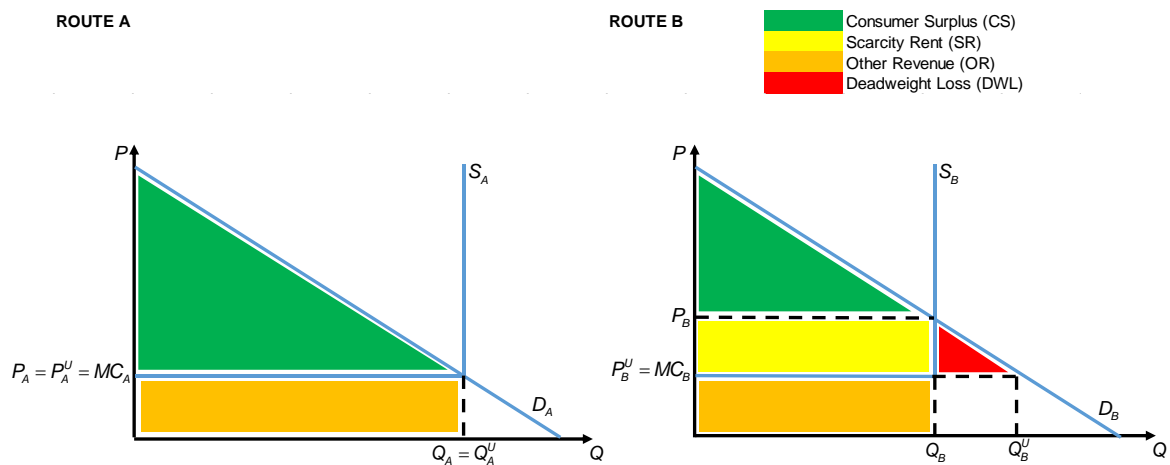
The essence of this is illustrated in Figure 3.1 below. Whereas the Frontier (2017) and FTI (2018) reports employ such diagrams to illustrate the prevalence and scale of scarcity rent, here we show the broader welfare impact on all relevant economic agents¹⁴. That is to say,

¹³ This stylised representation is also consistent with that shown in Appendix A of TAG Unit A5.2 Aviation Appraisal, available at: <https://www.gov.uk/government/publications/webtag-tag-unit-a5-2-aviation-appraisal-may-2018>

¹⁴ Appendix A of DfT TAG Unit A5.2 ‘Aviation Appraisal’ adopts a similar breakdown of welfare agents, available at: <https://www.gov.uk/government/publications/webtag-tag-unit-a5-2-aviation-appraisal-may-2018>

we also show the existence/scale of consumer surplus¹⁵, other producer revenue¹⁶ and deadweight loss¹⁷. Box 1 provides a broad definition of the concept of ‘rent’; whereas the Frontier and FTI reports attributed this rent entirely to scarcity of supply at Heathrow, it is possible that other factors may contribute to this rent, and we will return to this point later.

Figure 3.1: Two routes with different properties in terms of capacity constraint



For both routes, the price/output combination is determined where demand and supply intersect. On Route A, the unconstrained demand Q_A^U coincides with the equilibrium demand Q_A , and this generates both consumer surplus and producer revenue. By contrast, on Route B, the unconstrained demand Q_B^U exceeds capacity, such that supply is constrained at the equilibrium demand Q_B , and this additionally generates scarcity rent (since the equilibrium price P_B is above marginal cost) and deadweight loss (since there is a loss in efficiency relative to the perfectly competitive outcome where the airline prices at MC).

The purpose of showing the broader welfare impact is to highlight that the incidence and scale of any scarcity rent will inter-play with other components of welfare – namely other producer revenue, consumer surplus and deadweight loss.

¹⁵ ‘Consumer surplus’ is the excess of willingness-to-pay for a given quantity of a good over its price.

¹⁶ This is revenue which exactly covers the cost of supplying a given quantity of a good.

¹⁷ ‘Deadweight loss’ is the loss of economic efficiency that arises when the free market equilibrium for a good is not achieved.

Box 3.1: Definition of Rent

$$\text{Rent} = (\text{Market-Clearing Price} - \text{Marginal Cost}) * \text{Market-Clearing Quantity}$$

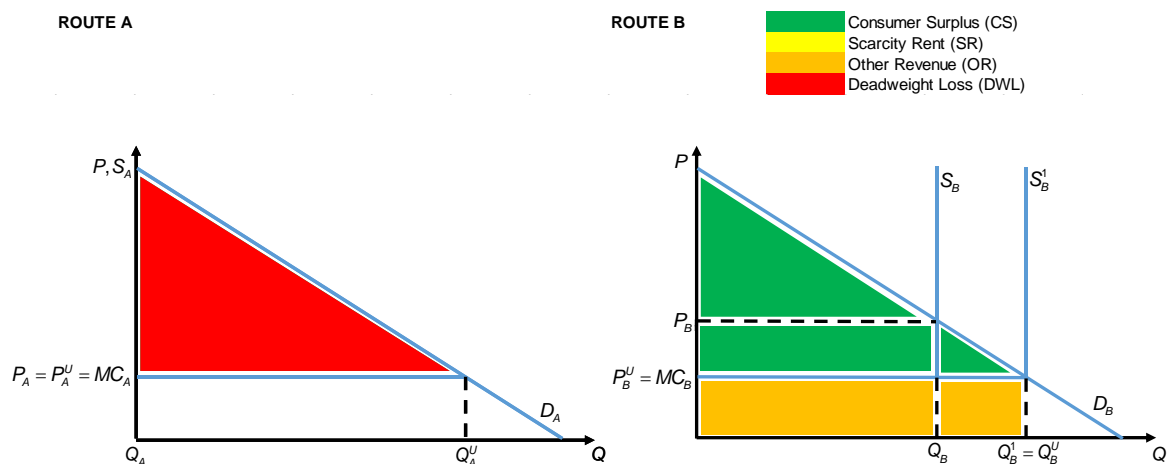
Figure 3.2 is a progression from Figure 3.1, which is presented in the FTI (2018) report. Here, the airline operating Route A observes the existence of scarcity rent in Route B, and responds by switching capacity from A to B. However, the overall level of capacity available to both airlines remains fixed, and this might therefore be seen as a short run position.

It follows that Route A is no longer served, such that substantial deadweight loss now arises in this sub-market. Turning to the second sub-market, the additional capacity redirected to Route B stimulates a rightward shift in the supply function from S_B to S_B^1 , such that unconstrained demand now coincides with capacity¹⁸. Relative to Route B in Figure 3.1, it can be seen that both the scarcity rent and deadweight loss are now absorbed within consumer surplus, such that the latter increases. Since quantity increases but price falls relative to Figure 3.1, the net impact on producer revenue will be an empirical question.

The welfare changes from Figure 3.1 can be summarised:

- Route A: consumer surplus is eliminated and deadweight loss arises
- Route B: consumer surplus increases, both deadweight loss and scarcity rent are eliminated, and other revenue increases – the net impact on the airline will be given by the increase in other revenue minus the loss of scarcity rent


Figure 3.2: Concentration around route facing capacity constraint



¹⁸ Here we assume that the redirection of capacity to Route B incurs no additional marginal cost. If it did, then the supply function would shift upwards as well as rightwards (see Figure 3).

Figure 3.3 is an alternative progression from Figure 3.1, which we introduce to show the potential impact of increasing airport capacity in the long run. In contrast to the short run response in Figure 3.2 where capacity is transferred from Route A to Route B, here the supply functions for *both* Routes A and B (i.e. S_A^1 and S_B^1) shift rightwards – such that both airlines exploit the release of new capacity. Furthermore, both supply functions are shown as shifting upwards to reflect increased aero charges – since at Heathrow these charges include a contribution towards the cost of runway expansion (see Box 3.2).

Figure 3.3: Off-setting of rent by the costs of providing additional capacity



The welfare changes from Figure 3.1 can be summarised:

- *Route A: consumer surplus is reduced and deadweight loss arises – the net impact on other revenue (and the airline’s welfare position more generally) will depend on whether, in revenue terms, the increase in price offsets the reduction in quantity.*
- *Route B: consumer surplus increases, deadweight loss is reduced, scarcity rent is eliminated, and other revenue increases – the net impact on the airline will be given by the increase in other revenue minus the loss of scarcity rent.*

It should again be stressed that these outcomes are illustrative – the precise outcomes in terms of price, quantity and the various components of welfare will depend upon the differentials between the market clearing price/quantities and those that would arise in the unconstrained situation. However, this figure illustrates the relevant mechanics of demand and supply response to capacity increase.

Box 3.2: Definition of aeronautical charges at Heathrow

'Airport charges and passenger volumes', taken from Heathrow company website (<https://www.heathrow.com/company/company-news-and-information/economic-regulation/airport-charges-and-passenger-volumes>)

The level of airport charges that Heathrow levies each year is in accordance to the CAA's pricing formula, which is part of the CAA's decision for the Q6 period of Economic Regulation at Heathrow from April 2014. The formula set by the CAA determines the level Heathrow cannot exceed in charging its airline on a per passenger basis (passenger only), which is also referred to as the 'maximum allowable yield'. The maximum allowable yield is set following a period of formal consultation with the Heathrow airline community.

Heathrow Airport is also subject to the Airport Charges Regulation 2011, which was transposed into UK law from a European Directive for its member states. This regulation aims to ensure a common framework for regulating airport charges, which includes consultation and transparency of information when setting airport charges.

The CAA formula for the maximum allowable yield is calculated by taking the following into account:

- *Retail Price Index = April RPI is used*
- *Value of X = Fixed at -1.5%*
- *Cumulative Development factor = Adjusts the maximum allowable yield to reflect the cumulative changes associated with development CAPEX projects*
- *Capital triggers = Reduction in the maximum allowable yield when the airport has not achieved particular capital investment project milestones*
- *Bonus = Airport can earn a bonus if it exceeds targets for certain service quality measures*
- *Correction factor = Corrects for any under/overrecovery of actual airport charges against the maximum allowable yield*
- *Business rates factor = Adjusts the maximum allowable yield in 2018 to reflect any change in the actual revaluation of business rates undertaken by the Valuation Office Agency 2017 compared to the regulatory allowance*

The maximum allowable yield is then converted to an estimated total revenue that Heathrow can recover. This is estimated by using a forecast for total passengers (i.e. maximum allowable yield multiplied by total passenger numbers).

The total revenue from airport charges is then recovered from Heathrow's structure of airport charges.

The structure of airport charges has been subject to formal consultation with the Heathrow airline community and has resulted in the total airport charges revenue being recovered through three categories; Landing Charges, Departing Passenger Charges and Aircraft Parking Charges. The three categories (excluding ANS) represent 21%, 75% and 4% of the total airport charges revenue respectively.

In our view, an important distinction which does not come through clearly in the Frontier (2017) and FTI (2018) reports is that between the short and long runs. More specifically, we would conjecture that, in practice, three permutations are relevant:

- Plane size/route structure fixed and infrastructure fixed – which we would characterise as the ‘**very short run**’;
- Plane size/route structure variable, infrastructure fixed – ‘**short run**’;
- All factors variable – ‘**long run**’.

Trusting that the above typology is meaningful in operational terms, then the previous diagrams can be interpreted such that:

- Figure 3.1 shows the very short run, since all aspects of supply are fixed.
- Figure 3.2 shows the short run, since route structure is varied but infrastructure is fixed;
- Figure 3.3 shows the long run, whereby SRMC shifts right on both routes, and for each such route the LRMC will in principle be given by the envelope of the two SRMCs

Within this typology, it might be noted that, where infrastructure is fixed, there will be the opportunity to optimise usage of this capacity – but this will serve only to equalise rents over opportunities within time slots, rather than eliminate rents altogether.

It is perhaps obvious that the market for air travel is not one market but a series of overlapping sub-markets with different characteristics and elasticity properties. For example, the sub-market for Mediterranean destinations might be a relatively large group in nature and involve competition between airports as well as airlines, while the long distance city-pair sub-market might tend to be a small group. Devices such as frequent flier programmes, agent commission structures and slot dominance at congested airports are argued to influence market competitiveness (Levine, 1987)¹⁹. The market conditions facing holiday travellers wishing to go to Malaga in two months’ time and a business traveller wishing to go to Mexico City next Monday are probably quite different. Against this background, the nature of the interaction between Routes A and B from a capacity-constrained airport such as Heathrow could be quite different depending on the characteristics of their respective sub-markets.

3.2 Additional insights from the Cournot model

Diagrammatic representation of the problem at hand can be a very powerful tool, but inevitably calls for a degree of simplification in the description of the problem. However, even if we acknowledge and accept the need for some expositional simplification, we feel that the Frontier (2017) and FTI (2018) reports – and IAG’s response – overlook a key feature of the problem. This is because the framework as outlined does not account for the competitive structure of the market of interest – a point which is also made by SEO/Cranfield (2017) .

¹⁹ Levine, M.E. (1987) Airline Competition in Deregulated Markets: Theory, Firm Strategy and Public Policy. Yale Journal on Regulation, 4, pp.393-494.

3.2.1 *Theoretical rationale*

Reflecting upon the structure that prevails in the case of Heathrow, we are drawn to the well-established theoretical finding by Kreps & Scheinkman (1983)²⁰ that, in a capacity (i.e. fixed slot capacity) then price game with efficient consumer rationing (i.e. yield management of aircraft seats) and with linear demand and constant MC, the chosen capacity will be equivalent to that which arises in a standard Cournot model²¹.

Adding recent empirical evidence to these considerations, Lijesen & Behrens (2017)²² investigated whether alternative city- or airport-pairs are viable substitutes and the extent to which they impact upon airline competition between the UK and continental Europe. Using monthly airline-route seat capacity levels, Lijesen & Behrens estimated airlines' strategic reaction to competitors' capacity levels, including competitors on other routes. They found that *'the vast majority of the city-pair markets is less competitive than a symmetric Cournot duopoly'* (p11). This finding contrasts (to some extent) with Koopmans & Lieshout (2016)²³ – cited by SEO/Cranfield (2017) – whose analysis encompassed all city-pairs worldwide (as opposed to focussing on the UK), and concluded that *'most aviation markets can be described as oligopolies with product differentiation'* (p5).

Reconciling these theoretical and empirical insights, our view is that the competitive structure of most UK sub-markets probably falls somewhere between monopoly and Cournot – but possibly with monopolistic competition arising in some sub-markets accessible to new entrants (especially low cost carriers). Moreover, in any of these cases, there is good reason to believe that the market clearing price will be above MC – such that scarcity rent of some scale will in principle arise. In practice, however, it is quite likely that – even at a heavily congested airport such as Heathrow – any scarcity rent will arise to different degrees across different city-pairs, at different times of day, and at different points in the year. That is to say, scarcity rent will apply in a variable fashion across Heathrow's slot portfolio.

3.2.2 *Diagrammatic illustration*

Building upon our theoretical assertion that scarcity rent could arise at Heathrow, we now introduce a further diagram of our own, Figure 3.4, which we believe better reflects the prevailing competitive structure. Note that rather than compare two routes (which might be

²⁰ Kreps, D.M & Scheinkman, J.A. (1983) Quantity Precommitment and Bertrand Competition Yielding Cournot Outcomes. *The Bell Journal of Economics*, 14 (2), pp.326-337.

²¹ The 'Cournot model' presents an industry structure where a small number of companies (a pair of companies in the standard presentation) compete on the quantity of output produced, but these quantities are determined independently and simultaneously.

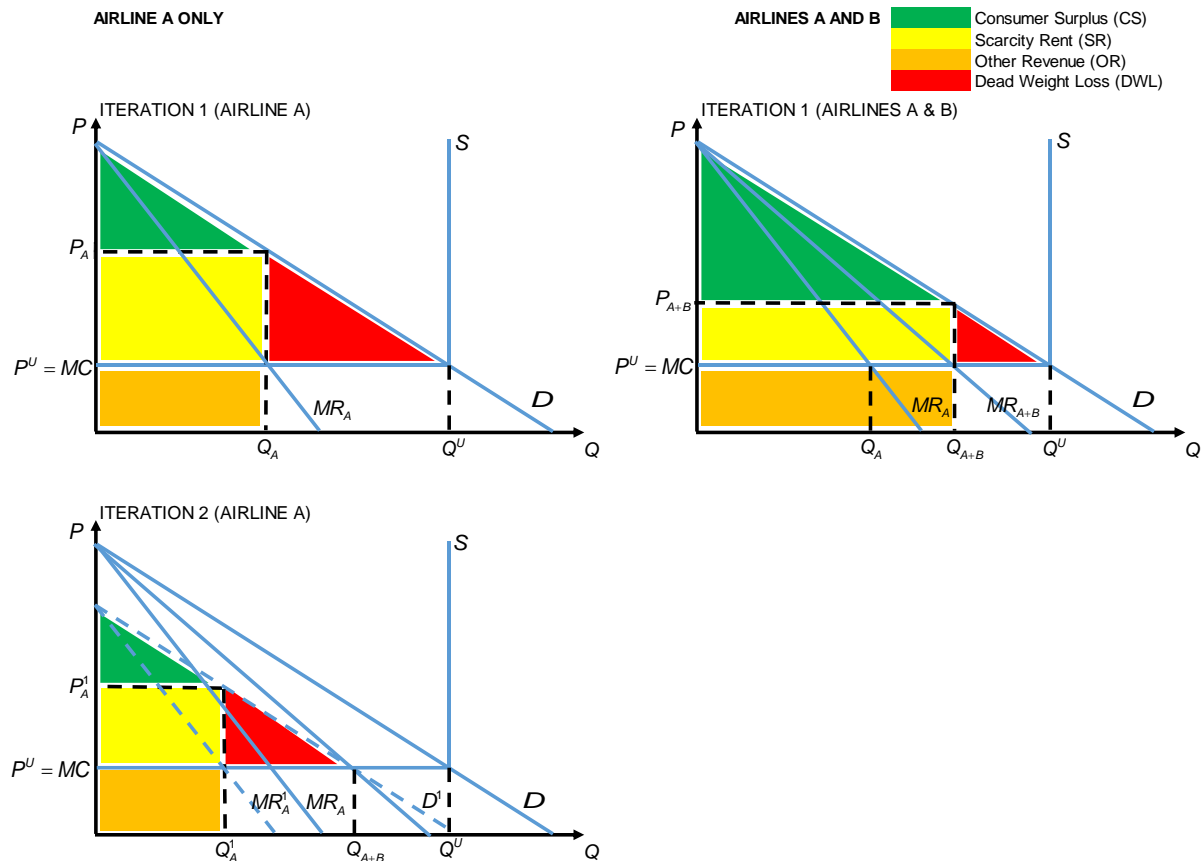
²² Lijesen, M. & Behrens, C. (2017) The spatial scope of airline competition. *Transportation Research Part E*, 99, pp.1-13.

²³ Koopmans & Lieshout (2016) Airline cost changes: to what extent are they passed through to the passenger? *Journal of Air Transport Management*, 53, pp.1-11.

considered two sub-markets), we instead compare two airlines competing for business on a given route (which might be a city-pair operating through Heathrow) via a series of iterations.

More specifically, we couch the analysis as a Cournot duopoly involving Airlines A and B, where Airline A has the greater market power. Capacity is determined at the outset via slot allocation and, given that capacity, both airlines operate at $MR=MC$ – which contrasts with $P=MC/S$ in Figures 3.1-3.3²⁴. Furthermore, both airlines are assumed to face the same linear demand and constant MC.

Figure 3.4: Capacity then price game with efficient consumer rationing



- ITERATION 1:
 - Top left diagram: Airline A claims 50% of the total capacity ($Q_A = 0.5 * Q^U$), on the expectation that Airline B will also claim 50%.
 - Top right diagram: Expecting that Airline A would claim a higher share of 75%, Airline B in practice claims only 50% of the residual capacity (i.e.

²⁴ Note however that, for simplicity of exposition, we show the market as a whole here as operating at less than capacity.

$Q_B = 0.5 * (Q^U - Q_A)$, such that 75% of the market capacity is now accounted for (Q_{A+B}) .

- ITERATION 2:
 - Bottom left diagram: Airline A revises its position by claiming 50% of the newly-defined market capacity Q_{A+B} (i.e. $Q_A^1 = 0.5 * Q_{A+B}$), such that Airline A's quantity is reduced relative to iteration 1 (i.e. $Q_A^1 < Q_A$), and so on.

The eventual outcome depends on the relative market power of the two airlines. If Airline A holds an effective monopoly then Q_A will be the equilibrium quantity. If on the hand both airlines hold equal power, then the outcome will be a third-share of capacity for each.

3.2.3 Conclusions from the Cournot model

Moreover, in the case shown here, it can be seen that if Airline A operates in isolation, then scarcity rent arises in tandem with deadweight loss, whereas if Airline B also operates, then the (marginal) value of the scarcity rent decreases, deadweight loss decreases, and these reductions are captured within increased consumer surplus. In other words, the competitive pressure unambiguously delivers social welfare improvement.

The duopoly example is illustrative and could be extended. Some city-pairs especially from secondary hubs are more competitive. The first two screens of a Manchester-Shanghai search yielded eight credible one stop options. Other city-pairs such as London to Cape Town are closer to duopoly because the alternative routeings involve more significant time penalties relative to direct non-stop service. Another simplification of the stylised representation is that it assumes a single market clearing price when in practice airlines deploy price discrimination to maximise revenue capture.

3.3 Additional insights from oligopoly theory

This section draws upon concepts from oligopoly theory to shine a different light on the prevalence and scale of scarcity rents at Heathrow. The broad proposition here is that airline markets are characterised by market power, which varies across sub-markets, and that the degree of market power is likely to be related to the level of scarcity governing supply conditions.

3.3.1 Theoretical rationale

The rationale behind this approach is based upon two steps, as follows:

- i) Recent consultancy reports have implicitly assumed that price is set at the market-clearing level. As already noted above, we would question this assumption, and believe that the Kreps & Scheinkman (1983) result offers a more defensible statement of market conditions. That is to say, in a capacity (i.e. fixed slot capacity) then price game with efficient consumer rationing (i.e. yield management of aircraft

seats), and with linear demand and constant MC^{25} , the chosen capacity **will be equivalent to** that which arises in a standard Cournot model²⁶.

- ii) If we accept i), and thus assume that the standard Cournot model is applicable to the case of a city-pair²⁷ operating out of Heathrow, and consider the simplest case of this model comprising two airlines (A and B) engaged in oligopolistic competition, then it can be shown (see Annex B) that the following identity will hold:

$$R_i = P - MC_i$$

$$= -P \cdot \frac{s_i}{\varepsilon_{Q_{A+B}, P}} \quad i = A, B \quad (1)$$

Where R_i is the rent²⁸ to airline i , P is the price (which is common to both airlines), MC_i is the marginal cost incurred by each airline, s_i is the market share of each airline, and $\varepsilon_{Q_{A+B}, P}$ is the elasticity of demand (for the city-pair market of interest) with respect to price.

Since in the present case, step i) is underpinned by the notion of a capacity-then-price game, it would not seem unreasonable to interpret the above statement of rent as an approximation to 'scarcity' rent. In this way, we arrive at a notion of scarcity rent whereby, given that capacity is binding, the resultant rent is determined by three factors working in combination, namely:

- The elasticity of demand with respect to price
- The market share of the airline
- The market-clearing price

3.3.2 Empirical illustration

Following Frontier (2019), let us assume average short haul and long haul fares at Heathrow of £237 and £1080 respectively. Then, with reference to the IATA elasticities given in Table 3.1, let us focus upon the Intra Europe market (i.e. row 2 of the table). In passing, note that different elasticities are given for short vs. long haul, and for different geographical market

²⁵ This implies $AC=MC$, which is admittedly a restrictive assumption.

²⁶ Relating this model to Heathrow, we consider that a capacity-then-price game is applicable, since the price adjustment lag will tend to be shorter than the output adjustment lag.

²⁷ Obviously, market segments can be defined in different ways, but there seems to be a reasonable level of consensus from the recent Heathrow consultancy reports that the city-pair sub-market is especially relevant to the issue of scarcity rent.

²⁸ That is to say, the difference between P and MC .

segments (covering both regions and spatial levels).²⁹ We defer to these IATA elasticities in preference to those given in WebTAG, since they show the significant heterogeneity which exists across market segments. By contrast, WebTAG elasticities would seem to be couched at the UK National level. It should come as no surprise that these different market segments exhibit quite different elasticities, for the following reasons:

- Route/Market level: in this case, fares increase on all carriers serving a given route (e.g. due to an increase in airport fees and charges).
- National level: fares increase on all routes to and from a particular country (e.g. due to a higher national departure tax), making demand less elastic than Route/Market level.
- Supra-National level: fares increase at a regional level across several countries (e.g. an aviation tax imposed on all member states of the European Union), making demand even less elastic.

Furthermore, short haul demand will always be slightly more elastic than long haul, reflecting greater opportunity for inter-modal substitution.

Table 3.1: IATA Price Elasticities of Passenger Demand³⁰

	Route/Market level		National level		Supra-national level	
	Short haul	Long haul	Short haul	Long haul	Short haul	Long haul
Intra North America	-1.54	-1.4	-0.88	-0.8	-0.66	-0.6
Intra Europe	-1.96	-1.96	-1.23	-1.12	-0.92	-0.84
Intra Asia	-1.46	-1.33	-0.84	-0.76	-0.63	-0.57
Intra Sub-Saharan Africa	-0.92	-0.84	-0.53	-0.48	-0.4	-0.36
Intra South America	-1.93	-1.75	-1.1	-1	-0.83	-0.75
Trans Atlantic (North America - Europe)	-1.85	-1.68	-1.06	-0.96	-0.79	-0.72
Trans Pacific (North America – Asia)	-0.92	-0.84	-0.53	-0.48	-0.4	-0.36
Europe-Asia	-1.39	-1.26	-0.79	-0.72	-0.59	-0.54

Now using the identity for R_i (1), we calculate the rent which arises at average fares for:

- Both short and long haul;
- Each of the three spatial market levels;

thereby yielding six calculations. We calculate the rent both in absolute terms (in £) and as a percentage of average fares. These calculations are illustrated in Figure 3.5 (for percentages)

²⁹ In the example here, market segments are defined in terms of city-pairs. In the real world, journey purpose, point-to-point versus network etc. might also be relevant differentiators.

³⁰ IATA (2008) Air Travel Demand, IATA Economics Briefing No. 9.

and Figure 3.6 (for absolutes), where we plot market share (on the horizontal axis) against rent (on the vertical).

With reference to the identity for R_i (1), it should come as no surprise that, for any given fare and elasticity, there will be a linear increasing relationship between market share and rent. However, what is stark is the differential between the rents that arise for short vs. long haul and for different market segments – this differential is driven by the elasticity, which shows considerable variability across these different parts of the market.

For example, if for a given short (or indeed long, since they have common elasticities) haul city-pair (i.e. Route/Market level), Airline A were to hold 50% of the market, then this would imply a rent of around 25%. By comparison, if Airline A were to hold 100% of the market, then the rent would increase to 50%. Repeating the same short haul example but using National level elasticities, then Airline A's rents would increase to around 40% and 80% respectively.

3.3.3 Conclusions from oligopoly theory

Drawing upon theoretical relationships associated with the oligopolistic behaviour of firms, this section has proposed a simple method for estimating an approximation to the scarcity rent – with an illustration based on Heathrow. In these terms, it has been shown that the level of rent is highly sensitive to:

- elasticity of demand – which varies across market segments;
- airline market share – which we conjecture is related to market power and supply scarcity in a symbiotic manner.

Moreover, this section has presented an alternative proposition for the existence of rent, whereby different airline sub-markets are characterised by different degrees of market power, and the latter is related to the scarcity of supply. Given its simplicity, this method is intended as a potentially useful means of validating other more detailed analyses of scarcity rent – rather than as a primary analysis in its own right. In order to operationalise the method for the case of LHR, data would be needed on fares and airline market shares for a sample of city-pairs; if available, city-pair-specific elasticities should be used, but failing that, elasticities should be taken from the most appropriate IATA segment.

Figure 3.5: Plot of market share (%) vs. rent (% of price)

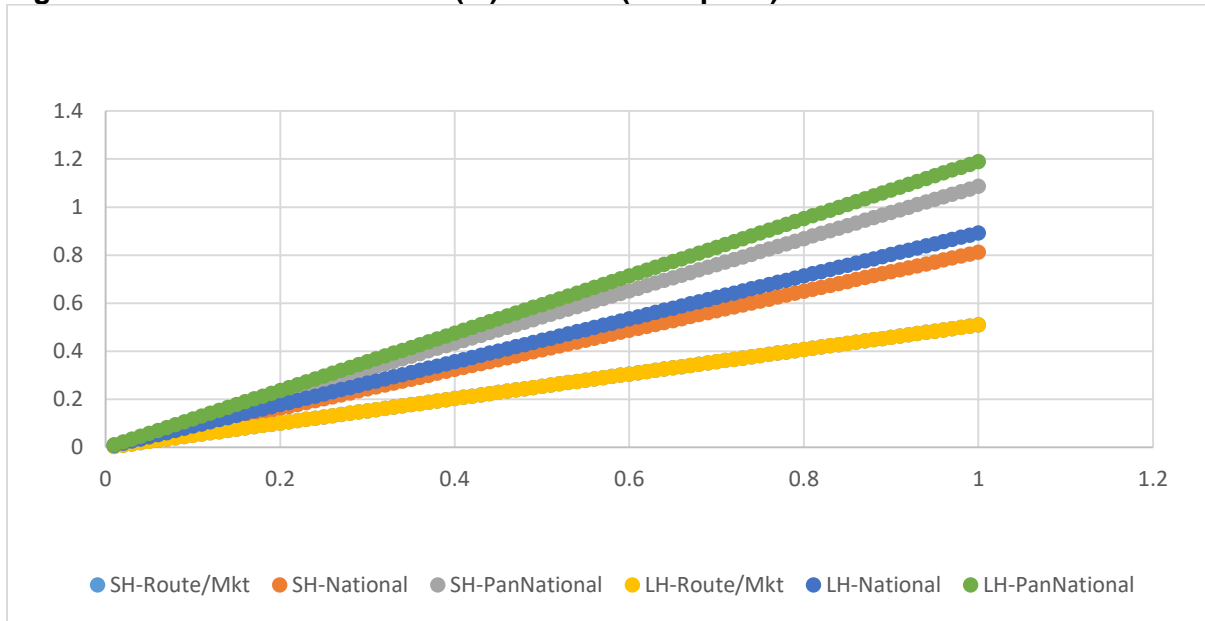
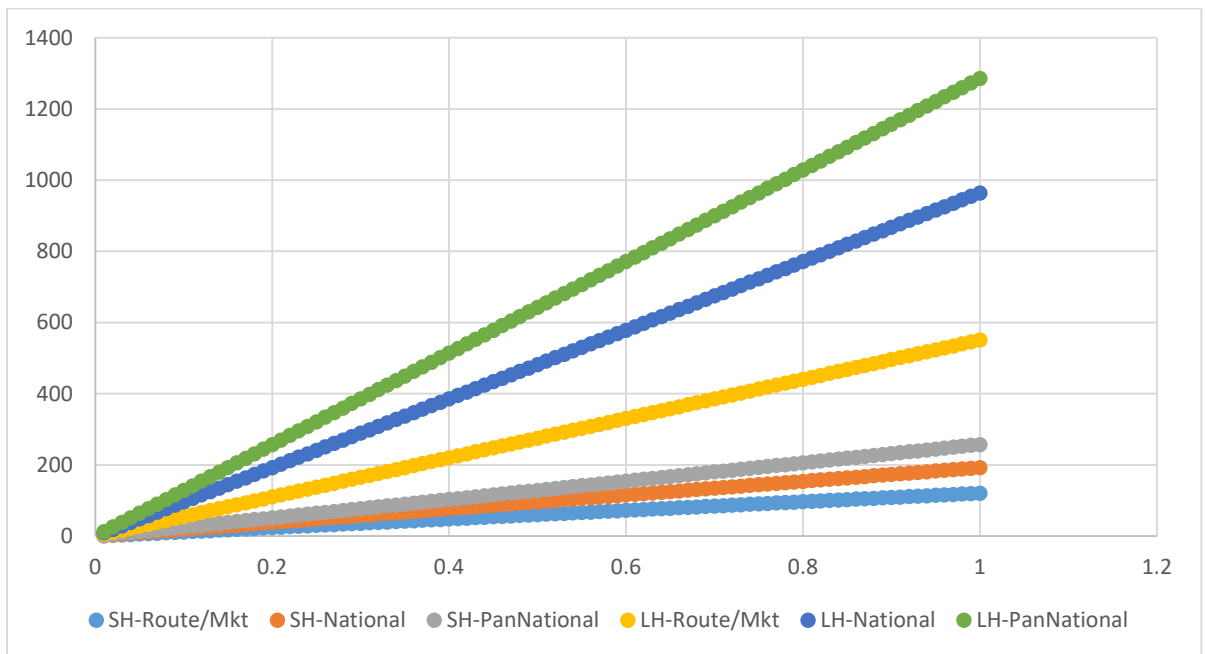



Figure 3.6: Plot of market share (%) vs. rent (absolute in £)



3.4 What would happen to rents if traffic growth were to continue but capacity not increased?

Finally, and for completeness, we should also countenance the scenario where traffic growth continues but Heathrow expansion does not proceed. In terms of the previous figures, this would manifest as an outward shift in demand, such that, all else equal, there will be an increase in any scarcity rent. This is illustrated below, for the case of Route A from Figure 3.1.

Figure 3.7: Reproducing Route A in Figure 3.1, but assuming increased demand with no increase in capacity



The welfare changes from Figure 3.1 can be summarised:

- *Route A: consumer surplus is reduced and scarcity rent increases.*

Although somewhat outside our terms of reference, our appreciation is that the real world equivalent of this figure is policy-critical. Starting from the capacity utilisation position today and fast-forwarding ten years, what do we really expect the relevant demand and supply elasticities to be at the constrained airport? The functional forms proposed by the consultants and developed here are convenient expositional devices but should not be treated as gospel.

4 DOES THE EMPIRICAL EVIDENCE SUGGEST THAT THERE COULD BE A SCARCITY RENT? – INSIGHTS FROM FARES ANALYSIS

Whereas Chapter 3 has drawn upon theory to inform the question of whether there could in principle be a scarcity rent at Heathrow, the present chapter will bring to bear empirical evidence to assess:

- whether a rent has in practice accrued;
- if it has accrued, the scale of that rent;
- any other empirical features of the rent.

As was noted at the end of Chapter 2, our review of empirical evidence will be limited to the econometric modelling of fares data. Across the portfolio of consultancy reports on scarcity rents and LHR, econometric modelling of fares would seem to represent the predominant approach to empirical analysis. Some reports have drawn upon other empirical approaches, but these analyses have been modest in scale, and they have generally been seen as means of validating econometric modelling – as opposed to analyses which stand on their own two feet. Given this remit, the key empirical contributions to the evidence base have been:

- Frontier (2017) – footnote 5
- Frontier (2019) – footnote 9
- SEO/Cranfield (2017) – footnote 12

4.1 Introductory comments

Before offering a range of observations on the econometric modelling approaches employed in the aforementioned studies, it may be helpful to make some introductory comments aimed at setting these empirical contributions within the context of the theoretical background covered in Chapter 3.

Whilst we have issued a range of qualifications, and suggested some enhancements, both Frontier and SEO/Cranfield would seem to subscribe to the essence of the theoretical framework presented diagrammatically in Chapter 3. It would therefore seem reasonable to seek to rationalise the econometric modelling approaches within this context.

In order to translate the theoretical framework to an empirical one, one would ideally need data which allows revelation of the relevant constructs, namely demand, supply and the associated cost functions (marginal and average) at an appropriate level of detail. Indeed, in order to definitively establish the scale and existence of any scarcity rent, one really needs to understand the dynamic interactions between demand and supply, as the capacity constraint in any given sub-market is encountered and agents on both sides respond to those conditions. Equipped with such understanding, one can draw interpretations regarding the price/cost margin (i.e. the rent) with a degree of authority.

Unfortunately, the available data allows only partial revelation of the relevant constructs – indeed most of the available data relates to the demand-side, and we have limited data at our

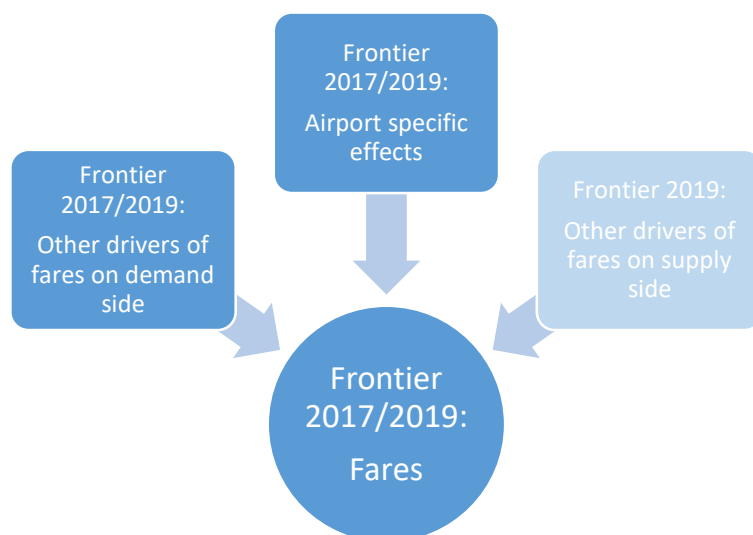
disposal relating to supply-side. The latter is restricted to capacity and competition, and we have little insight on the cost base of the relevant agents.

Against this background, the econometric modelling approaches employed by Frontier and SEO/Cranfield might be seen as an attempt to directly measure the price/cost margin – taking care to control for the key drivers on demand- and supply-side, but stopping short of full representation of D/S interaction. This approach relies upon a degree of pragmatism, but in the absence of data on the cost base, it is arguably the best (perhaps only) available approach to estimating scarcity rents using econometric modelling approaches.

4.2 Comparative assessment of the Frontier and SEO/Cranfield econometric analyses

Using data on flights departing from London and European hub airports, Frontier (2017) estimated a linear regression model to explain fare (or more specifically the logarithm of fare) against various characteristics of travel demand (or more specifically characteristics of the trips consumed). The latter included distance (and whether the distance amounted to long haul), frequency, journey purpose, whether passengers were transferring, whether the flight was by low cost carrier, as well as an indicator variable for Heathrow (which was the focus of Frontier’s attention in terms of quantifying scarcity rent). As we shall see in due course, a significant enhancement of the follow-up Frontier (2019) study was to explain fare against not only characteristics of travel demand, but also characteristics of travel supply.

Figure 4.1: Simplified representation of the rationale underpinning Frontier’s econometric estimate of scarcity rents



Frontier’s (2017) approach to the analysis has been criticised. In undertaking a review of the approach, we have tried to divide this critique into two categories. The first (1 in the list below) category encompasses relatively fundamental issues regarding the conceptual validity of the approach. This includes issues with the model set-up and variables included, as well as the quality of the data – especially the fares data. These considerations are central to the analysis and have implications for the potential usefulness of the work. However, we have also tried to be pragmatic in our assessment, recognising that there is a dearth of research in this area –

and that the Frontier studies are the only econometric analyses which have sought to discern a specific 'Heathrow' effect. Therefore, rather than fixate on the weaknesses of Frontier (2017), we instead seek to present a balanced view as to the likely weight that can be placed on this work (or enhancements to this work) in reaching conclusions on the prevalence and scale of scarcity rent at Heathrow.

The second (2 in the list below) category encompasses more technocratic issues associated with the econometric modelling. These include diagnostic testing and estimation issues. In general, we do not consider these issues to be fundamental to the validity of the approach and with further work could be dismissed or the econometric approach amended to accommodate. However, we do of course recognise that the magnitude of the estimated effects may be impacted by the resolution of these issues.

To summarise our approach, we have considered:

- 1a) The form of the model and missing variables
- 1b) Quality of data used – particularly the fares data
- 2) Estimation and testing issues

In what follows, we will focus on 1a) and 1b).

4.3 The fundamental limitation of the 'fares' data – it is actually average revenue per route

The raw demand/revenue data which underpins the analysis of fares in both the SEO/Cranfield (Box 4.1) and Frontier (Box 4.2) studies derives from essentially the same source, but does exhibit some minor differences. The source of the data is travel agents; the data is considered reasonably reliable, but a significant limitation is that this data accounts for just under half of passenger bookings – meaning that the sample could be unrepresentative. That said, SEO/Cranfield explicitly mentioned that the OAG platform used to extract this information adjusts the raw data so that it is reflective of the broad sales patterns across other channels e.g. sales direct from airlines. To the best of our understanding, there is no other source of fares data that is readily available.

During our review, it has become apparent that a fundamental issue with the data is that it represents average revenue per route rather than fare per se. Whilst air is not unique in exploiting this simplification in demand/fares/revenue modelling³¹, the lack of resolution in the dependent variable is in this case compounded by not distinguishing between fares paid at different times of day or for different travel classes³².

³¹ Indeed, standard industry methods of demand modelling/forecasting on GB rail – which are highly developed, and prescribed to industry users via the so-called 'Passenger Demand Forecasting Handbook' (PDFH) – rely upon average revenues for point-to-point (aka station-to-station) journeys.

³² This distinction *is* made by rail.

The implication of this comes when it is considered how airlines respond to congestion and capacity constraints. As discussed in both the SEO/Cranfield and Frontier reports, there is evidence that airlines do change the service mix, cabin mix and travel times in response to capacity/congestion. This presents a challenge for any econometric model which seeks to isolate the ‘scarcity rent’ arising from congestion and capacity. On the one hand, in explaining average revenue, it would seem important to control for the extent of business class patronage – so class of travel should naturally appear in an econometric model. On the other hand, if in response to congestion/capacity constraints, airlines deliberately target business patronage, then the scarcity rent measured in another part of the model would ignore this distortion. We return to this point below.

Box 4.1: Provenance of SEO/Cranfield’s Fares Data

‘Average monthly booked fares for specific pairs of origin and destination airports were derived from the Marketing Information Data Transfer (MIDT) dataset, as provided by OAG Traffic Analyser. Each average booked fare in the dataset also contains information on the published airline, as well as the points of origin and destination. The dataset also details connecting airports (if any, up to two intermediate stops) and the number of passengers. The average fare paid does not include additional charges or ancillary revenues.

The original sources of information for the MIDT dataset are Global Distributions Systems (GDSs) such as Galileo, Sabre, and Amadeus. According to ARG (2013), 44% of all bookings of major airlines were processed through GDSs in 2012. The proportion is 55% for network airlines, while low-cost carriers, that prefer direct sales, only sell 16% of their bookings through GDSs. In order to correct for that, the data provider (OAG Traffic Analyser) adjusts the market figures using mathematical algorithms based on frequencies and supplied seats in each flight sector. The reliability of these adjustments has been confirmed by Suau-Sanchez et al. (2016). Adjusted passenger bookings coverage is 100% of the market. Passenger ticket price coverage is around 40% of the market’. SEO/Cranfield (2017), p23.

Box 4.2: Provenance of Frontier’s Fares Data

‘Our source of data on ticket prices is IATA’s Airport Intelligence Service (AirportIS). This is a dataset built on IATA’s Billing and Settlement Plan (BSP) data, a consolidated payment system for participating travel agents and airlines.

This is subtly different from the similar datasets produced from a Global Distribution System (GDS), since it functions on payments rather than reservations. They nonetheless cover a similar selection of the passenger market, as the two systems work in tandem to allow the purchase of tickets through travel agents.

The IATA MarketIS product is advertised as covering over 80,000 travel agents, 400 airlines and 29 Global Distribution Systems (GDS)²⁶. In our data, the reported passengers account for 45% of the total number of estimated passengers. Our data on fares is therefore based on a minority of the total number of passengers flown. Fares for the passengers not covered by the GDS system are estimated by IATA on the basis of those that are covered, so the accuracy of our models is dependent on the internal estimating procedures of IATA. To minimize the possible effects of this issue, we have excluded routes from our sample without any recorded passengers (i.e. those that are entirely estimated)². Frontier (2019), p70.

4.4 Distinctions between modelling approaches

Using data on specific city-pair departures from across all European countries, SEO/Cranfield (2017) offered two key developments over and above Frontier (2017), as follows:

- By regressing fares on factors influencing fares, SEO/Cranfield in essence estimated a reduced form of the demand and supply relationships shown in the diagrams of Chapter 3. As a result, supply-side considerations were brought to bear in the regression. Importantly, this introduced a measure of capacity – which was lacking from the Frontier (2017) analysis – as well as concentration indexes of airlines at airport and route levels to represent the extent of competition (Table 4.1)³³.
- Importantly, it was the uplift in fares from the impact of capacity that was quantified and interpreted as scarcity rent, and not an airport origin dummy variable as per the Frontier analysis.

It does seem that Frontier’s (2017) analysis was lacking in explicit controls for supply-side differences between airports. Such controls could include measures of capacity (e.g. CUI) and/or the concentration of airlines at a given airport/on a given route, although we do note that Frontier’s latest work (2019) tested such factors as sensitivities; we return to this below.

Table 4.1: Representation of market competition or concentration in Frontier (2019) and SEO/Cranfield (2017)

	Airline competition or concentration for given route	Airport competition or concentration
Frontier	Dummy for routes served by single carrier Herfindahl-Hirschman Index (HHI) score	Proportion of seats between two cities served by airport-to-airport pair
SEO/Cranfield	HHI Number of airlines Share of dominant airline	HHI at airport Share of dominant airline at airport

4.4.1 Approach to quantifying scarcity rent: dummy variable approach versus explicit measure of capacity

Turning to how scarcity rent is derived, we note the following differences between Frontier’s and SEO/Cranfield’s approaches:

- A key feature of the Frontier (2017) approach, which was essentially replicated in Frontier (2019), was to represent fare differences between Heathrow and other airports

³³ It is conceivable that there could be endogeneity between some of these variables (e.g. scarcity of capacity could reinforce concentration), but it is the job of econometric modelling to try to expose such covariance.

via the inclusion of a dummy variable for departures from LHR, and interpret such differences as being entirely due to scarcity rent.

- In reality, this dummy variable will capture the influence of demand and supply characteristics not included in the model in addition to any scarcity rent. Such characteristics could include:
 - Differences in aero charges across departure airports
 - Differences in scarcity rents across destination airports (e.g. JFK)
 - Inherent traveller preferences for some departure airports over others reflecting differing quality of service by the airport (as opposed to different cabin mixes offered by airlines).
 - Differences in other operating costs across departure airports
- Whilst the SEO/Cranfield work is not perfect, their representation of scarcity rents via a measured variable (capacity – see Box 4.3) does seem useful to separate out the influence of this factor from other persistent unexplained factors (which is what an airport-specific dummy potentially captures).

Box 4.3: Two alternative measures of capacity considered by SEO/Cranfield (2017)

*'The **Capacity Utilisation Index (CUI)** estimates capacity utilisation relative to the 5% busiest peak hour. In other words, we derive an indicator measuring the extent to which an airport operates at the maximum capacity, an approach proposed by Berster et al. (2011). The 5% peak hour capacity of an airport is derived as follows: all the operational hours of an airport on a monthly basis are ranked in terms of the total number of flight movements, where the 1st hour is the busiest hour of the year. The 5% peak hour capacity is then defined as the capacity in the 5% busiest hour, or in other words the 95th percentile of the hourly frequency in one operational year. In order to prevent miscalculations due to limited capacity and demand during night hours, only the 16 busiest operational hours of each day are considered'. (p20).*

*'The **Average Number of Aircraft Movements** per runway allows for comparing airports on their average throughput per runway. For each airport worldwide, the number of runways was collected from data in the public domain (ourairports.com), providing detailed information of runways at all airports worldwide. To obtain a fair comparison between airports, we need to define a number of runways at an airport that can be operated simultaneously. Only paved runways of over 1500m are taken into account. We follow the simple rule that runways are considered to be independent if they are in parallel direction and the centre lines are at least 760m apart. For a subset of airports with a more complicated runway layout the number of runways that can be used simultaneously was adjusted manually. The average number of movements per day is obtained by dividing the annual number of scheduled movements by 365 and by the number of simultaneously operated runways. In turn, the average number of daily movements per runway is distributed over operating hours during daytime (as most airports are not operating 24 hours per day), by dividing the number of average movements per day by 16'. (p71)*

- The capacity impact on fares (in SEO/Cranfield) was estimated using data on city-pair departures from across Europe, whereas the Heathrow dummy variable (in Frontier) only captured differences between Heathrow departures and other London/European hub departures. Since SEO/Cranfield specified a continuous relationship, variations across city-pairs could in principle be used to estimate the marginal effect on fare of increasing capacity, holding all other factors constant including other persisting differences between Heathrow and other airports.
- There remains a question concerning the functional form of the relationship between capacity utilisation and fares. The diagrams of Chapter 3 suggest a non-linear, and possibly discontinuous, relationship between capacity utilisation and scarcity rent. This could in turn make adoption of the SEO/Cranfield approach sensitive to the choice of functional form. The extent to which this issue impacts on any particular airport depends on the extent to which the airport is at the extreme of the sample in terms of congestion. In the case of Heathrow, which is, by many measures, at the upper extreme, then the functional form could be expected to have a large impact – such that the choice of functional form becomes non-trivial. The issue is further compounded if there is excess demand (i.e. substantive unserved demand), although as noted below, many airports operate with this constraint, at least for times of the day/year.
- Finally, whilst the majority of the (and indeed the preferred) model(s) still represented capacity via a LHR dummy, Frontier (2019) also considered an explicit measure of capacity utilisation as a sensitivity test. This utilised an estimate of the level of demand suppressed by the capacity constraint – which was deemed to be above zero only in the case of LHR. As Frontier themselves acknowledged, this is not dissimilar from introducing a Heathrow indicator variable (just that the pattern of excess demand implicitly ‘fixes’ the relative magnitude of the scarcity rent observed over time rather than letting this be freely estimated in the dummy variable approach). Therefore, it is not surprising that this approach gave similar estimates to Frontier’s favoured dummy variable approach.
- Whilst it may be true that Heathrow is the only London airport where physical throughput capacity is reached for a large proportion of the day, it also seems unrealistic to suggest that other London airports do not suffer from the substantive congestion peaks that imply some scale of scarcity impact. As such, the capacity utilisation measure that Frontier adopted does not seem appropriate for the task of fully isolating scarcity rent. We return to this point when considering how the econometric approach can be improved – section 6.3.

4.4.2 What does the measured ‘scarcity rent’ in each approach capture?

In defending their adoption of a continuous function for capacity utilisation, SEO/Cranfield acknowledged that a linear relationship is simplifying, but consider that scarcity rent could arise from either or both of:

- increase in fare for a given class of ticket for a given route (which seems analogous to the definition in the Frontier work);

- distortion to the routes and cabin class allocations selected by airlines in response to constrained capacity.

On this basis, SEO/Cranfield interpreted scarcity rent as any increase in average revenue partly from the response of airlines to focus on higher yield routes and cabin classes – and such a relationship may well be better approximated by a continuous relationship.

To illustrate this issue further, Figure 4.1 below extends a diagram (Figure 40) from Frontier’s latest report (2019). In our version, we introduce some additional categories and colour-coding to show the divergence in the interpretation of what constitutes ‘scarcity rent’ in each approach.

- With reference to the yellow factor, both Frontier (2019) and SEO/Cranfield sought to isolate the potential influence of congestion on fares (i.e. scarcity rent).
- The blue factors were included as additional explanatory variables in both the Frontier and SEO/Cranfield models.
- The orange factors were included by Frontier but excluded by SEO/Cranfield.
- The green factors were excluded by Frontier but included by SEO/Cranfield.

This then provokes the question of whether such ‘product’ distortions (which might be seen as short run supply-side responses to capacity) should be considered in-scope for the measurement of scarcity rent. In other words, from a conceptual point of view, should scarcity rent be defined simply as the premium on a given fare because of capacity, or should supply-side distortion because of capacity also be factored into the measurement of rent?

A further question is whether the fare premium supposedly associated with capacity accurately estimates the welfare loss to consumers, when we only observe fares aggregated over different classes of travel, and part of the supply-side response is to ration through price/quality by offering more business seats. This question is posed for two reasons. Firstly, the traveller in business class presumably derives extra utility from travelling business rather than economy. As such, relying solely upon fare as the measure of welfare could be ignoring the benefit of substituting between business and economy. Secondly, if the overall market remains relatively stable, any distortion towards business at the capacity-constrained airport would logically lead to distortion towards economy at another airport. As such, failure to acknowledge that air travel operates as a complex system could lead to erroneous conclusions regarding the prevalence and scale of scarcity rent.

Recognising the implicit constraints with the aggregation of fares data, the above seems to provide a mandate for controlling for the extent of different cabin classes served in the regression. This should better isolate the impact of capacity constraints on fares on a like-for-like basis, and avoid concealment of any welfare distortions from changing product mix in response to capacity. This is what Frontier (2019) has tried to do and, in this respect at least, we consider this to be a more appropriate approach than that of SEO/Cranfield.

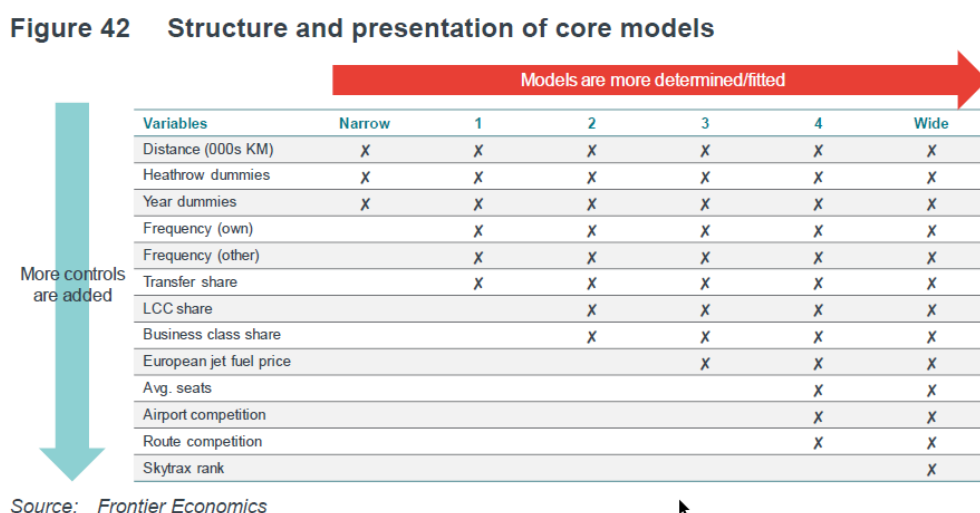
Figure 4.1: Candidate variables to explain average revenue (aka fare) per route



*% of pax visiting friends and relatives was actually included by Frontier (2017) but discontinued in Frontier (2019)

Finally, we reflect on another substantive development in Frontier’s (2019) report vis-à-vis (2017) – in that they undertook extensive sensitivity testing of their model (see Frontier’s Figure 42 reproduced below). From the perspective of the discussion above, they tested both a ‘narrow’ model specification (more like SEO, but without some key supply-side elements) and a ‘very wide’ specification (which included both SEO-type supply-side measures, but also fare mix variables). Whilst this sensitivity testing is welcome, we would advocate further testing, such as the specification corresponding to Frontier’s ‘wide’ specification less business class share and transfer share. That said, perhaps the most important finding from Frontier’s sensitivity analysis is that the estimate of scarcity rent is relatively stable across the various tests.

Figure 4.2: Reproduction of Frontier’s (2019) Figure 42



4.5 Cross-sectional data and panel data

In terms of data structure, SEO/Cranfield used panel data (on a monthly basis), whilst, in all but the most recent report, Frontier used a cross-section. There seems to be clear advantage in using panel data if this is available. It is recognised that collecting such data is not trivial, but this data structure does permit better disentangling of persistent factors that impact on Heathrow *per se*³⁴ versus those that impact on Heathrow operating close to capacity.

Turning to Frontier’s latest approach (2019), we do note that whilst panel data is analysed, there are no controls (known as effects) for time-invariant unobserved heterogeneity. Failure to include such controls ignores any persistent differences between airports or routes which are not captured by measured variables in the dataset, potentially leading to misleading estimates of scarcity rent. These controls should be tested for by the use of fixed or random effects estimators. The appropriate estimator can be established by use of a Hausman test or a Wu-Hausman test. Two potential panel stratifications could be considered. The first is origin (or destination) airport stratification – thus all routes which fly out of a given airport would have a common panel data control for all years (airport effects). Another approach is route-level panel data stratification, where each route has a dedicated control (route-level effects).

³⁴ This phenomenon is known as ‘time-invariant heterogeneity’.

5 TO WHOM DOES ANY SUCH RENT ACCRUE?

5.1 Background

If, for any given level of capacity on a city-pair at Heathrow, scarcity rent is generated, then the answer to the above question will depend upon the actions of five agents, namely:

- The passenger
- The airline
- The airport
- The regulator
- The Government

The passenger's willingness-to-pay to travel will be revealed as the demand function, driven especially by fares (of the flight in question as well as relevant competitors), income and quality (again relative to the competition).

Given capacity, marginal costs will be driven by the operating costs of the airline, in tandem with aeronautical charges levied on the airline by the airport (though the latter are tightly regulated by regulator to take account of the airport's substantial market power in this context) and taxes levied by the Government.

If marginal costs (incorporating aero charges plus other running costs incurred by the airline) were at the efficient (i.e. market-clearing) level (e.g. $P_B = MC_B$ in Figure 3.1), then no scarcity rent would be realised by the airline. However, the actions of the regulator serve to cap aero charges at below the marginal cost of accessing the airport. Since there is no incentive on the part of the airline to similarly cap air fares – and in any case, some mechanism is needed to manage excess demand – then an excess of P over MC (i.e. scarcity rent) will in principle arise.

On this basis, it seems clear that – on a strict definitional basis – any scarcity rent will accrue to the airline and not to the airport or Government. However, the strict definition may fail to give a rounded picture of the total rent and its distribution. Since the MC function is not readily observable to anyone outside the airline, the practical questions are really:

- what portion of fare is covering costs (e.g. the orange area in Figure 3.1) and what portion is rent (i.e. the yellow area)?
- within costs, what are the relative proportions of airline operating costs, aero charges and taxes – since this dictates the revenue shares claimed by the airline, airport and Government, respectively?

Of course, there will always be argument about what shares of revenue should be claimed by the airport and Treasury, although in their defence, these parties are subject to regulatory and/or public scrutiny. Notwithstanding this point, which we will return to in the next section, the barrier to answering the above questions is that, whilst we can (with qualification) observe

the volume of sales and fares paid (giving us insight on total revenue), we cannot similarly observe the airline's cost base (which would give us insight on total cost).

5.2 Rent capture by the airline

What can be said with confidence is that where an airline is less efficient than it could be (say through operating an ageing fleet with relative high fuel consumption, or where labour is less flexible (e.g. due to historical labour practices) and/or more expensive (e.g. due to pension fund legacy) than it could be, then scarcity rent (yellow area) will be relatively smaller and the other revenue covering costs (orange area) relatively greater, all else equal. In other words, some degree of 'scarcity rent' could be absorbed as cost inefficiency on the part of the airline – as opposed to being realised in the form of profit. Indeed, should this situation exist to the extent that an airline is returning little or no profit, then the 'capture' of potential scarcity rent by the cost base should not necessarily be seen as a defence of the airline's position. This is because, in some contexts, it is not inconceivable that an airline would view scarcity rent as a 'buffer' to the cost base, and this could possibly weaken the incentives to invest and/or innovate with the objective of improving efficiency.

5.3 Rent capture by the airport and government

Sometimes the argument is made that the actions of the airport and Government via charges and taxes serve to capture some portion of scarcity rent – contributing to higher air fares if these are passed on to the passenger. In principle, any such actions will be reflected in terms of MC rather than the rent per se – but since any rent is the residual of revenue after subtracting costs, this is really just a different way of looking at the same thing.

In determining the validity of the above argument, the practical question is whether there exists a relationship between charges/taxes and airport capacity. SEO/Cranfield (2017) investigated this question by regressing taxes and charges against the Capacity Utilisation Index (CUI) for outbound intra-European and intercontinental flights from a sample of European airports, with the results suggesting *'that governmental taxation may be capturing some of the scarcity rents through taxation, but that airports are less likely to capture rents through their charges'* (p51). SEO/Cranfield further investigated the existence of a relationship between charges and air fares, but this proved inconclusive.

Whilst qualifying that the SEO/Cranfield analysis was focused on European rather than UK departures specifically, the aforementioned analysis lends some support to the argument that the Government may in practice be capturing some portion³⁵ of scarcity rent at LHR via taxes, but offers no support for that the argument that the airport may be similarly capturing scarcity rent via charges (which is perhaps to be expected given the regulation of charges by CAA).

³⁵ It should be acknowledged that taxes such as Air Passenger Duty (APD) apply across the majority of UK airports (the exceptions are Northern Ireland and the Scottish Highlands and Islands region), including those with no capacity constraints, and not just LHR. It would be difficult to argue that APD is *intended* to capture scarcity rents.

5.4 Synthesis

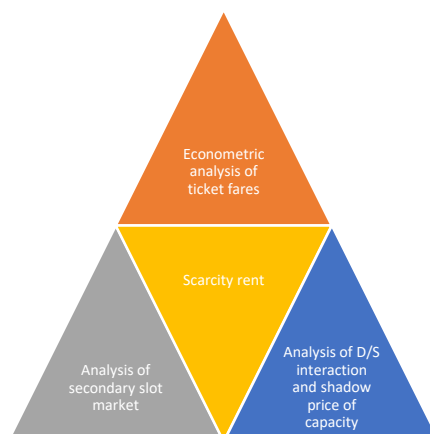
More generally, it is perhaps necessary to distinguish between two things. The first is whether the airport operator is earning supernormal profits out of its regulated activities (aero charges) or its unregulated activities (shopping, parking etc.) and whether the regulator is doing a good job in setting the regime. The second is whether there are scarcity rents in the system. On the second point, we would make the following observations: (a) slot pairs at LHR do change hands for appreciable sums in the secondary market, admittedly a thin market; (b) a key concern of the Airports Commission was the future likely trend in scarcity rents if air traffic grows by x per cent per annum and London capacity does not grow; (c) in constrained markets, rents might be accruing to airline inputs (e.g. labour) as well as to airline owners.

6 WHAT CAN BE DONE TO STRENGTHEN ANALYSIS OF SCARCITY RENT AT HEATHROW?

6.1 Background

Broadening the scope of Frontier's (2017) analysis of scarcity rent, Frontier's (2019) analysis of scarcity rents no longer focussed attention solely upon econometric analysis of fares – but instead entailed a 'triangulation' between three approaches: i) econometric analysis of fares; ii) analysis of demand and supply interaction and the shadow price needed to 'choke off' excess demand; iii) slot price analysis. The econometric analysis itself was somewhat more developed than their previous effort, by employing a panel dataset across several years as opposed to a single year, and seeking to represent capacity explicitly within the econometric model instead of relying upon a LHR dummy as a proxy for capacity. Whilst i) remained the primary analysis, Frontier positioned ii) and iii) as alternative analyses of scarcity rents that serve to corroborate i).

Figure 6.1: Triangulation of methods for analysing the prevalence and scale of scarcity rents at Heathrow



The rationale behind the econometric analysis has been covered in some detail in Chapter 4. The rationale behind the other two approaches can be summarised as follows:

6.1.1 Analysis of secondary slot market

In simple terms, this line of analysis involved the following steps:

1. Gather available data on the price paid for Heathrow slots in the secondary market.
2. Recognising that slot prices will vary by time of day, calculate an average slot price.
3. Using a central discount rate of 7.5%, calculate the discounted cash flow over a given asset life which exactly recovers the slot price paid at the outset.

In essence, Frontier interpreted the cumulative cash flow as the scarcity rent. Justifying this interpretation, they asserted: '*...economic logic suggests that the upfront price an airline will be prepared to pay for a slot will not exceed the expected net present value of the annual premium of income from the use of the slot compared to fares in an uncongested market. In our analysis, we assume that the net present value of the cashflow premium on each slot is*

equal to the upfront price paid, i.e. that the airline is just breaking even from the slot transaction' (p44).

6.1.2 Analysis of D/S interaction

In simple terms – and drawing reference to Route B in Figure 1 as an illustration – this line of analysis involved the following steps:

1. Forecasting the unconstrained demand (i.e. Q_B^U in Figure 3.1)
2. Forecasting the constrained demand (i.e. Q_B)
3. Subtracting 2. from 1. to get the excess demand (i.e. $Q_B^U - Q_B$)
4. Inferring what price increase, from the unconstrained demand situation (i.e. P_B^U), would be needed to 'price off' the excess demand
5. Inferring the scarcity rent as the difference between the constrained and unconstrained prices (i.e. $P_B - P_B^U$)

6.2 Strengthening the triangulation approach

From an intuitive perspective, we feel that each of methods deployed by Frontier has potential merit. However, we are also of the view that each of the methods faces its own set of limitations, and for this reason we strongly advocate some form of 'mixed methods' approach with 'triangulation' of the insights gleaned from each method.

We are supportive therefore of Frontier's broad approach, but this is not to obscure misgivings that we harbour in relation to how Frontier has implemented the individual methods. Chapter 4 has covered Frontier's econometric method in some detail, and it is outside our remit to conduct a similarly detailed critique of the secondary slot market and D/S interaction approaches. However, as an indication of our concerns we would highlight the following fundamental points.

6.2.1 Secondary slot market

We are not persuaded by Frontier's assertion that the slot price paid can be directly interpreted as the scarcity rent. In particular, they conducted a discounted cash flow analysis where the slot was assumed to have zero residual value. Whilst the latter obviously represents the lower limit, we believe that, even when capacity is increased at capacity-constrained airports, slots will embody some degree of residual value (i.e. they will carry some inherent value as a tradeable commodity, in much the same way that London taxi licences continue to be traded at substantial value despite the advent of Uber). Thus, the key to understanding slot prices is to dissect the potential for generating fare revenue from the residual value of the slot per se.

At the very least, we would advocate sensitivity testing – another test being the assumption that the residual value equals the upfront price in real terms (i.e. the use value is the interest charge only at say 5% real or 7.5% real, as appropriate). On this basis, a premium slot pair price of £50 million at 5% real would imply a use value of £2.5 million. Spread over 350 days and 600 seats per day gives a use value of £12 per ticket (which could of course be £5 for economy and £50 for club class). Other tests could consider an intermediate scenario where LHR expansion imposes a degree of downward pressure on the residual value.

Notwithstanding the above concern, Frontier's analysis is reliant upon a very limited dataset of secondary slot prices, and for this reason alone, the results should be interpreted with caution.

6.2.2 D/S interaction

Whilst we are comfortable with the essence of the approach, Frontier's implementation is too simplistic. In 'pricing off' demand at LHR, passengers have in principle two options: either switch to another airport (in London, the UK more generally, or another European hub) or not travel by air. Once a given shadow price is imposed on passengers using LHR, this will in practice change the relative generalised costs on all journey options (i.e. other airports or not travelling by air) available to all passengers, and the system must therefore be allowed to iterate until it reaches an equilibrium in relation to D/S at all significant airports (including LHR) within the system. In other words, Frontier has conducted a partial analysis focussed upon LHR, which takes no account of the impact of imposing a LHR shadow price on the behaviour of passengers across the system, and the assignment of passengers has then been halted at the first iteration. It should be highlighted that the DfT's NAPDM/NAPAM modelling framework follows essentially the same methodological approach – but is considerably more complex than Frontier's approach, and correctly views the air passenger market as a system of D/S interactions.

6.3 Strengthening the econometric approach

If theory and intuition suggest that there could be scarcity rent at LHR, then we do need an empirical framework to test this proposition and quantify the magnitude. What emerges from the discussions in the literature is that, partly because of data limitations, there is no single definitive way of going about this. Instead, the likelihood is that we must develop partial insights using a range of methods, and develop an overall view by triangulating the various results.

We do consider the econometric approach of Frontier to have merits as part of the evidence base. Indeed, coupled with the insights in the SEO/Cranfield work, the modelling could be re-examined and enhanced particularly by including measures of the supply-side of airports. The latest work by Frontier (2019) does represent a move towards that through inclusion of more demand- and supply-side variables. What is missing however is an adequate measure of capacity utilisation in Frontier's model and we return to this in section 6.3.1. As discussed earlier in Chapter 4, we do not believe that Frontier's (2019) sensitivity analysis using excess demand is an appropriate quantifiable measure of capacity utilisation, as due to the way Frontier constructed it, it is only ever positive for one airport – Heathrow.

In tandem, there needs to be clearer statement of what any increase in fare associated with congestion actually represents. The variables listed in the table below, drawing on those contained in the Frontier and SEO/Cranfield analyses, would potentially inform the following interpretation of the CUI coefficient:

The coefficient on the capacity variables would indicate the fare premium (or otherwise) at an airport on a like-for-like basis after controlling for factors that differ at the airport compared to other airports outside of the controls of airlines, holding airlines' fleet and cabin layouts and the competitive concentration of each route market

constant. As such, the measured fare differences exclude supply-side responses to capacity constraints.

The reason for giving such a lengthy and careful definition is because we observe only coarse aggregations within the existing fares data (across different classes of travel etc.). It may seem counterintuitive to try to strip out the supply-side response, but without such controls there is no way to disentangle whether the observed fare change as capacity increases is due to an increase in a like-for-like fare versus a shift to a higher 'class' fare – which is the implied supply-side response.

Moreover, given the use of an aggregated revenue per passenger measure as the dependent variable, all we can ever hope to yield from these models are quantifications of the response of fare to:

- A pure congestion/capacity effect as defined above (holding airlines' fleet and cabin layouts and the competitive concentration of each route market constant);
- A product differentiation effect (i.e. the fare impact associated with a shift to more business passenger share);
- A competitive market effect (i.e. the fare impact associated with greater or lesser market concentration).

Whilst the pure congestion/capacity effect is naturally central to our interests, we believe that, in order to properly understand the prevalence and scale of scarcity rents at LHR, greater emphasis needs to be given to the roles played by the product differentiation and competitive effects.

6.3.1 What should a capacity index capture?

It is important to note that any capacity utilisation index will reflect both the peak and shoulder peak periods, because the definitions of those periods will themselves be a function of capacity utilisation. In other words, if hard capacity constraints are encountered during the peak, then this will serve to extend the peak period.

We consider that the key way that the econometric analysis can be developed further in the short term is to invest in the development of empirical measures of capacity utilisation. More specifically, we consider that the following features of a capacity index are essential to robustly estimate scarcity rent:

- The index should be able to capture instances of 'over-capacity'. On the face of it, these might appear to be the cases of most interest, but it is not clear that in, say CUI, they are captured by CUI=100%. Indeed, the nature of 'peak spreading' implies that <100% could still be consistent with the existence of scarcity rents.
- Likewise, the measure should not just focus on strict over-capacity, but be able to capture operational disruptions associated with being at or close to capacity for short periods of the day.
- A continuous measure is only useful within a certain range, as below the lower limit congestion/scarcity impacts will be minimal. This raises the question, which would

ideally be answered *a priori* to estimation of any model, of what is a sensible lower limit in terms of impacting on practical airport operations?

- Any scarcity rent will be associated with airport pairs – and that could stem from congestion and capacity constraints at either (or both) airports in the pair (assuming fares are pooled across directions). This would suggest that a capacity index needs to reflect conditions at both airports. SEO/Cranfield employ a capacity index averaged across airport pairs – but this could mask significant under/over-capacity at either member of the pair. Instead, whilst a one-for-one weighting might not be appropriate, a measure which reflects the overall scarcity constraints would be preferable. In the extreme (counter-example to SEO/Cranfield), a maximum capacity for the pair could be used, although this could still be enhanced by considering additional constraints arising at the other airport e.g. London Heathrow to New York JFK.
- With respect to the above, a broad stakeholder consensus should be to establish the most appropriate (and feasible) measure of capacity and the appropriate operational threshold.

6.3.2 What variables should be included in the model?

Based on the discussion in section 4.4.2, we consider that a suitable starting point for the specification of variables to be included in a fares regression is as given in Table 6.1.

Table 6.1: Suggested variables to explain fare differences between routes

Variables in the Frontier analysis (to be retained)	Variables in the SEO/Cranfield analysis (to be added)
Distance between origin and destination	CUI
Indicator variable for a long haul route	HHI indexes at airport and route level or lead airline market share as determined empirically
Frequency of service from airport	Fuel price (representing costs)
Frequency of service from other airports	GDP summed over both the origin and destination country
% of pax business class	Sum of the hinterland population around the origin and destination airport
% of pax visiting friends and relatives ³⁶	Hub airport indicator variable (one if either the origin or destination airport is a hub)
% of transfer pax	
% of LCC pax	

³⁶ We note that this variable was not included in Frontier (2019), but was included in Frontier (2017). We consider this to be a useful characterisation of travel demand

We recognise that some of the measures may be refined given empirical findings, however these broad categories represent a useful starting point. We suggest that the inclusion of these variables should be tested using a general-to-specific modelling strategy. That is to say, the modelling should begin with as many of these variables included as possible, and then the model should be reduced depending on the statistical significance of each variable. This is because we recognise that some of these variables may pick up similar aspects of the determinants of fare and so should be tested empirically.

In addition, we suggest that panel data is used and that route-specific random effects are included. The alternative fixed effects estimation approach should be tested using a Hausman test to verify that the random effects approach is the more appropriate.

Finally, whilst analysts should present a 'preferred' model, they should also test the sensitivity of the estimated scarcity rent to different model specifications.

6.4 An alternative approach: analysing generalised cost instead of fares

We believe that the market for air trips could be characterised not by considering the quantity/fare relationship, but instead by considering the quantity/generalised cost (GC) relationship³⁷. This could be a very useful device if the market equilibrium service quality changes when approaching capacity. In particular, there is evidence to suggest that airlines do adjust their seating mix to favour higher quality services such as business passengers in response to capacity constraints. There is both greater cost of provision and greater value to consumers from this shift. As such, it makes sense to consider not just the fare paid, but also the implicit valuation of service quality in the analysis. Instead of fare, a GC index could be constructed and used as the dependent variable in a regression. This could in turn help to reduce the ambiguity in what the regression results show, as the impact of changing passenger mixes towards business would be accounted for not as a regressor, but as part of GC. Whilst we here offer initial thoughts on this proposition, it would be prudent to explore the scope for a GC-based approach through a small feasibility study.

³⁷ Indeed, this is the essence of the approach promoted in the Passenger Demand Forecasting Handbook used by GB railways (although strictly speaking, they use a combination fare and generalised journey time (GJT) as the measures of price and quality respectively). Furthermore, it should be noted that DfT's aviation modelling suite employs GC rather than fare.

7 SYNTHESIS AND RECOMMENDATIONS

7.1 Introduction

The prevalence and scale of scarcity rents at Heathrow and the potential for such rents to grow over time if demand grows relative to capacity is an important policy issue. From the perspective of the air traveller, whose interests the regulator represents, a key question is whether, if capacity is expanded at Heathrow, the air travel offer out of the London airport system in terms of fares, service quality and accessibility will improve. An ingredient of this is how much air fares will fall due to a reduction in scarcity rents, and whether that fall will more than offset any increase in fares from increases in aero charges required to fund the capacity expansion.

Several consultants' reports have been reviewed against a background of relevant applied economics literature. The headline conclusion is that there is going to be no silver bullet answer to the policy question and that the size of any price premium due to scarcity rents is going to be a matter for evidence-based judgement and interpretation. We have some suggestions for improvements to the econometric analysis which are worth trying, but we do not expect these to be game-changers. What might be a step change would be better, more disaggregate ticket price data such as LENNON data in the GB rail industry, on which better econometric models could be constructed. But our standpoint in this piece of work is that we are where we are, working with what we have on theory and evidence.

7.2 Theoretical insights

From a theoretical perspective, we are content with the Frontier (2017, 2019) and FTI (2018) reports' exposition of the market for seats as a tool for illustrating the effect of capacity limits on ticket prices. We have developed that apparatus slightly further in Chapter 3 above. There is clearly scope in principle for scarcity rents to exist where airports are at capacity and for them to grow as demand growth exceeds supply growth. However, we would counsel against using the vertical supply at capacity construct as a representation of reality. The demand and supply elasticities close to capacity are empirical matters which bear on the size and growth of scarcity rents.

We have two key points to make from a theoretical perspective. The first is to challenge the implicit proposition that the product is homogeneous. So, for example, a network airline might be offering a different product at least to part of its market from a point-to-point airline serving the same route. In that situation, price differences are not sufficient to indicate rents, because rents are prices less costs. The analysis would need to consider P-MC differences and changes/variations in them. This requires analysis of the cost-side as well as the demand-side.

Secondly, the proposition that rents are wholly associated with slot scarcity seems dubious. The literature suggests that parts of the market are oligopolistic in nature, in other words that there is market power. If the market is characterised by firms setting output then price, the Cournot model seems relevant, and in Chapter 3 we provide illustrative examples of the implications for P-MC margins using the IATA elasticities. Having said that, it is also likely that slot scarcity at hub airports is itself a significant source (among others) of market power, and

that greater hub capacity would dilute patterns of slot dominance and increase competitiveness. So slot scarcity rents and rents from other sources are not independent entities.

7.3 Empirical insights

Turning to the empirical evidence, our overall view is that both Frontier (2019) and SEO/Cranfield (2017) are pieces of work which should be taken seriously, and should form part of the evidence base on which informed judgement should be made. Although outside our terms of reference, we would add that the shadow costs in the DfT's NAPAM model is another relevant source of evidence to consider alongside. The key points about the econometric work are:

- The fares data is average route revenue, when ideally, finer resolution by time of day, business/economy etc. is desirable.
- The pattern of airlines' offer is influenced by the level of congestion and capacity restriction at hub airports. Airlines will naturally gravitate towards the least elastic traffics and the mix of what they offer in terms of service classes becomes endogenous. More disaggregate fares data would help the analysis.
- The SEO/Cranfield approach of representing scarcity rents by a measured variable (capacity utilisation) is useful in separating out the influence of this factor from other unexplained factors – we certainly prefer this approach to the use of airport specific dummy variables as proxies for capacity. However, SEO/Cranfield's capacity variable does not reflect the extent of excess demand, and may not therefore give a reliable measure of capacity utilisation³⁸.
- There is a clear advantage in using panel data rather than cross-sectional; if further work is done, we would consider the use of a panel with airport dummies constant over time.
- The fares data is travel agency ticket sales data which only accounts for about half the market, although there is some adjustment of the raw data to reflect sales across all channels. If work based on this data is to be used for policy purposes, we would recommend an audit of its strengths, weaknesses, representativity, fitness for this purpose etc., so that an informed judgement can be made of the weight it can bear.

In a ticket price of £x, rents accrue to a combination of Government, airports and airlines, and the factors they employ, as a product of regulation, taxation and market power. The position of scarcity rents has to be seen in that context, but with regulated aero charges one would expect any scarcity rents to accrue to airlines (which is why they are prepared to pay for slots). The question of whether such rents accrue to the airlines themselves or partially to factors of production in imperfectly elastic supply or to so-called legacy actors is not one we can answer.

³⁸ To be explicit, it is not the inclusion of airport-specific dummy variables that is problematic – since this is the essence of the fixed effects approach – but rather the absence of further effects (notably capacity) that, if significant, will be confounded with fixed effects.

We agree with the approach taken by Frontier (2019) in adopting a triangulation approach built around analysis of fares data, analysis of slot trading including their accounting treatment, and shadow cost approaches such as NAPAM which have the advantages of representing the London and national airport system and including some geography to represent the changing boundaries between airports as demand grows. The value added comes from interrogating each of the approaches in the light of the others – what is the interpretation of the results for scarcity rents and why is the result from approach A an order of magnitude different from approach B?

Finally, in policy terms, we would draw attention to the need somehow or other to forecast forward from the analysed position in the base year, say 2018, to the position in 2030 and 2040 during the life of the airport system both with and without enhanced capacity. Building a more secure description of the base year is one thing; using the model in predictive mode is another.

7.4 Recommendations

How to move forward from the position reached in summer 2019 is partially dependent on the timelines and decision points within which the scarcity rents issue is being addressed. But we think the following commissions should be considered by the regulator:

- Further work to enhance the existing regression analysis by Frontier (2019) by better treatment of congestion and capacity as outlined in section 6.3.
- A review and audit of the fares data, its representativeness of the market as a whole, and its fitness for purpose in this application, together with any recommendations for improved data, relevant disaggregation and validation.
- A small scoping study to assess the scope for re-couching the econometric model in terms of generalised cost rather than fare, especially with a view to achieving better representation of supply-side responses to congestion.
- Further work on the airline cost side to help determine whether price differences are partly a reflection of cost differences.

Alongside these suggestions for further commissioned work, the core of the way forward lies in policy work to develop the triangulation approach used by Frontier (2019) and described in Chapter 6 above. It would be useful if the fares approach and the analysis of the slot market were taken forward. More attention could be given to the insights from the NAPAM model since this incorporates demand/supply interaction which is helpful in representing what happens close to capacity. A crucial question for policy is how to forecast scarcity rents forward to 2030 and 2040. Sound judgement and interpretation of the different strands of evidence and their robustness will be essential.

ANNEX A: RESPONSES TO THE FRONTIER, FTI AND IAG CRITIQUE

This section summarises the substantive arguments presented by the various parties and gives our brief response in each case. In what follows, these arguments are organised around four key policy questions.

A1 Given the current regulatory regime in the UK, does theory suggest that there could be a scarcity rent at Heathrow?

Argument IAG-1: The focus of competitive activity is around city-pair markets.

- IAG: Airlines compete to transport passengers in city-pair markets. FTI notes that the *'Frontier report inherently considers Heathrow airport to be a single market'*. FTI further remarks that *'with numerous airlines operating out of Heathrow, and not all airlines serving the same routes, air travel to and from Heathrow is perhaps more accurately represented as a collection of individual markets'*. This implies that the assessment of airport congestion should focus on airline competition on city-pair markets. This also implies that such an assessment should take into account the extent of congestion in the greater London area, as airlines are able to add seats from other London airports.

Response: We broadly agree with IAG's proposition and this is reflected in our rationalisation of the problem in Chapter 3.

Argument IAG-2: If there is constraint on slot capacity, it does not necessarily follow that there is constraint on seat capacity.

- IAG: ...the shortage of available slots does not imply that airlines cannot offer more seats on city-pairs on which they compete. Indeed, the FTI report notes that *'passengers purchase seats, whereas the direct constraint in the market is on the slots available to airlines'*. In its report, Frontier (2017) fails to establish a clear link between the shortage of slots and a shortage of seats on a given city-pair market, nor does it explain how such a link could work in theory. As will be discussed in Section 3, taking into account that the capacity constraint is on slots and not on seats has important implications on (i) the extent to which airlines' seat capacity on city-pair markets is constrained; and (ii) the extent to which airport congestion gives rise to higher ticket prices to passengers.

Response: At face value, we do not disagree with IAG's proposition – although in practice, the supply of airline seats will be 'lumpy', since it will be conditioned by the supply of slots and planes (of varying sizes and configurations) – both of which are inherently discrete in nature. However, even if one were to accept IAG's proposition, this does not invalidate the Kreps & Scheinkman (1983) result. In our view, there is good reason to believe that the constraint on slots is intrinsically linked to the operation of the market for seats – and this is reflected in our rationalisation of the problem in Figure 3.4.

Argument IAG-3: If there is constraint on slot capacity, it does not necessarily follow that there is pressure on ticket prices.

- IAG: Slots are a necessary input for airlines to serve passengers on city-pair markets. Airlines have to obtain slots for landings and take-offs in advance (at least at slot-controlled airports). This is particularly the case for city-pair markets involving London, where airlines have to use a slot that they may request from slot coordinators at the various London airports. When slots are in shortage, this may affect airlines' operations, but this would not give rise to a ticket price increase, contrary to what Frontier (2017), and to a lesser extent FTI, predict. There are several reasons why Heathrow congestion is unlikely to cause ticket prices to go up:
 - Even if Heathrow runs at full slot capacity, the supply of seats on a given city-pair market can be increased, by increasing the number of seats per flight; by using or acquiring additional slots to serve that city-pair market; or by using slots at other airports;
 - While slot congestion may raise the cost to airlines of using slots, this cost increase is not expected to be borne by passengers through higher ticket prices.

Response: In our view, there is good reason to believe that each of the aforementioned supply-side responses to capacity constraint would impose some cost on the airline. As illustrated in Figure 3.3 for the case of $P=MC$ (although the same effect would apply analogously to $MR=MC$), theory dictates that this should manifest as an upward shift in the supply function such that, all else equal, price will increase.

Of course, if price were to remain at MC after the acquisition of additional capacity, then there would be no contribution to the fixed costs of securing that capacity – which would ultimately be unsustainable.

The social welfare outcome will depend on the net effect of two countervailing demand- and supply-side forces – namely the stifling effect of the price increase on demand versus the increase in capacity. We think it is not either-or; slot constraints at Heathrow will lead to a combination of larger plane sizes, squeezing out less profitable routes to other airports and increased scarcity rents.

Argument FTI-1: Airlines gain from network benefits, and this will encourage them to serve a range of routes even if some routes are less capacity constrained than others.

- FTI: Airlines are likely to gain from the network effects associated with serving multiple routes. Suppose a given airline has two slots, and can choose to assign:
 - both slots to serve Route A (for example, JFK – Heathrow);
 - both slots to serve Route B (for example, Heathrow – Hong Kong International Airport); or
 - one slot to each of Route A and Route B.

Network economies of scale can be said to exist when an airline that operates on both Route A and Route B (Airline 1) faces more favourable demand or cost conditions than two other airlines (Airline 2 and Airline 3) serving Routes A and B separately.

- FTI: Airline 1 could time its Route A and Route B flights such that a passenger intending to fly both routes (in the example above, from JFK to Hong Kong via Heathrow) would be more likely to purchase tickets from Airline 1, rather than purchase one ticket from Airline 2 and one ticket from Airline 3. Airline 1 could also schedule its Route A flights at times that would be unpopular for Route B passengers, allowing it to optimise its flight schedules. Airline 1 therefore gains an advantage from serving both Route A and Route B. Even if, say, Airline 1 observes that Route A has higher scarcity rents associated, then it may not choose to switch both slots to Route A since it would lose the benefits of the network effects described above. Therefore, it is unlikely that airlines at Heathrow would make their route allocation decisions based only on consideration of scarcity rents on each route. Indeed, it is likely that airlines at Heathrow (and in general) make their route allocation decisions based on a variety of factors that affect profitability, and not simply on the scarcity rents on each route.

Response: FTI's argument is an intuitive one rather than a theoretical one – but we find it plausible. Drawing more specifically from theory, an additional argument would be that network structures bring the potential for cost economies – especially for the major hub operators with substantial fixed assets in situ at the likes of Heathrow (e.g. aircraft maintenance, baggage handling, etc.). Both arguments concerning network effects would encourage some airlines (especially those operating out of the major hubs) to serve a range of routes rather than 'cherry pick' those which generate a rent. However, there are many situations; if Route A is Hong Kong to Heathrow and Route B Heathrow to Leeds/Bradford, there comes a point at which the scarcity rent forgone is too high to justify the short haul flight; one of the issues is the alignment between airline commercial interests and broader social interests in a heavily constrained situation.

Argument FTI-2: The regulatory framework may restrict the ability of airlines to target routes which generate a rent.

- FTI: We understand from discussions with the CAA that airlines themselves are constrained in route choice. While there are some regions which have adopted a more liberal approach to market access in air services (such as the intra-EU single aviation market, or the EU-US agreement that any EU or US airline can serve any route between the EU and US), the international norm is for countries to agree air services bilaterally. We further understand that these Air Services Agreements (ASAs) vary greatly in the restrictions they place on market access. Virtually all confine access to airlines which are majority owned and controlled by nationals of each country (hence the prevalence of single national 'flag' carriers in some markets). The most illiberal ASAs restrict: (i) the number of airlines that each side can designate; (ii) the routes they can fly; (iii) the frequency at which they can operate; and (iv) the fares they can charge. In such markets, therefore, an airline may not be able to serve a route even if it wished to.

Response: The aforementioned features of 'the most illiberal' ASAs would have different implications for the social welfare outcome: (i) could discourage competitive pressure, such that quantity may be lower than otherwise, and price higher; (ii) could have the same effect; (iii) could constrain capacity, such that quantity may be lower than otherwise, and price higher; (iv) could limit the scarcity rent arising in any given sub-market, but at the same time could

provoke a supply-side response on the part of the airline, perhaps encouraging them to re-direct capacity elsewhere.

Argument FTI-3: Scarcity rents can exist even when airlines are free to choose which routes to operate on if the airport in aggregate is capacity-constrained.

- FTI: IAG's worked example of an airline deciding to switch from one route to another fails to consider if its initial route was capacity constrained in the first place. Suppose Airline 1, operating on Route A, observes large scarcity rents being earned by Airline 2 on Route B, due to Route B's capacity constraint. Following IAG's logic, Airline 1 has an incentive to shift its capacity to serve Route B, competing away the scarcity rents earned by Airline 2. However, suppose Route A is similarly capacity constrained, and therefore earns some scarcity rents. If the whole scarcity rent at Route A is greater than the proportion of the scarcity rent it is likely to earn at Route B, Airline 1 has no incentive to switch to Route B. One could generalise this example (and continuing to follow IAG's logic of quantity competition) to Heathrow's case of many airlines and many routes. In principle, in equilibrium, the scarcity rents on each route would equalize. Therefore, scarcity rents can exist even when airlines are free to choose which routes to operate on if the airport in aggregate is capacity-constrained. As an example, we understand from the CAA that the London to Cairo route is served by just two airlines: British Airways (daily) and Egyptair (twice daily). The ASA limits airlines of each side to no more than 14 round-trip services per week. Therefore, new entry on this route is currently confined to an airline designated by the UK (which must be EU-owned and controlled) operating up to seven services per week.

Response: We agree with this argument – although we would put things differently. In assessing the social welfare impacts of capacity constraint – and the impacts of relaxing that constraint – it is vital to conduct the analysis in the round – comparing the Do-Something with the Do-Minimum across all relevant sub-markets, and disaggregate the impacts by all relevant agents, i.e. consumer surplus, producer revenue, scarcity rent and deadweight loss.

Argument Frontier2019-1: London Heathrow is highly differentiated from other London airports and this will contribute to scarcity rent.

- Frontier (2019): If passengers (and airlines) were to view the other London airports as almost perfect substitutes for Heathrow, then airlines would not be able to set higher prices at Heathrow, however large the theoretical demand to use the airport, because at the application of a tiny premium demand would simply switch to other London airports with spare capacity. But if passengers (or airlines) consider the airports to be significantly differentiated then it will take a much larger premium on fares for traffic to be diverted to alternative airports. In practice there are strong reasons for expecting to see a degree of this differentiation.

Response: We broadly agree with this argument, but suggest that this highlights our propositions that: 1) factors other than scarcity could contribute to rent (or, put differently, scarcity could be confounded with other factors that give rise to rent) and 2) any scarcity rent at LHR should be analysed in the context of the London (and possibly) GB demand/supply system as a whole.

Argument Frontier2019-2: The underlying driver of scarcity rent is that constrained capacity obstructs free entry and exit.

- Frontier (2019): Competition constrains prices to cost. It is the fact that free entry and exit is obstructed by the capacity constraint that allows prices to rise. For this reason, arguments that any observed premium at Heathrow may represent a different valuation of its service may be true, but it is wrong to conclude therefore that this premium would persist if the capacity constraint were alleviated. The higher value, the differentiation, needs to be there for a premium to be charged. But it can only persist if a capacity constraint obstructs the process of competition between airlines.

Response: Whilst lack of capacity will certainly be an impediment to competition, we do not believe that increased capacity will in itself guarantee competition – since there exist other barriers to entrants (e.g. slots at destination, monopoly or oligopoly power of the incumbent(s), etc.). That is to say, the introduction of additional slot capacity may impose downward pressure on any rent, but it remains to be seen whether the rent would be eliminated entirely.

Argument Frontier2019-3: In the event that additional capacity is not added at LHR, competition from other London airports will have a limiting effect on further growth in scarcity rent at Heathrow.

- Frontier (2019): It is unreasonable to think that any congestion premium would simply continue to rise exponentially over time if capacity at Heathrow remains unchanged. If the size of the premium is limited by the extent of the differentiation between airports then there must come a point where the premium is large enough to make services from other airports competitive, trading off passenger or airline preferences. For instance, hub economics may make airlines wary of providing many long haul services from point to point airports. But if the fares at Heathrow were to rise sufficiently there could come a point that makes those services profitable. Note though that fares in total would continue to rise, the 'premium' between Heathrow and elsewhere might not. Indeed this shows that, if the extent of the constraint at one airport becomes large enough, fares may start to rise at other airports as well.

Response: We agree that, if capacity at LHR were to remain unchanged, then competition from other airports would have a limiting effect on the rent that accrues at LHR. This highlights the importance of looking at the London (and GB) air passenger market as a system – rather than LHR in isolation.

Argument Frontier2019-4: Airlines have limited flexibility over their aircraft choices in the short to medium term, based on their fleet mix and seat configurations – and this inertia will create conditions where scarcity rents could accrue.

- Frontier (2019): It is often noted that the constraint at an airport like Heathrow is in aircraft movements, not in seats, but the argument here about a premium relies on there being a shortfall in seats. This is correct, but the reality is that airlines have limited flexibility over their aircraft choices in the short to medium term, based on their fleet mix and seat configurations. And while average load factors are less than 100%, there are limits to how much these can be increased given the need to maintain operational frequencies and the fact that load factors at peak times are often close to 100%.

Response: This argument is consistent with our distinction between the ‘very short’ and ‘short’ run in Chapter 3.

Argument Frontier2019-5: Scarcity rents will be variable across routes, short vs. long haul and time of day.

- We recognise that talk about a single ‘congestion premium’ involves some simplification of reality:
 - Routes: In practice, not all routes may be constrained, and some may be more constrained than others. The congestion premium will vary from route to route depending on the precise levels of capacity and demand on each individual route.
 - Switching: Airlines can replace some short haul flights, which are typically operated with smaller aircraft, with long haul flights, which are typically operated with larger aircraft. This increases the total number of seats flown at the airport level. But while this increases capacity on the long haul route, which helps to ease congestion on that particular route, it can only be achieved by reducing capacity on the other route, where congestion is therefore worsened. Within a capped system, switching of this kind can only go so far. Also, at a hub airport, airlines need to strike a balance between short haul and long haul flights. Therefore, switching of this kind may also be to the detriment of an optimal hub and spoke model.
 - Time of day: On a given route, there may be differences by time of day too. Passengers – and by extension airlines too – typically prefer to travel in the morning and return in the evening. Airports therefore see a peak in demand in the morning, followed by quieter spells during the late morning and afternoon and a second peak in the evening. This is a pattern seen at most airports. At Heathrow however, because it is fully constrained, it is effectively operating at peak capacity throughout the whole day. However, given that demand tends to be greatest in the morning, we would still expect the congestion premium to be greater then, as higher prices during peak times would be needed to encourage more price-sensitive passengers to switch to off-peak times, or to be priced out of the market altogether.

Response: We agree with the proposition that the prevalence and scale of scarcity rents will show considerable heterogeneity across different sub-markets.

A2 Does the empirical evidence suggest that there could be a scarcity rent? – insights from fares analysis

Argument Frontier2017-1: Flight prices from Heathrow exhibit a scarcity rent of 23%.

- Frontier (2017): We have found that in 2016, ticket fares at Heathrow were on average 23.3% higher than at other London airports and 24.4% higher than at other European hub airports, despite controlling for other factors that affect fares. This translates to a mark-up in one-way ticket fares of approximately £59. Furthermore, we did not find a mark-up on ticket fares from Gatwick in comparison to the other London airports.

Response: Given the Frontier (2017) set-up and interpretation of the Heathrow dummy variable, we feel that the above inference makes some sense. However, the Heathrow dummy variable is a catch-all for any systematic differences of Heathrow from other (origin) airports, such that the fare premium cannot be solely attributed to scarcity rent. We also take issue with the above inference regarding Gatwick. In particular, the model from which this inference derives includes a dummy variable for Gatwick but not for Heathrow (see Frontier (2017), page 64). Given that Heathrow has a clear fare premium relative to the sample as a whole, in the absence of a Heathrow dummy variable alongside the Gatwick dummy variable, the comparator for Gatwick is an amalgamation of Heathrow and other airports. Such a comparison will not isolate the scarcity rent for Gatwick if there are scarcity rents at Heathrow.

Argument Frontier2019-6: Flight prices from Heathrow exhibit a scarcity rent of around 25%.

- Frontier (2019): Across our models, we find an estimate of the congestion premium of around 17% of fares for short haul and 25% of fares for long haul. This translates into a congestion premium of c. £2.4 billion, and over £200 paid on top of an average long haul return fare. This is even when controlling for characteristics like seat class mix, LCC share or route-level competition, which are arguably themselves symptoms of the capacity constraints at Heathrow.

Response: We refrain from verifying that the Frontier calculations are correct, however we do agree with the last comment of Frontier. Indeed, this echoes our earlier statements that there is a pure congestion effect as well as competition and product distortion effects.

Argument IAG4: Frontier (2017)'s framework is over-simplistic, and the claims regarding scarcity rent are therefore unreliable.

- IAG: The Frontier (2017) report introduces a simplistic framework to analyse the impact of airport congestion. The framework finds that when demand for seats exceeds seat capacity, this results in excess profits that would not have been earned by airlines in the absence of the alleged binding constraint on capacity. In Frontier's view, the shortage of slots at Heathrow would lead directly to a ticket price increase, as airlines would be unable to supply more seats (i.e. more flights) to transport more passengers. However, Frontier's framework is too simple, and as a result its prediction that congestion would give rise to a premium on ticket prices at Heathrow is not reliable.

Response: We disagree with IAG. We believe that the econometric evidence from analysis of fares data can yield useful insights into the extent of scarcity rent at Heathrow, as part of a wider triangulation process. That is not to say we agree that the Frontier (2017) or (2019) work is the best that can be done with the data.

Argument FTI5: Frontier (2017)'s measurement of average fare fails to weight for patronage on each route.

- FTI: Frontier (2017) has used the annual average fare (for 2016) per OD route as the dependent variable and uses OLS to estimate the coefficients for these models. OLS treats each observation with equal weight, but each observation of the dependent variable is an aggregation using a different number of total passengers for each

observation. The fares on the least busy routes impart a greater individual impact on the model results than fares on the busiest routes.

Response: There seems to be two issues at play here. The first is whether the fare measure for low volume flows has more imprecision in its measurement. This would point to a weighted least squares estimation approach to acknowledge that the underlying data has more or less error depending on the overall demand level. Alternatively, and as Frontier (2017) implemented, robust standard errors can be used. Therefore, this issue is not a substantive criticism. The second (more subtle) issue is whether pooling fares across different sized markets is reasonable. The argument is that very small markets (which might have very specialised (high) fares) have a disproportionately large influence on the overall results. The answer here is less clear, but we note two supplementary issues. One is that service frequency was included in Frontier (2017)'s model. That should capture some differences resulting from the size of markets. Another is that Frontier included a lower limit on the number of passengers using a route as part of their selection criteria for including routes in the analysis. Further controls, different functional forms (of say service frequency) in the model and sensitivity analysis of excluding smaller flows could be tested to check the influence of small flows.

Argument FTI6: Frontier (2017) omitted to present or describe diagnostic testing of their econometric model.

- FTI: Frontier have not provided results for post-estimation testing that assesses the 'correctness' of each model. Following estimation of a model's coefficients, it is common practice to run diagnostic tests, which is to assess the reliability of the coefficient estimates.

Response: Whilst this issue is important from the perspective of verifying a preferred model, we do emphasise that there are more fundamental issues regarding the suitability of the Frontier approach and the variables used within. As such, we see this issue as one to be resolved after improvements to the Frontier approach have been implemented. Furthermore, we note that many models used for regulatory purposes do not 'pass' all diagnostic tests. Such tests should help guide model selection and then, once the best (or least-worst) model has been selected, diagnostic tests (and especially failures thereof) should inform the overall assessment.

Argument FTI7: Frontier omitted to present or describe testing of the functional form of their econometric model.

- FTI: Frontier has considered the relevant functional form of the explanatory variables when specifying its models by looking at two-dimensional scatterplots of some explanatory variables and the dependent variable. However, formal testing of a chosen functional form can objectively measure the appropriateness of a chosen functional form based on the results of the estimation. Frontier have not provided any results of diagnostic tests of its chosen functional forms.

Response: This is a valid criticism. However, we again see this as an issue for further modelling once the function to be estimated has been agreed. Thus more fundamental is whether the dummy variable approach in the context of clear missing explanatory variables representing the supply-side of the relationship is a valid starting point. With respect to

functional form, if these changes were implemented, we would highlight that the SEO/Cranfield work estimated a linear relationship between capacity and fare. This may be inappropriate and a non-linear relationship should be considered, as economic theory suggests the impact of capacity increases exponentially as the capacity constraint becomes binding.

Argument FTI8: Frontier omitted to present or describe testing for measurement error within their econometric model.

- FTI: Frontier has chosen to include explanatory variables in its models that are not statistically significant. Measurement error is a possible explanation for a lack of statistical significance of an explanatory variable that is thought to impact a dependent variable. The Frontier report suggests that the data has been inspected in 2-dimensional charts, which can reveal significant, simple measurement errors. However, Frontier has not provided any results of formal measurement error testing. If one or more variables suffer from measurement error, where the size of the error is correlated with the value of the variable, then the coefficient estimate will suffer from attenuation bias.

Response: FTI's criticism of Frontier (2017) is valid, but this has been addressed (to some extent) in Frontier (2019) (which post-dates FTI's report). Whilst we believe that the analysis in general could be improved, this sensitivity testing does give comfort that Frontier's model is not overly sensitive to inclusion or exclusion of marginal explanatory factors. So even if such factors had measurement error within them, this does not seem to impact upon the results. That seems to be the primary point of FTI8, so we do not support it.

Argument FTI9: Frontier omitted to present or describe testing for measurement error within their econometric model.

- FTI: Typically, variables that are not statistically significant should be removed from a model because any common variance they share with the other explanatory variables may be wrongly attributed to them. This would bias the coefficient estimates of the redundant variables and the associated explanatory variables. Frontier states that its regression models suffer from multicollinearity and considers this may reduce the 'congestion premium' estimate using the London Heathrow dummy variable

Response: Whilst FTI8 could be thought to have merit – although empirically this does not seem to impact upon the results – we do not recognise FTI9 as sound statistical reasoning. Including extraneous variables does not bias other coefficient estimates, it makes their estimates less precise. Again we would point to Frontier's (2019) sensitivity analysis to demonstrate that the conjectured potential empirical result of a reduced congestion premium does not substantially arise when variables are excluded (and indeed the effect may be in the opposite direction).

Argument FTI10: Frontier omitted to present or describe testing for outliers within their econometric model.

- FTI: Frontier does not indicate whether it has tested for outliers that may have major effects on the coefficient estimates.

Response: FTI's criticism of Frontier (2017) is fair, but Frontier's latest (2019) analysis included outlier analysis in Annex A. That said, this analysis was fairly cursory – especially as Frontier's definition of an outlier is open to debate – and we agree that more work could be done in this area.

Argument FTI11: Frontier omitted to account for omitted variable bias within their econometric model.

- FTI: Frontier's econometric analysis likely suffers from omitted variable bias. For example, Frontier's econometric models fail to account for differences in within-airport route competition. Routes that are serviced by more carriers at a single airport are likely to see more price competition than routes serviced by fewer or even a single carrier. Therefore, without controlling for within-airport route competition, differences in annual average fares across airports may be overly attributed to excess demand.

Response: FTI's criticism of Frontier (2017) is fair, but Frontier's latest (2019) analysis addressed this criticism.

Argument IAG5: The IATA data may be unrepresentative because of its reliance on travel agency bookings.

- IAG: IATA fare data may not be representative of the majority of air travel bookings. If it is true that the IATA fare data relies on 'traditional travel agency booking', but that 'direct/online bookings' now represent the majority of air travel bookings, then IAG would be correct to criticise Frontier's analysis on this basis.

Response: We believe that the quality of the fare data is very important. The SEO/Cranfield used a similar source to Frontier and there are references in that report that the data has been amended to reflect prices offered in border revenue channels. However, we consider this to be something that the next phase of this work needs to explore further.

A3 To whom does any such rent accrue?

Argument FTI12: The only party (in the value chain described above) that is free to adjust prices is airlines.

- FTI: The Frontier Report explains that, when capacity is constrained, prices must be adjusted upwards in order to equate quantity demanded with quantity supplied. However, as Frontier note, when an airport is regulated (such as Heathrow or Gatwick), the aero charges cannot be adjusted in the same way, as they are fixed by the regulator. Therefore, the only party (in the value chain described above) that is free to adjust prices is airlines. Indeed, the Frontier Report states that '*airports cannot adjust their pricing to ensure that demand equals supply in the constrained case*'. It further states that '*competition in the airline market plays an important role in adjusting prices so that demand equals supply*' and that '*the restricted capacity leads to rising ticket prices so as to match passenger numbers to the seats available*'.

Response: We agree with the caveat that there might be rents to the airport operator if economic regulation is inefficient and/or there is some form of monopoly rent in ancillary activities.

Argument Frontier2019-6: The congestion premium we are estimating – the increase in average fares relative to an unconstrained world – is primarily a form of scarcity rent.

- Frontier (2019): If the demand from passengers to fly from Heathrow exceeds the number of seats that can be made available (because of the capacity constraint) fares will inevitably rise – we refer to this increase in fares at Heathrow as the ‘congestion premium’. The scale of this increase is an empirical matter, depending on the viability of alternative choices available to passengers in different market segments. The congestion premium we are estimating – the increase in average fares relative to an unconstrained world – is primarily a form of scarcity rent. It is not the additional cost of operating at a busy airport, although this will probably be part of it. It is also not the extent to which passengers value Heathrow’s location and amenities, although this is a necessary condition for a premium to exist. It is also not the presence of ‘peak pricing’, which instead arises from the long-run marginal cost patterns facing airlines. It is rather the extent to which prices must rise to ‘choke off’ the excess demand at Heathrow.

Response: The implication of Frontier’s argument is that the rent accrues entirely to the airline. We believe that the airline will be the primary recipient but, with reference to Chapter 5, it is important to also acknowledge the roles played by other agents – namely the airport, the regulator, and the Government.

Argument IAG6: Since HAL has been paying considerably higher dividends to shareholders than the airlines, any scarcity rents must be accruing to HAL.

- HAL’s geared-up returns on RAB equity are a multiple of airline equivalents and HAL can afford to pay its shareholders more in dividends than it invests in LHR. By contrast, British Airways (BA) could not afford to pay dividends between 2001 and 2015; Virgin Atlantic struggled to realise any profits whatsoever; and BMI went bankrupt. Meanwhile, shares in HAL (reportedly) currently trade at a significant premium to RAB equity, while those in IAG currently trade at a discount to assets and at around 50% of the FTSE valuation. So whilst interested shareholders undoubtedly recognise the economic reality of LHR, it is those of HAL (not airlines) who collect a congestion premium. So it is unsurprising that FE was unable to provide any analysis to support its assertion that airlines are generating supernormal profits, by collecting a £59/passenger fare premium at LHR. As HAL appears to be seen to be far more valuable than IAG to shareholders, it must be the case that any scarcity rents that do exist at Heathrow accrue to HAL rather than to airlines.

Response: As discussed above, this elides various issues. The empirical questions are: are LHR aero charges higher than the efficient level (and if so, what is the source and scale of inefficiency)?; how big does the NAPAM model predict scarcity rents at LHR to be today and in 2030 in the DM?; who is getting those rents – bearing in mind that LHR (infrastructure provider subject to economic regulation) and airlines (commercial companies operating in a competitive global market) are quite different businesses, and that airlines’ financial results will reflect the entirety of their global businesses and not just their routes from LHR?

A4 What would happen to the rent if capacity were increased?

Argument Frontier2017-2: If Heathrow were expanded today, ticket fares would decrease by 23% relative to other London airports as a result of removing the capacity constraint.

- Frontier (2017): We have analysed how ticket prices are affected by capacity expansion at both Heathrow and Gatwick and have undertaken detailed econometric analysis to estimate the cost of the congestion premium today. We conclude that expanding Heathrow Airport provides significantly greater benefits to passengers than expanding Gatwick Airport. In particular, we demonstrate that:
 - Expanding either airport is likely to have an impact on ticket prices at both airports in the long term. Overall, however, the reduction in ticket prices caused by expansion of Heathrow Airport is significantly larger than the impact on ticket prices of Gatwick expansion. This is because excess demand at Heathrow Airport is substantially higher than at Gatwick Airport and Heathrow is unique compared to other London airports because it is a hub offering a substantial long haul network.
 - The reduction in ticket prices from expansion at Heathrow is substantially larger compared to Gatwick. If Heathrow were expanded today, ticket fares would decrease by 23% relative to other London airports as a result of removing the capacity constraint. On a return flight basis, this means that over the course of 2016, the congestion premium cost passengers at Heathrow roughly £2 billion. By 2030, this would result in a reduction in one-way ticket prices of £64 and £247 for short and long haul flights respectively compared to a reduction in ticket prices as a result of Gatwick Airport expansion of £24 and £83 for short and long haul flights respectively. Therefore, we conclude that expanding Heathrow provides a much greater reduction in ticket prices for passengers than expanding Gatwick.

Response: In Chapter 7, we highlighted a number of considerations which we believe are pertinent to understanding the prevalence and scale of scarcity rents at Heathrow. Key among these are the heterogeneity of services (and the pattern of congestion therein), the dynamics of supply-side response to congestion, and the likelihood that any rent cannot be solely attributed to scarcity (independently of market power etc.). Thus, whilst we are comfortable with the thrust of Frontier's econometric analysis (to the extent that we believe that their methods are – with some qualifications – defensible), we would caution against making such definitive and generic claims about the impact of capacity expansion at LHR on air fares.

Argument FTI12: Increased capacity will eliminate any scarcity rent only if all demand is satisfied in the Do Something and aero charges remain constant.

- FTI: The Frontier Report then offers the theoretical conclusion that *'if the capacity constraint is removed, new airlines can enter existing routes and this increase in airline competition ensures that prices fall'*. Whilst we would agree that prices will theoretically fall as the capacity constraint is relieved, we note that Frontier conclude that, post expansion, ticket prices would in aggregate fall by the entire amount of the estimated scarcity rent. In our view, this would only be true if:

- the capacity expansion completely clears the capacity constraint; and
- the average aero charge faced by airlines stays at the same level.

Response: We agree with this rationale.

Argument FTI13: If instead expansion only relieves some proportion of the capacity constraint, then ticket prices will not fall by the full amount of the scarcity rent.

- FTI: The Frontier Report's conclusion assumes that Heathrow's planned expansion will eliminate the current capacity constraints at Heathrow. If instead expansion only relieves some proportion of the capacity constraint, ticket prices will not fall by the full amount of the scarcity rent. Airlines will continue to earn scarcity rents, albeit less.

Response: This is obviously correct.

Argument FTI14: If instead expansion only relieves some proportion of the capacity constraint, then ticket prices will not fall by the full amount of the scarcity rent.

- FTI: In relation to ticket prices, the Frontier Report does not discuss the effects of the costs of expansion. Assuming HAL delivers the expansion, the cost of the potential Heathrow expansion would be added into HAL's RAB, increasing the amount it can recover from airlines. However, the passenger numbers would also rise, meaning that the overall impact on the average aero charge per passenger is ambiguous. At this early stage, it is not possible to be definitive on how this metric would change post expansion (it is, of course, hugely dependent on the costs of expansion, the regulatory framework, and the change in passenger numbers). We would simply note that Frontier's conclusion assumes no change to average aero charges per passenger. If the aero charge per passenger rises because of expansion, ticket prices will not fall by the full amount of the previously identified scarcity rent. Conversely, if the aero charge per passenger falls, ticket prices may (but not necessarily will) fall more than the identified scarcity rent.

Response: We agree with this rationale.

Argument FTI14: Capacity expansion will affect ticket prices differentially depending on the city-pair and it is even possible that some may increase whilst others decrease.

- FTI: In its theoretical discussion of scarcity rents, the Frontier Report inherently considers Heathrow airport to be a single market. With numerous airlines operating out of Heathrow, and not all airlines serving the same routes, air travel to and from Heathrow is perhaps more accurately represented as a collection of individual markets. Each airline, and indeed each city-pair, is likely to have a unique cost to serve, and therefore a unique supply curve. Moreover, the overall capacity constraint at Heathrow is likely to affect each city-pair differently. For example, on some routes there might be no capacity constraint and the prevailing price is equivalent to the market clearing price. On these city-pairs there would be no scarcity rents. On other city-pairs, the capacity constraint may be more severe, and on these there would be very high scarcity rents. Therefore, even if the overall average effect of expansion is a reduction

in ticket prices, it is likely that capacity expansion will affect ticket prices differentially depending on the city-pair and it is even possible that some may increase whilst others decrease.

Response: We agree with this rationale – and this is effectively the situation shown in Figure 3.3.

Argument Frontier2019-7: If Heathrow had 50% additional capacity today, the current congestion premium would largely be eliminated.

- Frontier (2019): Expansion is forecast to increase capacity by over 50%. Given that we estimate unserved demand to be around 20% today, we believe that this implies that if Heathrow had that additional capacity today, the current congestion premium would largely be eliminated. What the situation will be when the runway finally opens in 2026 remains to be seen, but with greater capacity it follows that the premium will be much reduced, if not eliminated in the early years – even if the additional capacity is sufficient, it will take time to ramp up use of the new runway. It is likely that there will still be excess demand at peak times, but with spare capacity during other parts of the day, this can be alleviated in part. Additional capacity should also enable there to be increased competition, and scope for low cost carriers to provide a further downward pressure on prices.

Response: We do not believe that demand above capacity is an appropriate measure of congestion. It is likely that demand can grow in a capacity constrained environment by such factors as peak spreading. As such we are wary of this calculation.

ANNEX B: DERIVING THE LERNER INDEX AND SCARCITY RENT

Let the inverse demand function for both airlines be given by:

$$P = P(Q_{A+B})$$

where: $Q_{A+B} = Q_A + Q_B$

Furthermore, assuming that the airlines face constant but possibly different MC, let the total cost function for each airline be given by:

$$TC_i = MC_i \cdot Q_i \quad i = A, B$$

By definition, the profit function of each airline is given by:

$$\pi_i = (P(Q_{A+B}) - MC_i) \cdot Q_i$$

In a Nash equilibrium, each airline will maximise their profit, subject to the quantity supplied by the other airline. On this basis, the First Order Conditions (FOC) for this profit maximisation problem for each airline are given by:

$$\frac{\partial \pi_i}{\partial Q_i} = \frac{\partial P}{\partial Q_{A+B}} \cdot \frac{\partial Q_{A+B}}{\partial Q_i} \cdot Q_i + (P(Q_{A+B}) - MC_i) = 0 \quad i = A, B$$

It must be the case that:

$$\frac{\partial Q_{A+B}}{\partial Q_i} = 1$$

Substituting within the FOC and simplifying:

$$\frac{P - MC_i}{P} = - \frac{\partial P}{\partial Q_{A+B}} \cdot \frac{Q_i}{P}$$

Re-writing this as:

$$\frac{P - MC_i}{P} = - \frac{\partial P}{\partial Q_{A+B}} \cdot \frac{Q_i}{P} \cdot \frac{Q_{A+B}}{Q_{A+B}}$$

Re-arranging:

$$\frac{P - MC_i}{P} = - \frac{\partial P}{\partial Q_{A+B}} \cdot \frac{Q_{A+B}}{P} \cdot \frac{Q_i}{Q_{A+B}}$$

Or:

$$\frac{P - MC_i}{P} = L_i = - \frac{s_i}{\varepsilon_{Q_{A+B}, P}}$$

Where L_i is the Lerner Index of market power of each airline, and s_i is the market share of each airline.

Remembering that scarcity rent is given by the margin between price and MC, this can be re-written:

$$R_i = -P \cdot \frac{s_i}{\varepsilon_{Q_{A+B}, P}}$$

Of course, in order to confirm that the FOC does indeed deliver a maximum, the Second Order Conditions (SOC) must hold:

$$\frac{\partial^2 \pi_i}{\partial Q_i^2} = \frac{\partial^2 P}{\partial Q_{A+B} Q_i} \cdot Q_i + \frac{\partial P}{\partial Q_{A+B}} + \frac{\partial P}{\partial Q_i} \leq 0 \quad i = A, B$$

Since $\partial P / \partial Q_{A+B} = \partial P / \partial Q_i$, this simplifies to:

$$\frac{\partial^2 \pi_i}{\partial Q_i^2} = \frac{\partial^2 P}{\partial Q_{A+B} Q_i} \cdot Q_i + 2 \frac{\partial P}{\partial Q_i} \leq 0$$

**INSTITUTE FOR TRANSPORT STUDIES
DOCUMENT CONTROL INFORMATION**

Title	Independent Peer Review of Recent Research on the Existence of Scarcity Rents at Heathrow
Authors	Richard Batley, Phill Wheat & Peter Mackie
Editor	Richard Batley
Version Number	Final Report
Date	22 nd August 2019
Distribution	Beth Corbould, Andrew Walker, Stuart Holder
Availability	Restricted
Signature	