frontier economics

ESTIMATING THE CONGESTION PREMIUM AT HEATHROW

A report prepared for Heathrow

May 2019

Dan Elliott

dan.elliott@frontier-economics.com

Chris Cuttle

chris.cuttle@frontier-economics.com

CONTENTS

Executive Summary 4
- Introduction 4
- Our approach to estimating the congestion premium 5
- Our results 5

1 Introduction 10
- 1.1 Background 10
- 1.2 Our previous reports on the congestion premium 18
- 1.3 The scope of this report 19
- 1.4 The structure of this report 21

2 Top-down estimates 22
- 2.1 Introduction 22
- 2.2 Approach 23
- 2.3 Results 34

3 Analysis of slot transfer prices 36
- 3.1 Introduction 36
- 3.2 Approach 40
- 3.3 Results 45

4 Econometric analysis 47
- 4.1 Introduction 47
- 4.2 Data 50
- 4.3 Approaches to capturing congestion 52
- 4.4 Model specification 54
- 4.5 Results 56
- 4.6 Conclusions 64

5 Conclusion 67

Annex A Technical annex on econometrics analysis 70

Annex B Response to FTI Report 88

EXECUTIVE SUMMARY

Introduction

Heathrow has been full in terms of aircraft movements for a long time. Traffic first exceeded 90% runway utilisation in 1995 and has been effectively at 100% since at least the mid-2000s. Over the same period total traffic to and from the London airports as a whole has more than doubled. But with Heathrow full, that growth has largely taken place at other airports.

Figure 1. Heathrow is capacity constrained

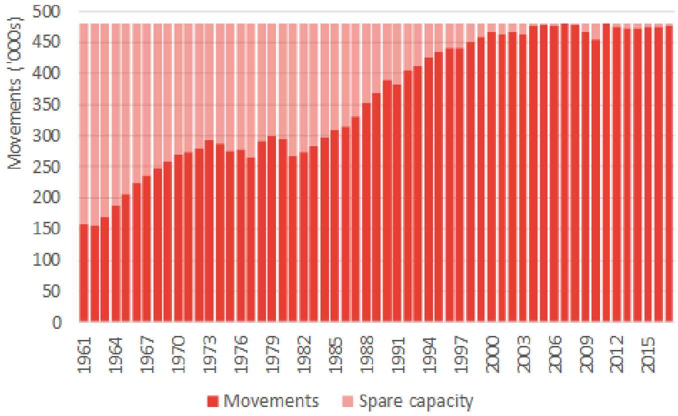

Source: Frontier analysis of CAA data

Figure 2. Heathrow is growing at a constrained rate

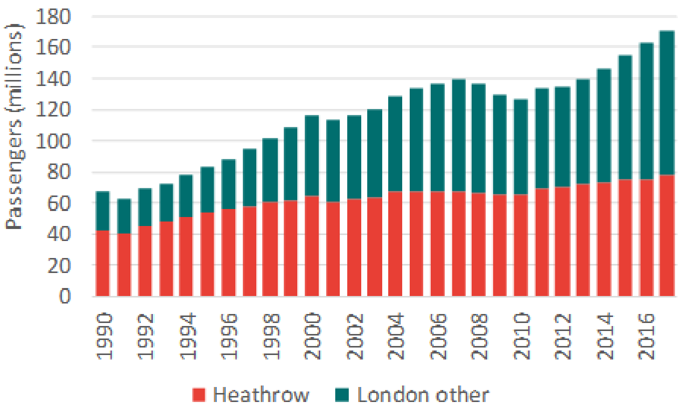

Source: Frontier analysis of CAA data

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.

We have estimated the congestion premium at Heathrow before on two occasions using econometrics. In both instances, we found it be around 25%. In response, in 2018 the Civil Aviation Authority (CAA) commissioned FTI Consulting to review our more recent work. They were generally positive, and we welcome the helpful comments and are pleased to see that the debate is continuing.

This is our third review of the congestion premium. This report builds on our previous work by both adding to our econometric analysis and producing new pieces of supporting analysis, each of which produce separate estimates of the congestion premium by way of corroboration of the statistical results. We consider the combined results of these approaches to be our most comprehensive estimation of the congestion premium yet.

Our approach to estimating the congestion premium

In the past, we have modelled route-level fares at Heathrow and comparable airports to estimate the size of the congestion premium. Our updated econometric analysis, builds on the approach of our previous reports, by adding more comparator airports, more years and explores more models.

We have also carried out two pieces of supporting analysis that produce separate estimates of the congestion premium:

- An analysis of slot transfer prices; and
- A top-down analysis of passenger volumes.

As a consequence we now have three separate approaches that all point towards the presence of a significant premium.

Our results

Econometrics

We use econometric analysis to compare ticket prices at Heathrow with those at other large airports in Europe, and quantitatively explain the drivers of any differences, and in doing so isolating the impact of congestion.

We model air fares for over 2,000 short haul and long haul routes over 8 years at 10 European hub airports. Our models control for a range of other variables, including distance, frequency, low cost carrier (LCC) share, the economy/business seat mix, route-level competition and airport quality which could otherwise explain variations in fares.

Figure 3 Structure and presentation of econometric models

Models are more determined/fitted

More controls are added

Variables	Narrow	1	2	3	4	Wide
Distance (000s KM)	✗	✗	✗	✗	✗	✗
Heathrow dummies	✗	✗	✗	✗	✗	✗
Year dummies	✗	✗	✗	✗	✗	✗
Frequency (own)		✗	✗	✗	✗	✗
Frequency (other)		✗	✗	✗	✗	✗
Transfer share		✗	✗	✗	✗	✗
LCC share			✗	✗	✗	✗
Business class share			✗	✗	✗	✗
European jet fuel price				✗	✗	✗
Avg. seats					✗	✗
Airport competition					✗	✗
Route competition					✗	✗
Skytrax rank						✗

Source: Frontier Economics

We also explore models where we control for congestion explicitly, to show that our congestion premium estimates are not a quirk of our core models, or the result of some other Heathrow-specific characteristic. Specifically, we use the top-down estimates of unserved demand as an estimate of congestion, along with an "out-of-sample" approach that compares actual Heathrow fares with predictions based on analysis of the other airports only.

Figure 4 Econometric estimates of the congestion premium

	Short haul	Long haul	Combined
Core models	14-22%	22-29%	20-27%
Unserved demand models	12%	22%	20%
Out-of-sample estimates	23%	27%	26%
£ per year	£430m	£1.9bn	£2.4bn
£ per return flight	£34	£217	£110

Source: Frontier analysis of IATA data

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.

In the interests of transparency and rigour, this report provides a detailed technical annex of our econometric results.

Top-down analysis of passenger volumes

Up to 1990, the growth in movements at Heathrow was broadly in line with the growth at other London airports. However, since 2000, growth at Heathrow has largely plateaued. Since then, the annual movements at Heathrow has increased by around 9,000 (less than 2% growth over nearly 20 years), whereas at other London airports, it has increased by over 100,000 (equal to around 16% growth over the period).

We estimate that, if Heathrow were not constrained, it is plausible that demand would have increased by around 2.3% per annum on average. Applying this higher growth rate to passenger volumes since 2000 implies that demand would be 17.5% higher at Heathrow today. Pricing out 17.5% of unconstrained demand, using a price elasticity of demand of -0.7 (the most conservative number used by the DfT in its forecasting analysis), implies a congestion premium of 25%. This means an additional £47 per return fare on short haul flights, and an additional £216 per passenger on long haul flights. In total, this amounts to £2.6 billion across all passengers[1] in 2018.

This is a high level calculation, insufficient by itself to make the claim that a congestion premium exists. But it serves to corroborate the hard evidence from our econometric study.

[1] For this calculation, we have conservatively assumed it applies only to point-to-point customers

Figure 5 The congestion premium over time estimated on a top down basis

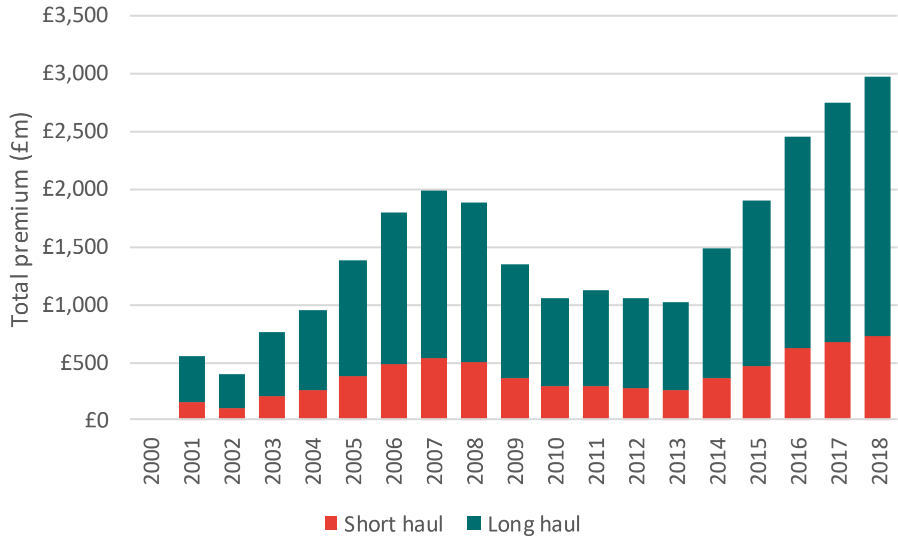

Source: Frontier analysis

Slot transfer prices

Despite likely continued growth in passenger demand at Heathrow, since the early 2000s there have been few to no new slots to allow new airlines or incumbents to add more flights due to capacity constraints. At the same time we have seen a rising number of financial transactions involving slots at Heathrow, most notably the record-breaking purchase of a daily pair of slots by Oman Air in 2016 for $75m.

Figure 6 Very few new slots become available in the pool at Heathrow

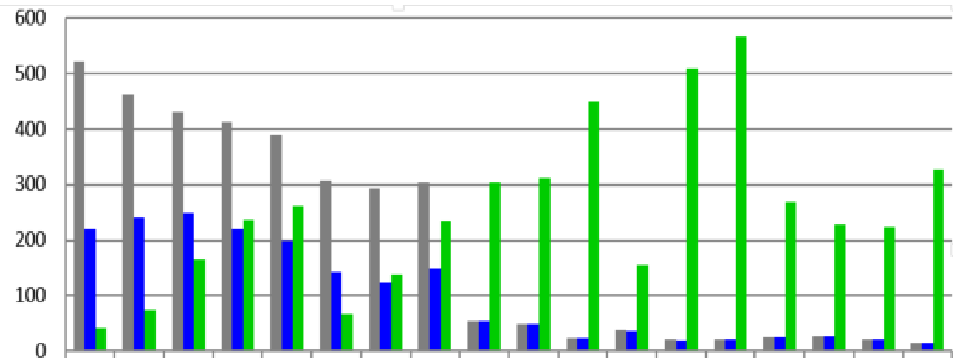

	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017
Pool Availability	520	462	432	412	390	308	293	303	54	48	23	38	22	22	26	28	22	16
Pool Allocation	219	241	249	221	199	143	123	148	54	48	23	36	20	22	26	28	22	16
Slot Transfers	42	74	166	236	262	68	139	235	304	313	449	156	508	568	268	228	224	326

Source: Heathrow Airport Holdings

At an airport that has spare slots at any time of day, slots will have little to no value. But at a congested airport, with no slots free in the pool, an airline can only acquire a slot by obtaining it from a rival airline. This scarcity therefore makes the slots valuable.

Details of slot transaction fees are usually not readily publicly available. Instead we have to rely on slot prices as they are reported in the media, a sample of which has been collated by HAL. Using this data we have estimated the total premium across all flights at Heathrow per annum.

The link between payments for slots and the congestion premium is that, in order for the transaction to make commercial sense, the premium an airline expects to earn when using the slot needs to be greater than or equal to the upfront fee that the airline pays for the slot.

We estimate that if we take the average price paid per slot in the sample above, which was £812,000, and assume an asset lifetime of 10 years and a WACC of 7.5%, the annualised value is equal to £2.0 billion. The implied premium would therefore be c. £2.0 billion.

Again this is a high level calculation, with some uncertainty about the appropriate assumptions to use, but it too serves to corroborate our econometric results.

Conclusions

There are caveats and limitations with each of these approaches when taken in isolation. For example, it is always possible to challenge the robustness of econometric analysis and technical issues associated with benchmarking prices at different airports. To address this we have used transparent, reasonable assumptions, and in the case of the econometric analysis, used a wide range of alternative approaches to demonstrate that our results are not an artefact of the assumptions we have made.

The estimated size of the premium is consistent across all three approaches, and we believe that this demonstrates that the premium is not a quirk of econometric analysis, but a clear pattern which can be observed across a variety of methodologies.

Figure 7 All three approaches point to a consistent premium per annum

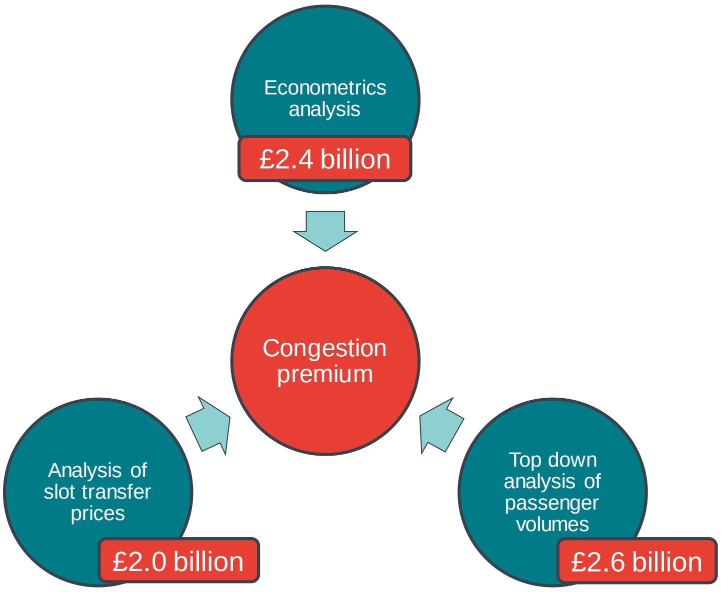

Source: Frontier analysis

Overall, we estimate a congestion premium of around 25% on the average fare, meaning around £2 billion per annum. For individual passengers, our analysis suggests they pay on average an extra £34 on a short haul return flight and (also on average) over £200 on a long haul return flight. This has important implications for consumer protection, and for the choices made during expansion and how they could alleviate the congestion premium.

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.

1 INTRODUCTION

1.1 Background

Heathrow has two runways and capacity for 480,000 flights a year. As shown below, traffic first exceeded 90% runway utilisation in 1995. And from the early-/mid-2000s traffic has been effectively at full capacity. This makes Heathrow one of the, if not the most congested airports in the world.

Figure 8 Heathrow is capacity constrained

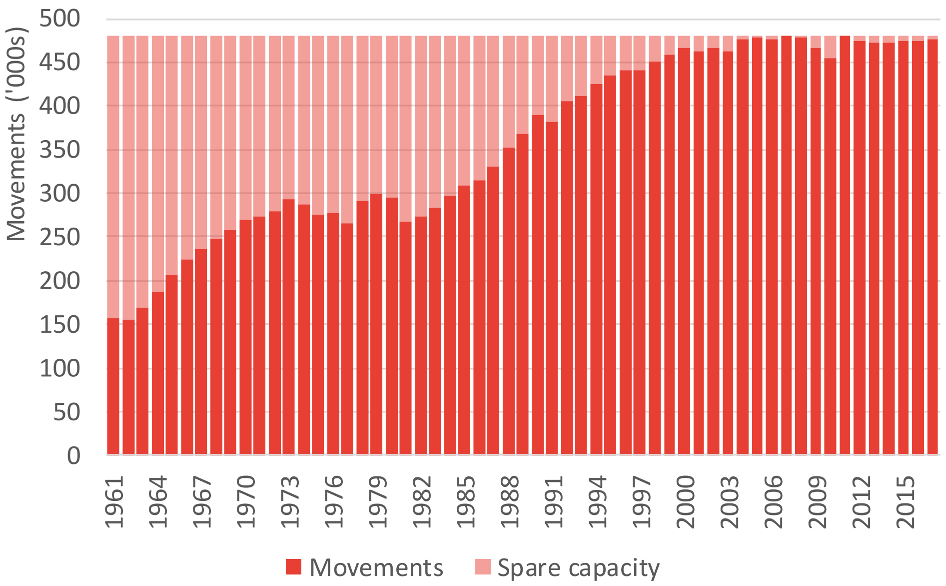

Source: Frontier analysis of CAA data.

Since the mid-1990s, total traffic to and from the London airports as a whole has more than doubled. But with Heathrow full, that growth has largely taken place at other airports.

Figure 9 Heathrow is growing at a constrained rate

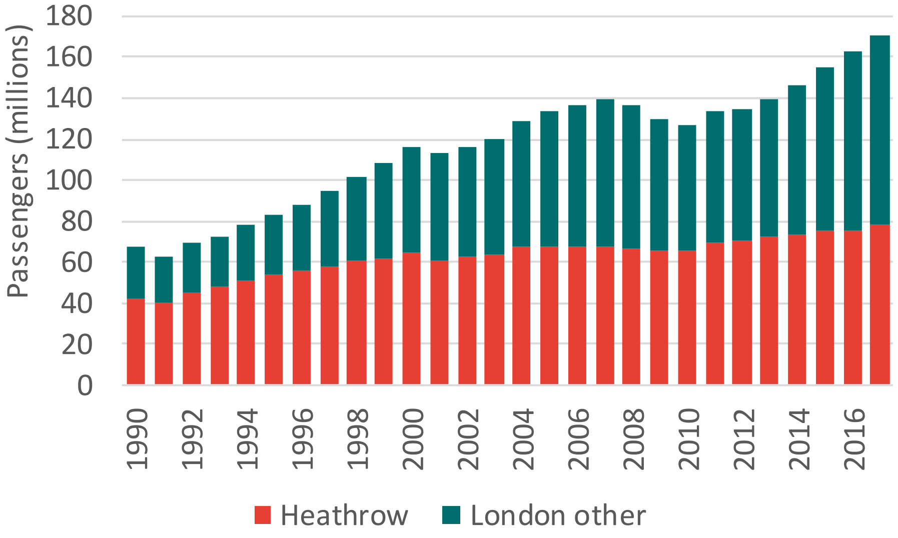

Source: Frontier analysis of CAA data

Furthermore, there is no end in sight to the growth in demand. The DfT forecasts that if there were no capacity constraints at UK airports, then total demand across the UK would more than double from around 260 million passengers per annum (mppa) in 2016 to as many as 535 mppa by 2050. This growth stems from a well recorded and persistent relationship between the demand for air travel and general economic growth. The DfT aviation forecasts expect demand to grow at a rate faster than income, as shown in its income elasticities of demand (see Figure 10).

Figure 10 The DfT's air traffic demand elasticities

		Elasticity with respect to	
Sector	**Share of base year demand**	**Income**	**Air fares**
UK business	7%	1.2	-0.2
UK leisure	47%	1.4	-0.7
Foreign business	6%	1.0	-0.2
Foreign leisure	19%	1.0	-0.7
International to international transfers	9%	0.5	-0.5
Domestic	12%	1.1	-0.5
Total	100%	1.2	-0.6

The Airports Commission change to the international to international transfer fares elasticity has been retained. The elasticity changed from -0.7 to -0.5, reflecting that it now relates to a broader market, following the inclusion of overseas hubs.

Source: DfT – UK Aviation Forecasts 2017

But existing airport capacity is insufficient to meet this growth. By 2050, it is estimated that demand could exceed available capacity by around 100 mppa (or around 20% of total forecast demand).

Figure 11 The DfT's passenger forecasts

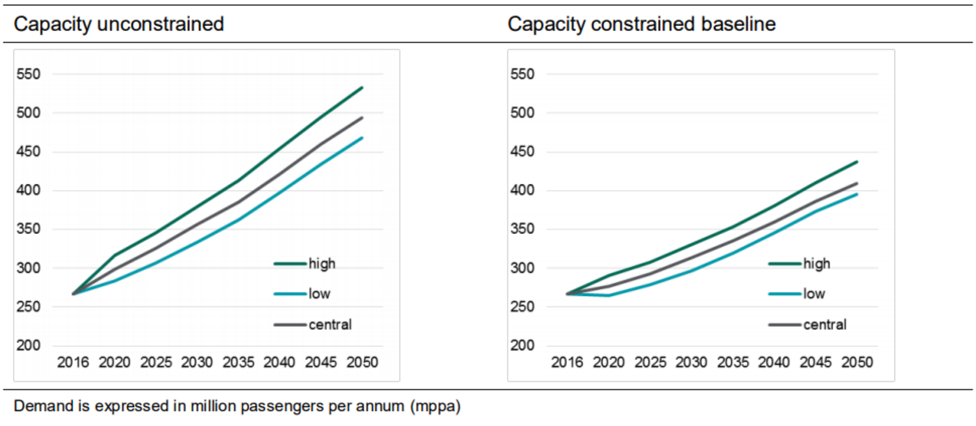

Source: DfT – UK Aviation Forecasts 2017

The third runway at Heathrow is currently expected to open in 2026 and is set to increase capacity by just over 50% to 740,000 movements a year. But until then, the only room for extra flights at Heathrow would come from very minor adjustments to operating procedures.

In any market, when demand exceeds capacity in the short run, prices must rise to the point where the excess demand is priced out of the market. Markets for air travel are no different.

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 – this is how markets work. In this case we refer to the increase in fares at Heathrow resulting from the shortfall of capacity as the "congestion premium".

It is important to stress at the outset that this premium does not stem from the airport simply being full. It is possible, indeed likely, that operating at a very full airport is more expensive for airlines, because of greater delays on the ground, more holding delays in the air and a general lack of operational flexibility. And in a competitive aviation market these costs, experienced by all operators, could be expected ultimately to feed into fares. The "congestion premium" we are referring to may encompass these costs but it is not limited to them. The wider premium stems from the inevitable economic logic of what happens when more people want to consume a product than the supply permits. In these circumstances, *if* supply cannot adjust, prices can be expected to rise in excess of cost to choke off demand.

Figure 12 Congestion premium

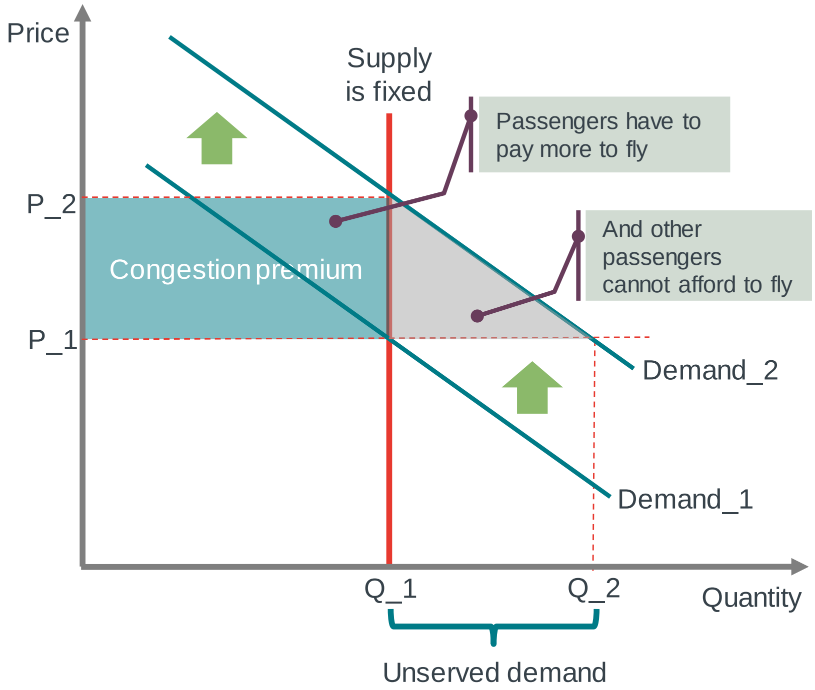

Source: Frontier illustration

Heathrow is not simply full, but rather has been full at all times of day and year for more than a decade while the demand for air travel in general has continued to increase[2]. This fact suggests there is a strong possibility that significant levels of demand to use Heathrow may be being frustrated.

However, we note that while economics suggests that a congested airport will experience a congestion premium, measuring the size of that premium is an empirical matter, and the main subject of this report.

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:

- Location: First, all airports have some degree of differentiation due to their location. We expect passengers to have a strong degree of preference for their local airport, all things being equal.
- Surface access: The surface access links between the London airports and central London are not all equally good. Heathrow is particularly well-connected

[2] Since 2000, total passenger volumes at London airports combined (defined by the CAA in its annual reporting as Heathrow, Gatwick, Stansted, Luton, London City and Southend) has grown by 2.3% per annum on average. And across the UK as a whole, it is has grown by 2.7% per annum. At Heathrow, however, the growth has only been 1.1% per annum.

with public transport and roads, and has the largest drive time boundaries compared to the other London airports.

- Service quality: Airport service quality may vary, which affects passenger and airline preferences. For instance, Heathrow is recognised as having the best service quality out of the London airports. For example, since 2011 it has appeared in Skytrax's Top 20 best airports in the world for service quality, outperforming the other London airports in all years.[3]
- Hub operations: Heathrow is the UK's main gateway for long haul destinations. In 2018, over 80% of all long haul flights departing from London airports departed from Heathrow.[4] This is because the economics of long haul operations typically require hub operations to provide feeder traffic for the long haul services. These hub operations create network effects that make it difficult for airlines to relocate long haul services to other airports. As a result Heathrow is the often the only airport in London where a direct long haul connection can be obtained.

It has been noted that this differentiation may explain variation in fares at different airports. And that, therefore, any observed "premium" at Heathrow may not be due to congestion. This is a misunderstanding of the economics of congestion. Differentiation is a *necessary* condition for the presence of a premium. It is not a *sufficient* condition.

While noting that competition *between* airports is inevitably differentiated (if only by location), competition between airlines *at* an airport is not.

If an airport is uncongested, with free entry and exit of airlines on any route they choose to serve, it would not matter if passengers valued using one airport more than another. Because fares could not persistently rise above cost without provoking further entry, driving fares back down again. This is what competitive markets do. They prevent prices rising to capture the value customers place on a product or service. 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.

As we have already stated, the extent to which any of the effects hold true in practice is an empirical issue which we explore in this report.

But a number of further observations stem from this discussion.

First, we note that 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

3 https://www.worldairportawards.com/the-worlds-top-100-airports-2017/

4 Frontier analysis based on OAG Schedules data

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.

Secondly, 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 limit 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%. We observe that on average Heathrow has managed to grow at about 1% p.a. in terms of passengers while movements have not increased. This reflects the flexibility airlines have been able to apply to this constrained situation and will, to some extent, have alleviated the impact of the congestion premium. Again, we must repeat that the extent to which this has been effective is an empirical matter.

Thirdly, 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.

Hence the main point still holds: average ticket prices at Heathrow would be lower if it were not constrained.

Finally we note that the congestion premium at Heathrow should not be confused with the more general peak pricing observed at many unconstrained airports during peak times of the day or peak times of the year. By way of illustration, the chart below, based on 2018 data, shows the average price of a one-way ticket from the UK to Spain and Italy – the two most popular destination countries abroad for UK passengers. This includes all airports in the UK (and not just Heathrow) to all airport in Spain and Italy.

Figure 13 The airline business model

Average ticket price from the UK to Spain and Italy in 2018

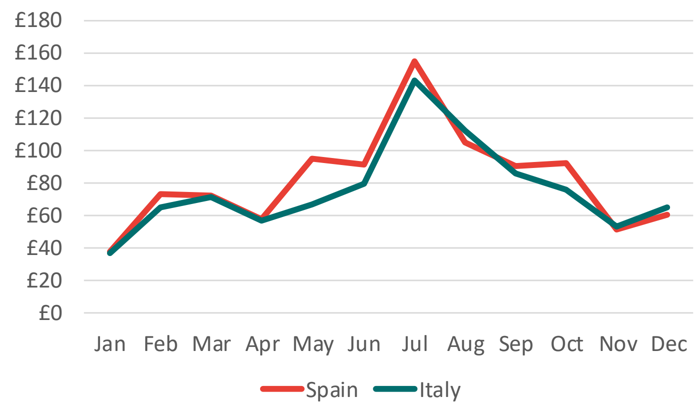

In the long run, airlines must recover all costs

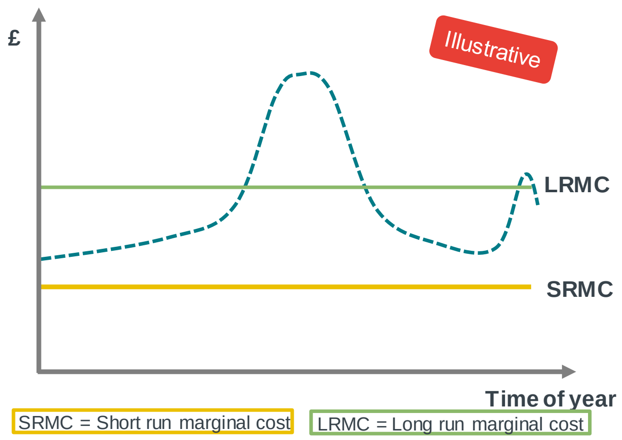

Source: Frontier analysis of OAG data

Note: The chart on the right is based on dummy data for illustrative purposes.

It shows that prices were lowest in January and highest in July. A broadly similar pattern is observed across many other destinations served from the UK. This pattern is a natural feature of the airline business. Airline capacity remains broadly fixed throughout the year, but demand varies. With a fixed amount of capacity, airlines need to lower prices during quieter periods to stimulate demand, and then increase prices during peak periods to suppress demand. While price may vary, operating costs remain broadly similar, aside from any seasonal variation in airport charges or increased costs during peak periods, etc.. This means that profitability varies over the course of the year. But this can still be consistent with a competitive market:

- During quieter periods airlines excess capacity results in lower fares being set in the market. Airlines will continue to operate these flights provided fares cover the short run marginal cost of operating these flights. This includes airport charges, taxes, fuel, and labour rates for cabin crew and pilots, etc.. While this makes a cash contribution to the business, it may make little to no contribution to the recovery of fixed costs such as aircraft or head office costs.
- During peak times seats are more scarce and fares increase. It is during these periods that airlines recover their fixed costs.

It should be noted that in neither case do airlines "set" fares. These are competitive markets in which fares adjust to fill the capacity made available. Airlines are adjusting the number seats they offer in response to the revenues they can generate, not the other way round. In particular, if they choose too much capacity, peak fares will be lower and the airlines will fail to recover their fixed costs, leading to subsequent retrenchment.

The higher prices during the summer months compared to lower prices in winter reflect a premium of sorts. But this pattern occurs at airports which are not congested, and simply reflects the balancing act airlines must perform to recover their costs in total over peak and off-peak times. But the premium that may be observed at highly congested airports like Heathrow goes beyond this pattern. Rather than demand exceeding capacity for just a few weeks of the year, the situation is more chronic.

In conclusion, the congestion premium is important for many reasons:

- Customer protection: Passengers need to pay more to fly and understanding how much more is important. The CAA has a responsibility to protect the interests of customers, and it could have some regulatory options to help alleviate the situation – at least in part.
- Expansion: In our view, alleviating capacity constraints and thereby reducing the congestion premium is perhaps the single biggest benefit of expanding Heathrow. But while Heathrow has been given permission to build a third runway, precisely what it builds in terms of terminal design and cost is still to be decided.

  There are many design options each varying in terms of timing, size, the mix of traffic that can be served, and clearly on cost too. If the extra capacity was used exclusively for short haul routes, this would do little to ease congestion on long haul routes. It is important to keep the congestion premium in mind when weighing up different options, and to evaluate them on the extent to which they can alleviate the congestion premium.

  Media reports have estimated that the cost of expansion could be between £14 billion and £17 billion, although these estimates are still speculative and subject to change.[5] While this is a large difference, we believe they should be viewed alongside estimates of the congestion premium. A more costly option (which we note would be a one off cost) could do more to alleviate the congestion premium (for many years).

The third runway will increase the number of runway slots at Heathrow by just over 50%, allowing incumbent airlines and new entrants to add many extra flights. With extra capacity added on many routes, airlines will need to lower prices to stimulate demand to fill those flights. This is especially the case for new entrants who may seek to enter the routes with the highest congestion premia and undercut incumbent airlines. At a congested airport, prices on constrained routes will reflect willingness to pay, however at an unconstrained airport, which additionally allows for new entrants and helps to facilitate competition, prices will tend towards the long run marginal cost of providing the service.

5 https://www.bbc.com/news/uk-england-london-42399840

1.2 Our previous reports on the congestion premium

We have estimated the congestion premium at Heathrow in two previous reports:

- Impact of airport expansion options on competition and choice (April 2014)[6]. This was based on 2012 data.
- Competition & Choice 2017 (December 2017)[7]. This was based on 2016 data.

In both instances our approach was to use econometric analysis to estimate the congestion premium. At a high level, this involves comparing average ticket prices on individual routes at Heathrow with those at other airports, and controlling for other factors which may also impact on price, such as flight distance and, for instance, the extent of low cost carrier competition on the route. We argued that any difference observed after making these adjustments is probably attributable to the fact that Heathrow is a significantly congested airport, whereas the other airports in the sample are not. Or in other words, if Heathrow had spare capacity, airlines – both incumbents and new entrants – could increase capacity on the most constrained routes which would see ticket prices fall and move closer toward the cost of operating the route.

In both reports we estimated that the congestion premium was around 20%-30% per passenger (depending on the model chosen). This equated to around £20-£30 per passenger for short haul flights, and £90-£100 per passenger for long haul flights. Across all passengers at the airport, this equated to around £2 billion in total per annum.[8]

Our second report was critiqued by IAG. In 2018 the CAA commissioned FTI to review our report and IAG's critique and to publish its view. A timeline is provided below:

Figure 14 Timeline

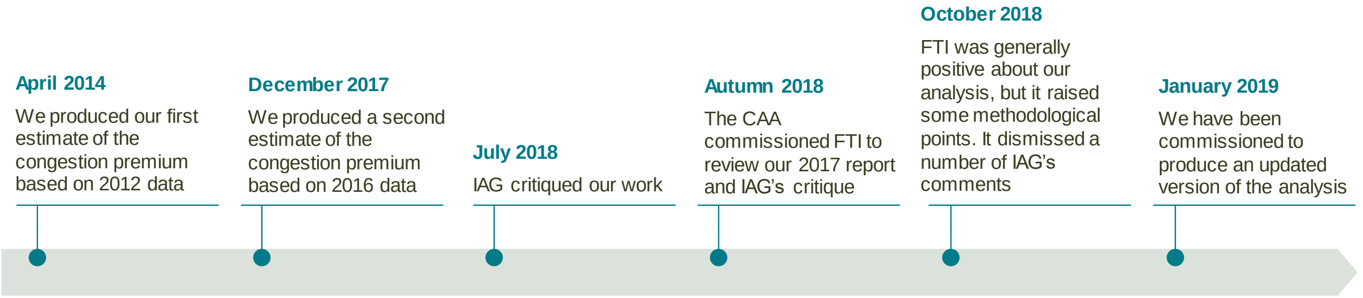

6 https://your.heathrow.com/takingbritainfurther/wp-content/uploads/2014/04/frontier-report-Impact-of-airport-expansion-options-on-competition-and-choice.pdf

7 https://www.caa.co.uk/uploadedFiles/CAA/Content/Accordion/Standard_Content/Commercial/Airports/HAL%20-%20Frontier%20Competition%20and%20Choice.pdf

8 We estimated the premium paid by passengers making point-to-point journeys only, as opposed to transfer passengers at Heathrow (e.g. those flying from Dublin to Heathrow and then on to New York) or 'beyond' / 'one-stop' passenger at Heathrow (e.g. those flying from Heathrow to Madrid and then on to Buenos Aires). As described in more detail in the rest of this report, this is because it is unclear from a theoretical perspective whether airlines would be able to charge a premium for these passengers given that by definition they face many options to fly – e.g. a transfer passenger could fly via Heathrow or potentially a number of other hubs. Whereas for point-to-point passengers, Heathrow may be the only option.

IAG generally dismissed our arguments:

- *"Airlines compete intensely at a route level and so there can be no fare premium"*[9]

We disagree with this point. It is not possible for new airlines to enter at Heathrow and compete on routes with a high premium because there are no slots available. Clearly there is route-based competition between airlines already at Heathrow, but as mentioned above, this can only go so far. An incumbent airline can switch to a route with a high congestion premium, where the extra capacity provided will help to alleviate the premium on that route. But this is only possible if capacity is removed from another route, where clearly the premium will be worsened. The fact that the overall constraint is binding is unavoidable.

FTI was generally positive about our work, whilst also dismissive of some of IAG's challenges:

- *"Frontier's analysis makes a valuable contribution to the debate around scarcity rents at Heathrow airport and makes valid arguments that are worth further consideration"* and
- *"the IAG Response reflects a misunderstanding of the details of the econometric analyses in the Frontier Report and fails to make a persuasive case for any bias in the results"*[10]

Ultimately, econometric analysis is not an exact science, and we believe it is not possible to produce one single model which can explain ticket prices perfectly. We welcome the dialogue which helps develop our approach to measuring the congestion premium, and we have reflected on the challenges and questions raised by both FTI and IAG.

We are pleased that our work has furthered the debate around the congestion premium, and that the CAA – who commissioned the FTI report – is exploring the issues in more detail. We think this is an important area of public policy with significant potential benefits to consumers and the UK economy.

1.3 The scope of this report

In light of the ongoing debate, we have been commissioned by Heathrow to update our estimate of the congestion premium. This updated analysis builds upon our previous work and we consider it to be our most comprehensive estimation of the congestion premium yet:

- **Econometric analysis:** building on the approach of our previous reports, and taking into account the comments made in the FTI report, we have extended our analysis in a number of ways:
  - More years: Rather than estimating the congestion premium at Heathrow in one single year, we have carried out the analysis on data over the period

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

[10] http://publicapps.caa.co.uk/docs/33/CAP1722b%20FTI%20Consulting%20-%20Report%20to%20CAA%20on%20Scarcity%20Rents%20-%20Final%20-%2030%20October%202018.pdf

2011 to 2018. This is useful to consider how stable any models of ticket fares are and how the congestion premium has evolved over time as demand continues to grow.

- More comparators: Our econometric analysis is based upon comparing ticket prices at Heathrow with those at other airports. We have significantly widened the number of comparator airports in our sample to increase the number of observations, in order to improve the overall robustness of our analysis.
- More models: Econometric analysis remains our main approach to estimating the congestion premium. The purpose is to explain ticket prices as a function of key explanatory variables. There is no single perfect model to explain ticket prices, and every model has its pros and cons. In this report we have carried out more tests and sensitivities. We have also carried out more model diagnostics to assess how well our models perform and the potential for statistical bias in our results.

- **Supporting analysis:** To support the econometric analysis, we have also carried out two additional pieces of analysis that produce independent estimates of the congestion premium. This means that we are not just relying on econometrics, but rather have three separate approaches that all point towards there being a significant premium. Alongside the econometrics, these additional approaches are:
  - Analysis of slot transfer prices at Heathrow: If there is no congestion premium at Heathrow, then why do airlines spend significant amounts to acquire slots at Heathrow? What is the economic rationale for these payments if there if fares relate purely to airline costs? We believe that this is clear evidence of the congestion premium: Airlines are prepared to pay significant sums to acquire slots. Without the expected premium this behaviour would not make commercial sense. We have analysed slot transfer prices to produce an estimate of the congestion premium – which does not require any econometric analysis or any comparison of fares at Heathrow versus those at other airports.
  - Top-down analysis of passenger volumes at Heathrow: Heathrow has been constrained for many years. In this time, the number of movements has hardly increased, and passenger volumes have only been able to grow at a constrained rate. We have estimated what passenger volumes might have been at Heathrow in a world where it were not constrained. We can then infer what the congestion premium might need to be in order for this excess demand to have been priced out of the market.

By having three separate approaches all pointing towards there being a significant congestion premium, we consider this report to be our most comprehensive estimate of the congestion premium yet, and demonstrates that the premium is not just a quirk of econometric analysis.

Figure 15 We have produced separate estimates of the congestion premium

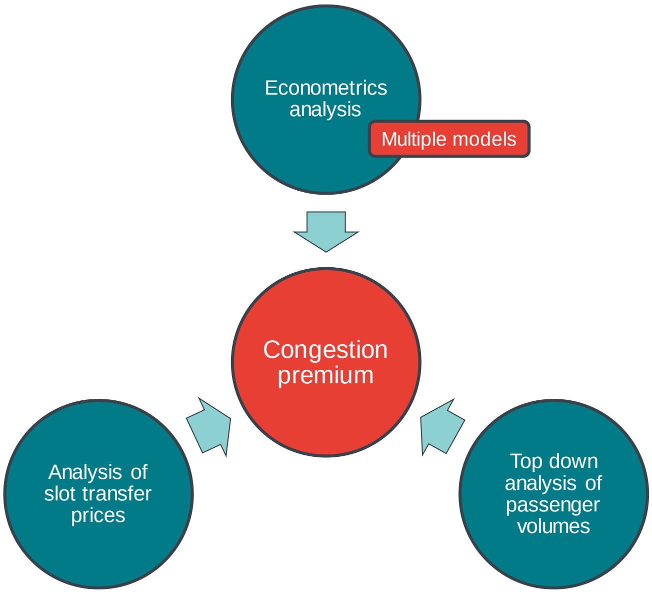

1.4 The structure of this report

The rest of this report is structured as follows:

- In **Section 2** we describe our **top down analysis** of passenger volumes at Heathrow.
- In **Section 3** we describe our analysis of **slot transfer prices** at Heathrow.
- In **Section 4** we describe the **econometric analysis** that we have carried out to estimate the congestion premium.
- In **Section 5** we provide our overall conclusions.

We also include the following annexes:

- In **Annex A** we provide more technical details of econometric analysis.
- In **Annex B** we provide a detailed breakdown of how this update to our analysis responds to each of the points made in the FTI report.

2 TOP-DOWN ESTIMATES

2.1 Introduction

In this section we estimate the congestion premium by analysing how passenger volumes at Heathrow have evolved over time.

Because movements at Heathrow have been constrained for many years, passenger volumes have only been able to grow at a constrained rate, reflecting primarily a gradual shift in fleet and route mix. However, if Heathrow had been unconstrained during this period, we expect that it would have been able to grow at a faster rate and that passenger volumes would be higher today. This is illustrated below:

Figure 16 Top-down analysis – illustration of approach

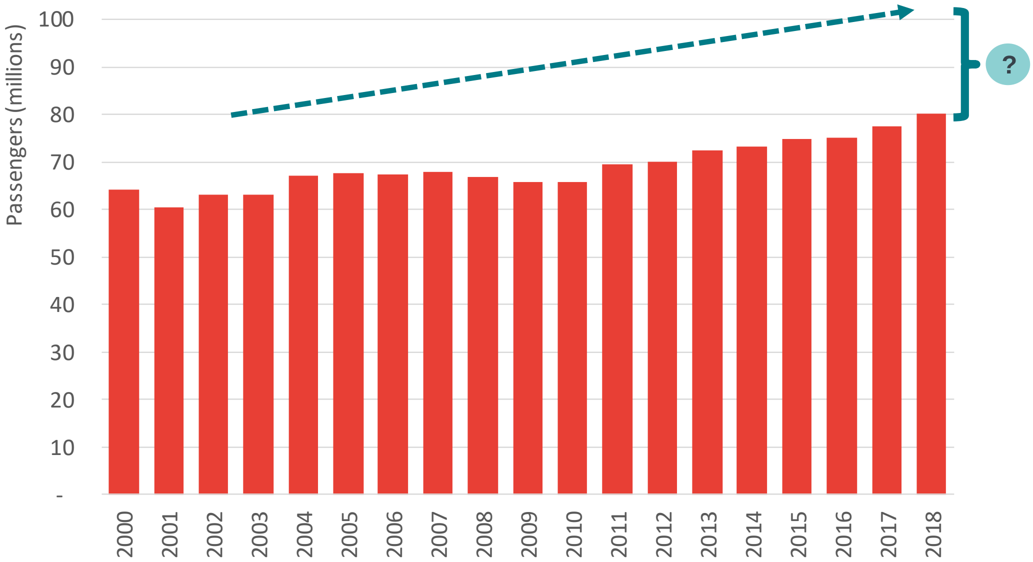

Source: Frontier analysis based on CAA data

If we can estimate what passenger volumes might have been under an unconstrained scenario, we can then consider what price increase (congestion premium) would have been required to price this excess demand out of the market.

2.2 Approach

A high level overview of our approach is set out below.

Figure 17 Top down analysis – overview

1 When did LHR first become constrained?	▪ First, we need to define when Heathrow first became constrained.
2 Unconstrained growth rate	▪ From this point, we can then consider what the unconstrained growth rate might have been.
3 Unserved demand	▪ By comparing the unconstrained passenger profile with the actual (constrained) passenger profile, we can then estimate the size of the unserved demand.
4 Congestion premium %	▪ We can then multiply the unserved demand % by an assumed price elasticity of demand to estimate the % increase in price required to price the excess demand out of the market.
5 Total congestion premium (£m)	▪ By analysing ticket fares and total passenger volumes, we can then estimate the size of the congestion premium in total.

We describe these points in turn in more detail.

2.2.1 When did Heathrow first become constrained?

It is challenging to estimate exactly when Heathrow first became constrained. Being capacity constrained is not a binary black-or-white condition, and in principle it can arise because of different constraining factors:

- Runway capacity: An airport may have plenty of spare terminal capacity but not enough runway capacity. This is driven by the number of runways and there being an efficient traffic management system on the aprons and taxiways.
- Terminal capacity: An airport could have plenty of runway capacity but not enough space within the terminal buildings to actually accommodate or process the passengers. For example, the area around check-in desks or security control could be very small such that there are very long queues and delays, and passengers may choose to fly less regularly to avoid the inconvenience.
- Gates: An airport may have plenty of runway capacity, and plenty of space in the terminal buildings but it may not have enough gates to accommodate more aircraft, or perhaps it does not have gates suitable to accommodate more wide-bodied aircraft.
- Surface access: The car park or the public transport arrival and departure hall could also be very small, or the trains and buses could be so infrequent or small that there is a hard operational limit to the number of passengers who can actually arrive at or depart from the airport.
- Planning restrictions: Local authorities may impose restrictions which also act as a binding constraint. For example, airports near urban areas are usually only

allowed to operate at a restricted level during the night. We note that Heathrow's limit of 480,000 flights per annum is not a physical or technical limit, but rather a planning restriction. We understand that Heathrow is able to handle up to 88 flights per hour[11]. This would imply a technical limit of around 770,000 flights per annum, of which over 560,000 would be possible during daytime hours (11.30pm to 6am).

For an airport to be unconstrained, there must be enough spare capacity across all parts of the airport, and not having enough capacity in just one single part of the airport could be enough to create a bottleneck which makes the airport constrained.

Furthermore, capacity bottlenecks will not bite equally at all times of day or year. As an airport becomes congested it first experiences that congestion at peak times. As demand increases these periods of congestion lengthen – in Heathrow's extreme case to cover virtually all operating hours.

We note that capacity may increase over time as airports improve processes and traffic control and become capable of achieving a greater throughput from their existing infrastructure. Airports may be able to change airline behaviours or passenger behaviours such that they use infrastructure less intensively. For example, if airlines reduce their turnaround time, the airport can free up gates for other airlines, thereby increasing overall throughput.

At Heathrow, the key bottleneck is runway capacity. It has been operating at over 90% runway utilisation since 1995. The chart below, based on data from the CAA[12], reports how the number of movements evolved over the period 1961-2017 at:

- Heathrow;
- London as a whole – defined by the CAA in its annual reporting as Heathrow, Gatwick, Stansted, Luton, London City and Southend;
- 'Other London' airports – which we define as London as a whole (above) minus Heathrow; and
- The rest of the UK – this covers all airports in the UK minus the London airports.

To compare growth over time, we have reported the data as an index where 100 equals the number of movements in 1961.

[11] https://ext.eurocontrol.int/airport_corner_public/EGLL

[12] https://www.caa.co.uk/Data-and-analysis/UK-aviation-market/Airports/Datasets/UK-Airport-data/Airport-data-2017/

Figure 18 Movements over time – index

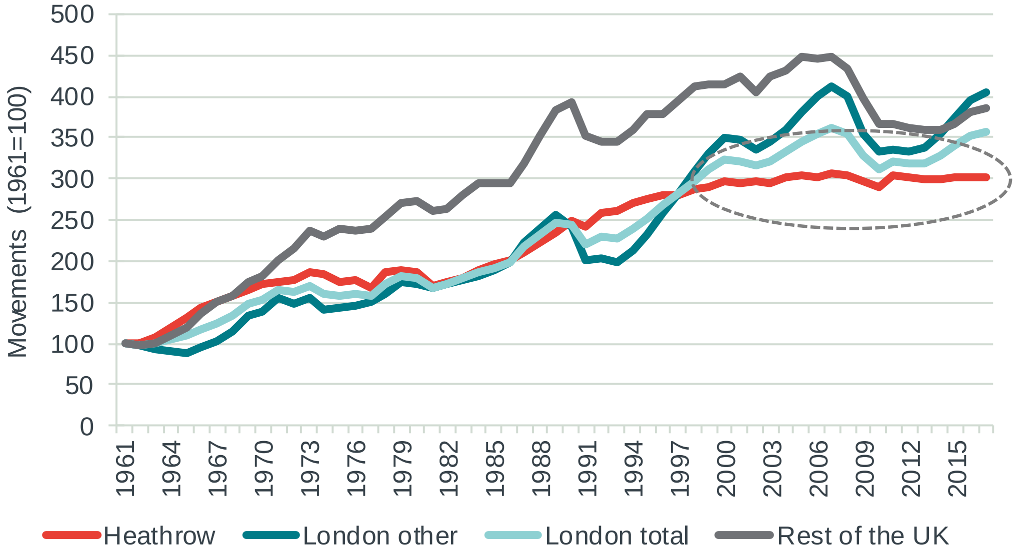

Source: Frontier analysis of CAA data

The results show that total demand has increased significantly over time. Up to 1990, the growth in movements at Heathrow was broadly in line with the growth at other London airports. However, since 2000, growth at Heathrow has largely plateaued. Since then, the total number of movements at Heathrow per annum has increased by around 9,000 in total (less than 2% growth over nearly 20 years), whereas at other London airports, it has increased by over 100,000 (equal to around 16% growth over the period).

As an aside, we note that Heathrow appears to have been more resilient to shocks in demand compared to other airports. For example, during the financial crisis, starting in 2008, Heathrow experienced a relatively modest dip in the number of movements, compared to much more significant reductions at other London airports and the rest of the UK. This is consistent with an airport that faces significant excess demand such that even during times of shock there is still enough demand to ensure that airlines will continue operating flights.

The chart below shows the growth in movements year on year at Heathrow, and also the average growth in movements per annum in each decade:

Figure 19 Growth in movements at Heathrow (year on year)

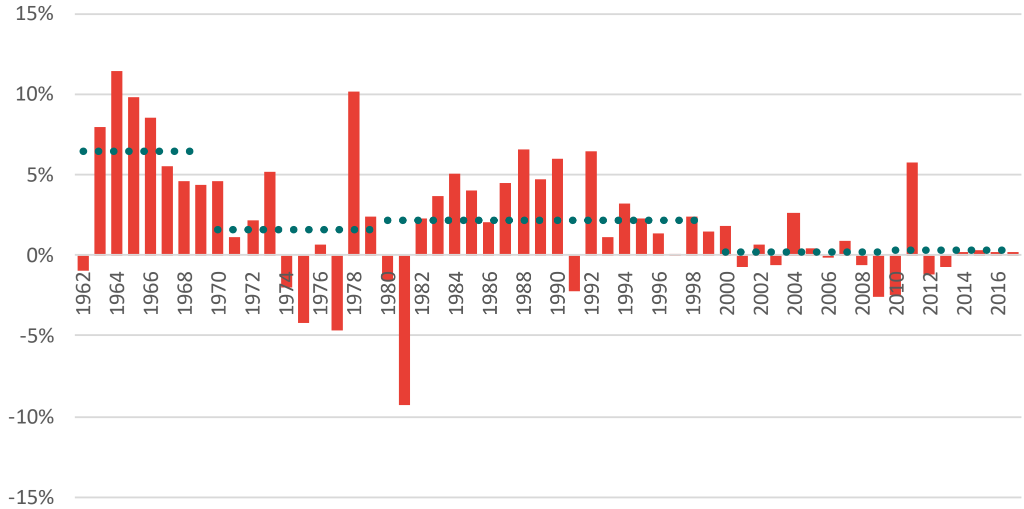

Source: Frontier analysis

Note: The dotted lines show the average growth rate in each decade

This chart also shows that since 2000 growth at Heathrow has slowed down considerably, and has been close to 0% per annum on average (0.02% on average in the 2000s and 0.03% on average in the 2010s). Based on this evidence, we consider it is reasonable to argue that Heathrow has been constrained since 2000.

2.2.2 Unconstrained growth

If Heathrow has been constrained since 2000, then how much would it have grown if it had been unconstrained?

Whereas in the previous step we focussed on the growth in movements – because the bottleneck at Heathrow is runway capacity – here we focus on the growth in passenger volumes. This is because any fare premium ultimately impacts on the ticket fares paid by individual passengers, and we need to estimate excess demand in passenger terms, not in movement terms.

The table below compares the average growth in passengers (CAGR) at Heathrow and other airports since 1975:

Figure 20 Passenger growth (CAGR)

	Heathrow	Other London	London total	Rest of the UK	Total UK
1970-75	6.7%	4.0%	5.9%	-8.3%	0.1%
1975-80	5.2%	10.1%	6.6%	6.9%	6.7%
1980-85	2.6%	7.1%	4.1%	3.9%	4.0%
1985-90	6.4%	8.1%	7.0%	9.3%	7.7%
1990-95	4.9%	2.7%	4.1%	6.3%	4.8%
1995-00	3.5%	12.5%	7.0%	6.5%	6.8%
2000-05	1.0%	5.0%	2.9%	8.1%	4.9%
2005-10	-0.6%	-1.4%	-1.0%	-2.5%	-1.6%
2010-15	2.7%	5.3%	3.9%	3.1%	3.6%
2015-17	2.0%	7.7%	4.9%	8.6%	6.4%
2000-2017	**1.1%**	**3.5%**	**2.3%**	**3.5%**	**2.7%**

Source: Frontier analysis based on CAA data

Since 2000, passenger volumes at Heathrow have only grown by 1.1% per annum on average. However, other London airports have grown by 3.5% on average. As mentioned, even at a constrained airport, passenger volumes can continue to grow by airlines operating larger aircraft with more seats, achieving higher load factors, and switching some short haul flights to long haul flights.

The higher growth at other London airports likely overstates the unconstrained rate of growth at those airports. This is precisely because Heathrow was constrained over the period, and we might expect that to some extent excess demand at Heathrow was overflowing to other London airports, supplementing their own organic growth rates.

At the combined London level, demand increased by around 2.3% per annum on average. Similarly, we might consider this to be an underestimate of the true unconstrained growth. This is because Heathrow is London's largest airport and it was not able to grow at an unconstrained rate over this period, and while some passengers may have overflowed from Heathrow to other London airports, others may have decided not to fly at all.

We believe that examining the growth at London airports as a system is a reasonable and conservative proxy for how Heathrow might have grown if it had not been constrained. First, we note that the DfT's National Air Passenger Allocation Model (NAPAM), which was also used by the Airports Commission when considering the question of expansion, considers London airports very much as a system. NAPAM is specified in such a way that any constraints at one airport will lead to passengers weighing up whether they still wishes to fly from their next most convenient airport instead, or whether on balance they decide to no longer fly because alternative airports are too inconvenient for them. This results in a dynamic whereby London airports gradually become constrained one after another as demand overflows from constrained airports to unconstrained airports, with some passengers choosing not to fly along the way because of the inconvenience of alternative airports. This gradually filling up is shown below. (It is through this choice dynamic that the DfT is able to forecast how many passengers will not be able to fly in the future because of constraints).

Figure 21 London airports are a system

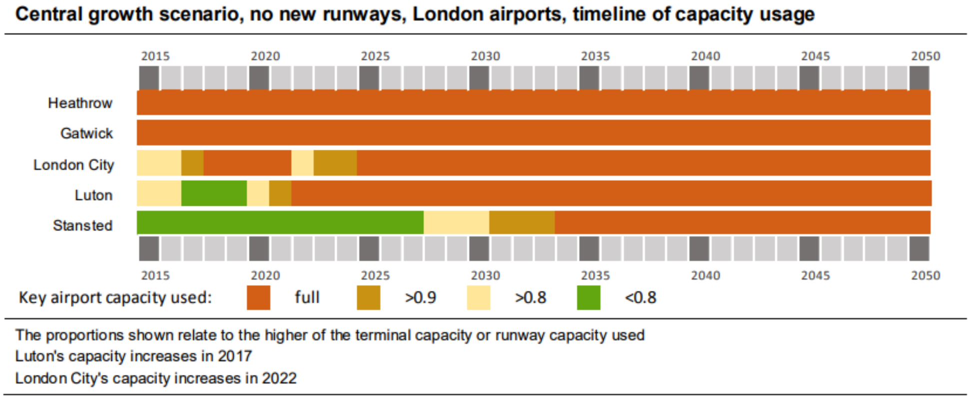

Source: Airports Commission Strategic Fit: Forecasts. https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/439687/strategic-fit-updated-forecasts.pdf

The large growth rates observed at some London airports in recent years may simply reflect that some of that demand was not able to fly from Heathrow. Or in other words, Heathrow would have grown faster than it did if it had had spare capacity, and other London airports would have grown less fast.

We believe it is conservative to assume for the purpose of this analysis that Heathrow would have grown at least as fast as London airports as a whole since 2000 (i.e. 2.3% CAGR). Heathrow has the best transport links and the highest service quality compared to the other London airports, such that if it had had spare capacity (and prices therefore did not have a congestion premium) it is reasonable to believe that many passengers would have preferred to have flown from there instead. This is also conservative because it does not take into account the passengers who would have flown to/from Heathrow but not to/from other London airports.

Therefore, our approach is to take passenger volumes at Heathrow in 2000 and to grow them each year by the growth in total passenger volumes across all London airports combined, instead of the constrained growth which was actually observed at Heathrow. The chart below compares the two growth profiles. The growth at London airports as a whole (in green) equates to an average of 2.3% over the period, while at Heathrow (in red), the average was 1.1%.

Figure 22 Comparison of yearly growth at Heathrow and London airports

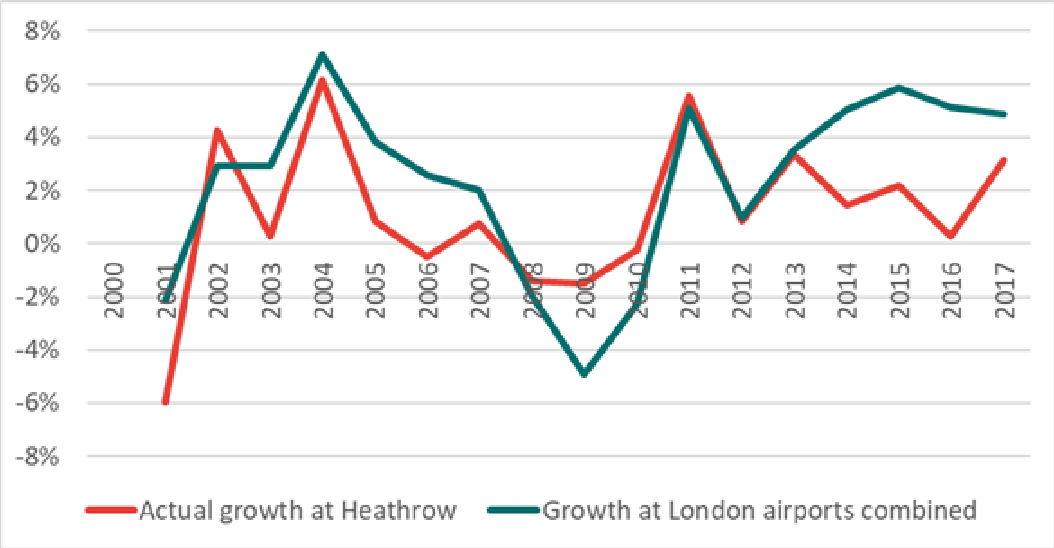

Source: Frontier analysis

2.2.3 Unserved demand

Applying the higher growth rate described above to passenger volumes since 2000 implies that demand would be much higher at Heathrow today if it were not constrained. On this basis, passenger volumes might have reached around 97 million in 2018, compared to just over 80 million in reality. This implies a level of unserved demand of around 17 mppa, which equates to around 17.5% of the total unconstrained demand in 2018.

Figure 23 Unserved demand could have been around 17 mppa in 2018

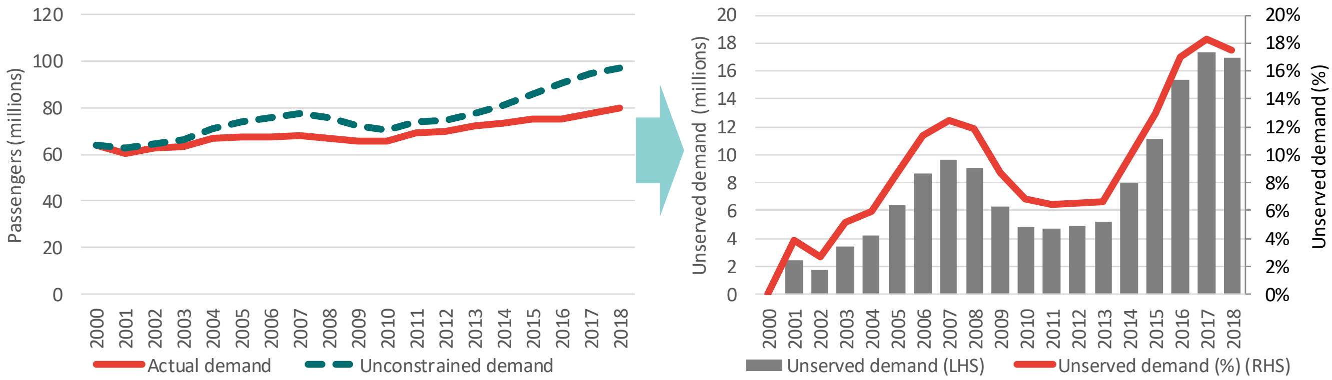

Source: Frontier analysis

The observation that our estimate of unserved demand decreases in some years is based on the fact that in some years Heathrow grew faster than London airports combined – e.g. it was more resilient during the financial crisis.

We note that not all of these 17 million passengers described above would have been O/D passengers. Applying the growth rate to total passengers at Heathrow, around one third of which is made up of transfer passengers, implies that some of this unserved demand (c. 6 million passengers) would have included transfer passengers. This means we are not suggesting that all of this excess demand is made up of O/D passengers.

As an aside, we note that this highlights that capacity constraints at Heathrow have limited not only the number of O/D passengers that can fly to and from Heathrow, but also the number of transfer passengers that connect via Heathrow. While some of the O/D demand may have switched to other London airports (although, in line with NAPAM, we note that some passengers will have decided not to fly), the transfer passengers cannot have switched to other London airports, because Heathrow is the only hub in the UK[13]. This demand must have therefore flown via other hubs instead, such as Schiphol, Paris Charles de Gaulle, and Frankfurt.[14] And transfer passengers are important for hub airports. They increase the total demand for all flights, which means that many routes benefit from an increase in frequency – which is to the benefit of O/D passengers who now face greater choice – and some routes may only be viable because of the transfer passengers. And clearly, having direct connections is beneficial for O/D passengers. In other words, not only is Heathrow smaller because of its constraints, it is also less well-connected.

We have applied the unconstrained growth rate to O/D and transfer passengers equally, meaning that by 2018, we estimate 17.5% of the total O/D demand that would have been at Heathrow under an unconstrained scenario was not able to fly.

2.2.4 Estimated congestion premium

The next step is to consider what increase in price (congestion premium) would have been required to price the excess demand, estimated above, out of the market. To do this, we have researched price elasticities of demand (PED). PEDs aim to quantify the relationship between price and demand. They are typically expressed as follows:

A 1% change in price leads to an x% change in demand

The table below sets out estimates of elasticities of demand with respect to price (reported as air fares in the table below) as used by the DfT when forecasting demand.

[13] Omitting the possibility of self-connection using two separate tickets, which occurs but is a very small share of the transfer market.

[14] In principle, some of these transfer passengers may have self-connected via other London airports. However, we believe this is likely to be very negligible. First, the other London airports have very limited long haul networks, limiting the connectivity for transfer passengers via those airports. And secondly, it is unclear how many passengers would choose to self-connect via other London airports – and run the risk of missing flights and losing bags etc. – when other hubs can provide more seamless connections with a through ticket.

Figure 24 The DfT's price elasticities of demand

		Elasticity with respect to	
Sector	**Share of base year demand**	**Income**	**Air fares**
UK business	7%	1.2	-0.2
UK leisure	47%	1.4	-0.7
Foreign business	6%	1.0	-0.2
Foreign leisure	19%	1.0	-0.7
International to international transfers	9%	0.5	-0.5
Domestic	12%	1.1	-0.5
Total	100%	1.2	-0.6

The Airports Commission change to the international to international transfer fares elasticity has been retained. The elasticity changed from -0.7 to -0.5, reflecting that it now relates to a broader market, following the inclusion of overseas hubs.

Source: DfT: UK Aviation Forecasts 2017 https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/781281/uk-aviation-forecasts-2017.pdf

The PEDs range from -0.2 for UK business passengers to -0.7 for UK and foreign leisure passengers. This implies that business passengers are less sensitive to price than leisure passengers. Or in other words, if prices were to increase, leisure passengers are more likely to stop flying compared to business passengers.

It follows that if the PED is greater in magnitude, then the price increase (congestion premium) required to price the excess demand out of the market will be lower. For example:

- Using a PED of -0.2: To price 17.5% of total demand out of the market, prices would have needed to have risen by 86%.
- Using a PED of -0.7: To price 17.5% of total demand out of the market, prices would have needed to have risen by 25%.

To be conservative, we have therefore used the PED of -0.7. By applying this rate in all years, the chart below shows how the congestion premium (in percentage terms) has evolved over time.

Figure 25 The congestion premium % over time

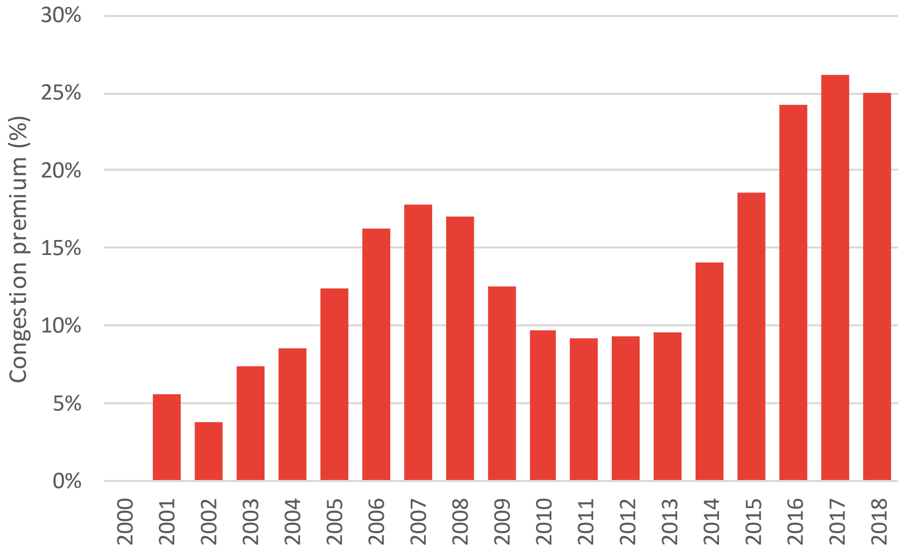

Source: Frontier analysis

The general pattern is that the implied premium has been increasing over time, as unconstrained demand growth has outstripped the actual historical growth on average. There are large decreases and increases in the premium in some years. As described above, this relates to the fact that in some years, Heathrow's actual constrained growth may have been greater than the total growth observed in London (e.g. following the financial crisis) while in other years, total London growth may have been significantly greater than the growth observed at Heathrow.

2.2.5 Total premium

We have also converted the percentages above into absolute figures on a per passenger basis and also in total across all passengers affected.

In doing this we have made a distinction between three types of passengers at Heathrow:

- 'Local' or point-to-point passengers: These are O/D passengers who fly to/from Heathrow on direct connections without any stopovers along the way – e.g. passengers flying from Heathrow to New York only;
- 'Beyond' or multi-stop passengers: These are O/D passengers who fly to/from Heathrow with at least one stopover along the way – e.g. passengers flying from Heathrow to New York and then on to Chicago; and
- Transfer passengers: These are passengers who neither start nor finish their journey at Heathrow, and instead connect at Heathrow – e.g. passengers flying from Dublin to Heathrow and then on New York.

The chart below shows the breakdown of total passengers volumes at Heathrow in 2018 across these three groups:

Figure 26 Breakdown of passenger volumes at Heathrow in 2018

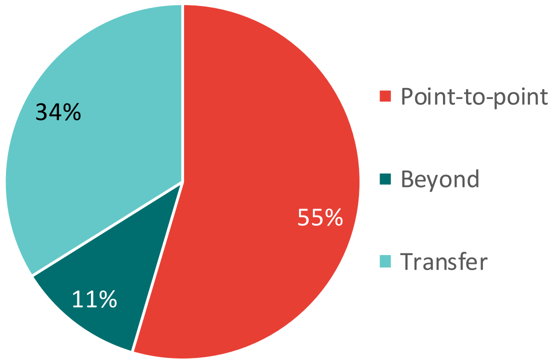

Source: Frontier analysis of passenger data from AirportIS and CAA data

Point-to-point passengers made up just over 55% of total passenger volumes in 2018. We have been conservative and have assumed that the premium only impacts on point-to-point passengers. In other words we assume there is no premium for beyond or transfer passengers.

The rationale is that transfer passengers may have multiple options for how they connect to their final destination. A passenger flying from Dublin to New York may be able to connect via Heathrow or a number of other hubs. Because passengers have multiple options, airlines offering a connection via Heathrow may find that any increase in their price will result in passengers choosing to fly via other hubs instead. This competition (and relatively limited differentiation) ensures that airlines may not be able to increase their prices even if there is excess demand at Heathrow. By extension, we conservatively assume that there is no premium for beyond passengers because they may also have multiple options to connect to their final destination, both between many different airlines at Heathrow, as well as possible options at other London airports too. However, for O/D passengers wishing to fly on direct routes, Heathrow may be the only option in London, and passengers may have no direct outside options. This is an illustration of the differentiation point made in our introduction. We are saying that Heathrow is less differentiated for transfer and beyond passengers than for O/D passengers.

We recognise that the picture may be mixed from route to route. A transfer passenger wishing to fly to particular end destination may only have one option if it is only connected to one hub, or if only one hub has a direct connection to the end destination – although two-stop connections could also be an option. Similarly, some direct connections at Heathrow may also appear at other airports in London, which may limit the ability of airlines at Heathrow to set higher prices – although as noted Heathrow still remains differentiated.

By assuming that the premium only applies to point-to-point passengers – just over half of the total passengers at Heathrow – we believe that our approach is conservative.

The table below sets out the average ticket fare for point-to-point passengers at Heathrow in 2018, split out by short haul and long haul. We have expressed the results as an average per return trip. By applying the premium percentages described in the previous step, we can then estimate the congestion premium. We have made a high level assumption that the premium percentages apply to short haul and long haul passengers equally. Finally, we have multiplied the estimated premium per trip by the number of trips.

Figure 27 Top down estimate of the congestion premium in 2018

	Average return ticket fare	Premium (%)	Premium per trip*	Total trips (m)	Total premium (£m)
Short haul	£237	25%	£47	13	£600
Long haul	£1,080	25%	£216	9	£1,990
Average / total	£593	25%	£119	22	£2,590

Source: Frontier analysis

**A premium of 25% implies that prices are 25% higher today than they would be if Heathrow were not constrained. The prices we observe today already include this premium, such that they effectively represent 125% of the unconstrained fare. The premium is therefore estimated as [£237 x 25%] / 125% = £47 for short haul, and [£1,080 x 25%] / 125% = £216 for long haul.*

Under this approach, the results suggest that the premium was equal to £2.6 billion in 2018.

2.3 Results

We have presented evidence to suggest that Heathrow first became constrained around the year 2000. Since then, passenger volumes have grown by 1.1% per annum. We believe it is conservative to argue that Heathrow could have grown at least as fast as the total growth observed across all London airports over the same period: 2.3% per annum on average. This means that by the year 2018, passenger volumes at an unconstrained Heathrow would have been around 97 million per annum, instead of the 80 million actually observed in 2018.

This implies that 17.5% of the total unconstrained demand being priced out of the market at Heathrow. Using a price elasticity of demand of -0.7 (the most conservative number used by the DfT in its forecasting analysis), we estimate that this implies a congestion premium of 25%.

We have conservatively assumed that only point-to-point passengers pay this premium. This represents just over half of all passengers at Heathrow. We believe airlines at Heathrow may not be able to increase prices for transfer or beyond passengers because, unlike point-to-point passengers, these passengers face many outside options.

We have therefore estimated that premium in 2018 is equal to:

- £47 per passenger making a return trip to a short haul destination, flying point-to-point;
- £216 per passenger making a return trip to a long haul destination, flying point-to-point; and
- £2.6 billion across all short haul and long haul point-to-point passengers.

The chart below repeats the analysis for the full period 2000-2018:

Figure 28 The congestion premium over time estimated on a top down basis

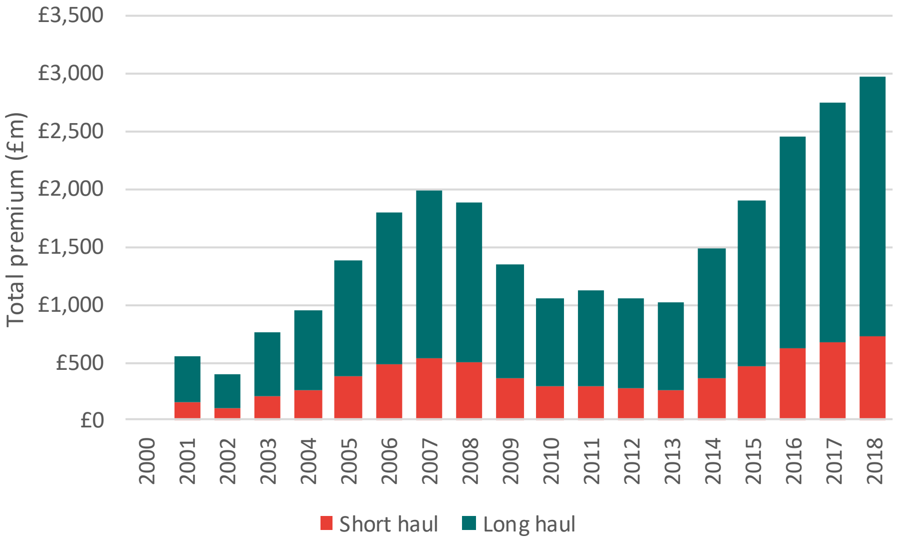

Source: Frontier analysis

3 ANALYSIS OF SLOT TRANSFER PRICES

3.1 Introduction

In this section, we estimate the congestion premium based on an analysis of the prices that airlines pay to acquire slots at Heathrow. As an introduction, we first describe:

- What is a slot?
- Slot transfers at Heathrow; and
- The conceptual link between slot transfer prices and the congestion premium.

We describe these points in turn.

3.1.1 What is a slot?

A runway slot is the right to use an airport's infrastructure to operate a flight at a particular time on a particular day of the week, for a full 'season'. There are two seasons in the year – summer and winter – where each season lasts 26 weeks. For example, an airline may hold a particular slot which gives it the right to take off or land at Heathrow at 9am every Monday for the summer season. An airline needs at a minimum a pair of slots to operate – i.e. one to land and one to take off again – as well as an equivalent pair of slots at the destination airport.

The total number of slots at an airport is determined by an interlocking range of restrictions, including runway capacity, terminal capacity and planning restrictions such as bans on night time flights. Heathrow has an annual cap of 480,000 movements a year, which equates to around 18,500 slots[15] in total.

In practice, operating a daily connection – a landing and a take-off every day of the year – technically equates to 28 individual slots, i.e. 7 slots for landing and 7 slots for take-off in the Summer season, and 7 slots for landing and 7 slots for take-off in the Winter season.

In the EU, slots are allocated to airlines based on the EU Slot Regulation[16], which in turn is based on IATA's Worldwide Slot Guidelines (WSG)[17]. This allocation is based on grandfathering/incumbency rights: an airline that presently holds a slot can continue to operate it in perpetuity, provided it uses it at least 80% of the time (the '80/20 rule'). Airlines do not pay for their 'primary allocation' of slots, nor do they pay to use the slot. This is not to be confused with airport charges, which are designed to recover the cost of the aeronautical services provided at the airport.

Free slots are placed in a slot 'pool'. This will include slots which may have never been used before, or which may have been released by a previous holder not satisfying the 80/20 rule. The pool will also include any new slots created following expansion. Airlines apply to the slot coordinator to obtain slots. This coordinator is

---

[15] 480,000 movements a year equates to around 9,200 flights a week, with 2 seasons: (480,000 / 52) x 2 = c18,000.

[16] https://ec.europa.eu/transport/modes/air/airports/slots_en

[17] https://www.iata.org/policy/slots/Pages/slot-guidelines.aspx

an independent body appointed by the airport. Scheduling guidelines require that 50% of slots in the pool are allocated to new entrants, with the other 50% going to incumbents. In the EU, each country has its own independent airport slot coordinator to facilitate the process.

With some limitations, the rules allow for secondary trading, i.e. transfers of slots between airlines, with side payments allowed to change hands in these transactions if the exchanged slots are perceived to have very different values.

3.1.2 Slot transfers at Heathrow

Heathrow has been operating at, or near to, 100% runway utilisation since the early 2000s. Since then, demand has continued to grow, but there have been few to no new slots to allow new airlines or incumbents to add more flights. As shown below, the number of slots entering into the pool has decreased considerably. This has coincided with a rise in the number of slot transfers between airlines.

At the same time we have seen rising incidence of financial transactions involving slots at Heathrow.

In 2016, Oman Air paid a “world record” $75m for a daily pair of slots at Heathrow[18]. In 2015, Virgin Atlantic completed a deal to raise £220m of investment capital by using its slots at Heathrow as collateral[19]. These transactions demonstrate that slots at Heathrow are have a significant financial value.

Peak slots at other major London airports may also valuable at some times. For example, in early 2019 Vueling purchased a bundle of slots at Gatwick airport for a fee of £4.5m. In 2013, Flybe sold a larger bundle of slots to easyJet for £20m.[20] However, the value per slot pair at Gatwick is significantly lower than that at Heathrow.

Figure 29 Very few new slots become available in the pool at Heathrow

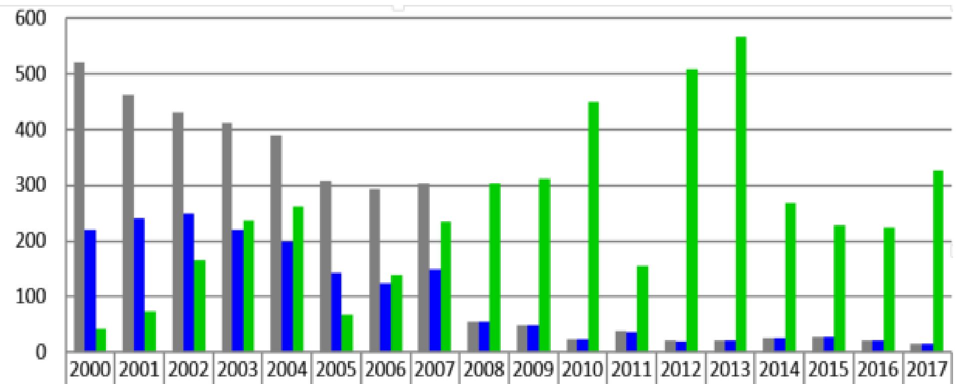

	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017
Pool Availability	520	462	432	412	390	308	293	303	54	48	23	38	22	22	26	28	22	16
Pool Allocation	219	241	249	221	199	143	123	148	54	48	23	36	20	22	26	28	22	16
Slot Transfers	42	74	166	236	262	68	139	235	304	313	449	156	508	568	268	228	224	326

Source: Heathrow Airport Holdings

[18] https://www.thetimes.co.uk/article/oman-breaks-heathrow-record-with-deal-for-slots-5mhdlzs23mn

[19] https://www.ft.com/content/40193024-9e8b-11e5-b45d-4812f209f861

[20] http://www.travelweekly.co.uk/articles/321156/flybe-sells-gatwick-slots-to-vueling-for-45m

Transferring a slot from one airline to another can benefit both parties. The seller may determine that it can earn more money from selling the slot to another airline than it can from operating the slot itself (given its discount rate and its estimated present day value of streams of profit in the future). The purchaser may believe that if it has the slot it can earn profits operating a route, and in principle it could be prepared to pay the selling airline a sum of money which is less than or equal to the net present value of its expected future profits on the route.

The table below shows that, when slots are traded, the purchasing airline typically uses the slot to fly larger planes, and on longer routes, than the selling airline. In this way slot transfers can lead to a more efficient use of limited capacity by increasing the number of passengers and passenger-kilometres flown from an airport, even with fixed runway capacity. The potential for reallocating runway capacity to more profitable uses is part of what makes the right to operate slots valuable.

Figure 30 When slots are traded, the new operator typically flies larger planes and longer routes

ACL

Economics of Slot Trading

	Before Trade	After Trade	Difference
Seats per slot	135	255	+90%
Sector length	575 km	6800 km	x 12
ASK per slot	77,625	1,734,000	x 22

- Based on a sample of Heathrow trades, excluding temporary lease agreements
- ASK = Available Seat Kilometers

Source: Airport Coordination Limited (2009).

While slots are often used more efficiently after trading, this does not imply that slot trading improves efficiency compared to the unconstrained situation.

3.1.3 The link between slot prices and the congestion premium

At an airport that always has spare slots at any time of day, slots will have little to no value. This is because there is plenty of spare of capacity and there are no restrictions preventing any airline from operating extra flights at the airport if it wished to do so: it can acquire a new slot from the pool for free. This is true even of slots at peak times when fares may be higher. Airlines may earn a greater contribution towards their fixed costs at these times, but that is not *because* of slot scarcity. It is because of more restricted seat capacity, as explained previously. Despite the higher fares, airlines have no reason to place a financial value on the slots *per se*, because more slots are freely available even at these times.

But at a congested airport, with no slots free in the pool, an airline can only acquire a slot by obtaining it from a rival airline. This scarcity therefore makes the slots valuable.

In a competitive market with spare capacity, we would expect average prices on a route to reflect the long run marginal cost of operating the route, which covers variable operational costs and the long run of capacity, including aircraft finance and route-related overheads. If an airline is prepared to pay additional sums of money to acquire a slot this must signify that the airline believes it can use that slot to earn a premium above this long run marginal cost. If there were no premium, then paying for the slot would make no sense. It would imply incurring an extra cost which the airline would be unable to recover.

Note that it is sometimes argued that airlines should not be made to pay for slots, because this would drive up airfares. This argument is without economic merit. The price paid for a slot has no effect on the forward-looking cost of offering more seats on a given route, or the passenger demand to fill them. So fares are unaffected by slot prices. The direction of causation is unambiguously the other way round. If demand is high and slots are scare, fares will rise. If possession of a slot is the only way for an airline to participate in this situation, airlines will pay for the slot. If there is competition to obtain the slot, airlines would be willing to pay up to the "excess value" they perceive they could earn from its possession.

In practice, however, not all slots are equal: demand for flights tends to peak in the early morning. A moderately congested airport could have no spare capacity at peak times of day, but plenty of spare slots at other times of the day. This is illustrated below. In this case the peak time slots become valuable, as airlines compete to operate the most popular slots, whereas at other times of day enough slots are available in the pool so that entrants do not find it worthwhile to make offers to buy incumbents' slot rights.

Figure 31 Demand peaks at certain times of day, making some slots more valuable

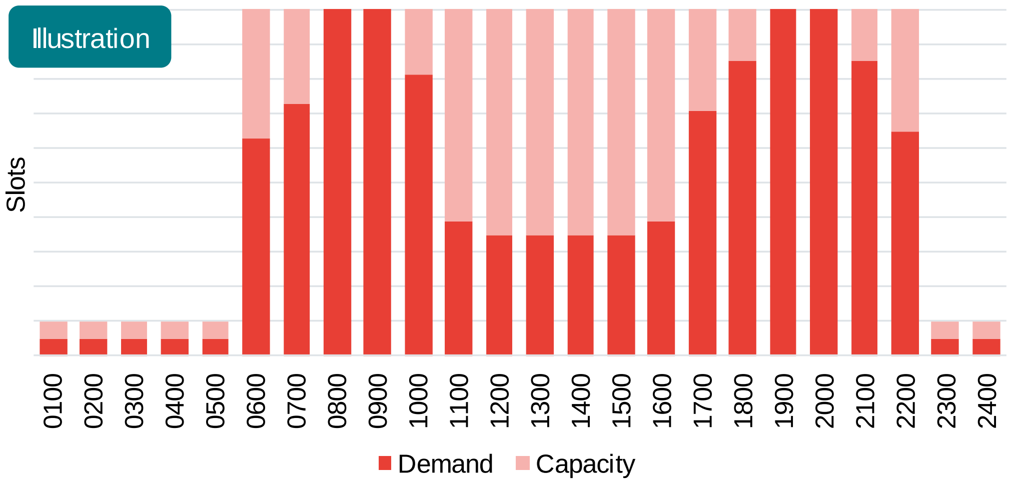

Source: Frontier

Heathrow is congested throughout the whole day. All slots are valuable, but excess demand will be higher at peak hours, and this means these slots will probably be worth more.

Our approach in this workstream is to ask what airlines must believe about the premium at Heathrow to be willing to pay the slot prices that we observe.

3.2 Approach

A high level overview of our approach is set out below.

Figure 32 Slot price transfer analysis

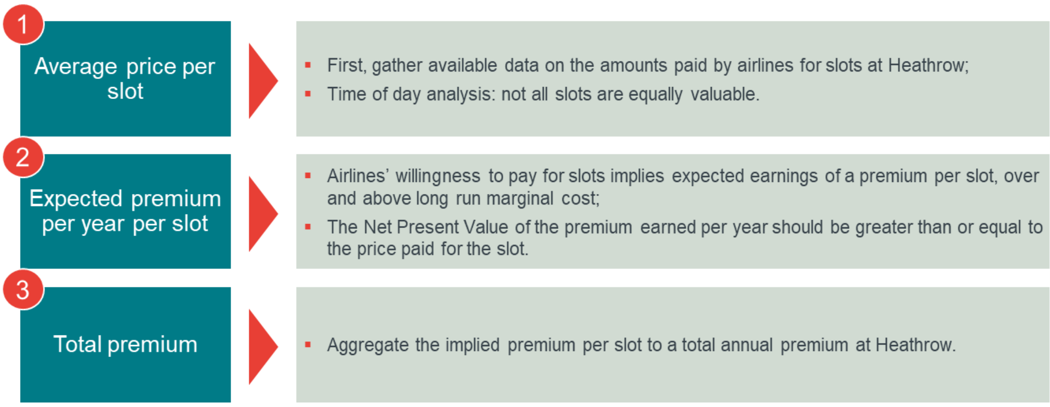

We describe these points in turn in more detail.

3.2.1 Slot transfer price data

Details of slot transaction fees are usually opaque and not readily publicly available. Airport Coordination Limited (ACL), the body responsible for slot allocation and coordination at Heathrow and a large number of airports, keeps a database of completed slot trades[21]. However, this does not contain slot transfer prices. Instead we have to rely on slot prices as they are reported in the media, a sample of which has been collated by HAL. We recognise that this sample may not be representative or always reliable, but it is the only information available. In any event its use in this context is intended to be in confirmation of alternative valuations of the congestion premium, not as a primary source of those estimates.

[21] https://www.acl-uk.org/completed-trades/

Figure 33 Slot price dataset

Data supplied by HAL						Frontier calculations		
Date	Purchaser	Previous Owner	Daily Slots	Total Price	Av. Price per daily slot	Number of slots	Average price per slot (£1000)	Total value of slots (£bn)
Feb-13	EY	9W	1x early am, 2x pm	$70m	$23.3m	84	£533	£10
Aug-14	EY	AZ	3x pm, 2x evening	€60m	€12m	133	£364	£7
Feb-15	TK	SK	1x pm	$22m	$22m	28	£514	£9
Feb-15	AA	SK	1x early am	$60m	$60m	28	£1,402	£26
Oct-15	DL	AF/KL	6x am/early pm	$276m	$46m	168	£1,075	£20
Feb-16	WY	AF/KL	1x early am	$75m	$75m	28	£1,983	£37
Jan-17	DL	OY	5/wk am	$19.5m	$27.3m	20	£757	£14
Mar-17	AA	SK	1x am, 1x early pm	$75m	$37.5m	56	£1,039	£19
						AVERAGE	£812	£15
						MINIMUM	£364	£7

Source: Heathrow slot price data, ACL completed slot trades database, Frontier calculations.

Note: The record Oman Air $75m purchase is excluded from the overall results as an outlier.

The observed transactions cover bundles of slots. For example, the Oman Air transfer was for a daily slot pair – which as mentioned technically amounts to 28 individual slots – while other transfers included multiple daily slot pairs.

For comparability, we have divided the total prices paid by the number of individual slots included in the deal. For example, in February 2013, Etihad (EY) bought a bundle of slots from Jet Airways (9W) equal to 84 individual slots for $70 million. This implies a price per slot of £533,000.

The sample contains a mix of slots at different times of day. The time of day of the flight is a strong explanatory factor for the value of the slot. As shown below, midday/early afternoon slots are approximately 30% cheaper than early morning slots, and evening slots are around 50% cheaper against the peak early morning slots.

Figure 34 Slot prices: time-of-day discounts

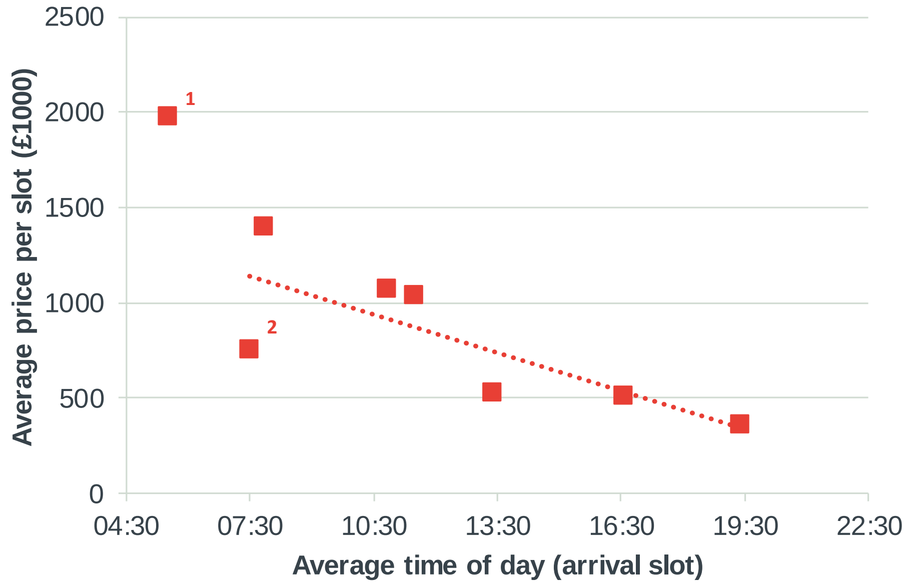

Source: Heathrow slot price data, ACL completed slot trades database, Frontier calculations.

Notes: 1. This is the record Oman Air $75m purchase, which we exclude from the analysis as an outlier (also excluded from trendline)

2. This is a weekday-only slot, whereas all other data points are for 7-days-a-week slots: higher prices at weekends could explain the lower price of this slot.

Using this data we seek to estimate the total premium across all flights at Heathrow per annum. This means that we need to estimate the price of the average slot at Heathrow and then scale this up to cover the 480,000 flights (c18,500 slots) per annum.

There could be a bias in selecting the average price paid if it is the case that airlines only buy slots which are especially valuable. However, the fact that the slot transfers listed above cover a wide spread in time of day, means that we do not consider this sample to be particularly biased. Using the data above, we have estimated two figures:

- The average price per slot in the sample above, excluding the Oman Air transfer which we consider to be an outlier: This is equal to £812,000 per slot. If we assume this is representative of all slots at Heathrow on average, then this implies that the total value of all slots at Heathrow adds up to around £15 billion; and
- The minimum price per slot in the sample above: This is equal to £364,000 per slot. Making a more conservative assumption that this is representative of the average value of all slots at Heathrow, this implies a grand total value of £7 billion.

As mentioned above, airlines can continue to operate slots in perpetuity as long as they satisfy the 80/20 rule. Therefore, the total value of slots does not give an indication of the premium paid by passengers each year, but rather the net present value (NPV) of the future premia paid by passengers.

3.2.2 Expected premium per year per slot

The premium that the airline expects to earn when using the slot is a cashflow over a number of years, which in order to make commercial sense needs to be greater than or equal to the upfront fee that the airline pays for the slot.

In this analysis we take the approach of calculating the Net Present Value (NPV) of the premium earned per year per slot. This is shown below.

$$Net\ Present\ Value = \sum_{t=0}^{N} \frac{P}{(1+i)^t}$$

Where:

- *N* is the number of time periods over which the airline expects to operate the slot and earn the premium (the asset lifetime – which in principle could be infinite if the airline satisfies the 80-20 rule);
- *P* is the annual premium;
- *i* is the discount rate on premia earned in future time periods.

Figure 35 Slot transfer price analysis: Net Present Value calculation

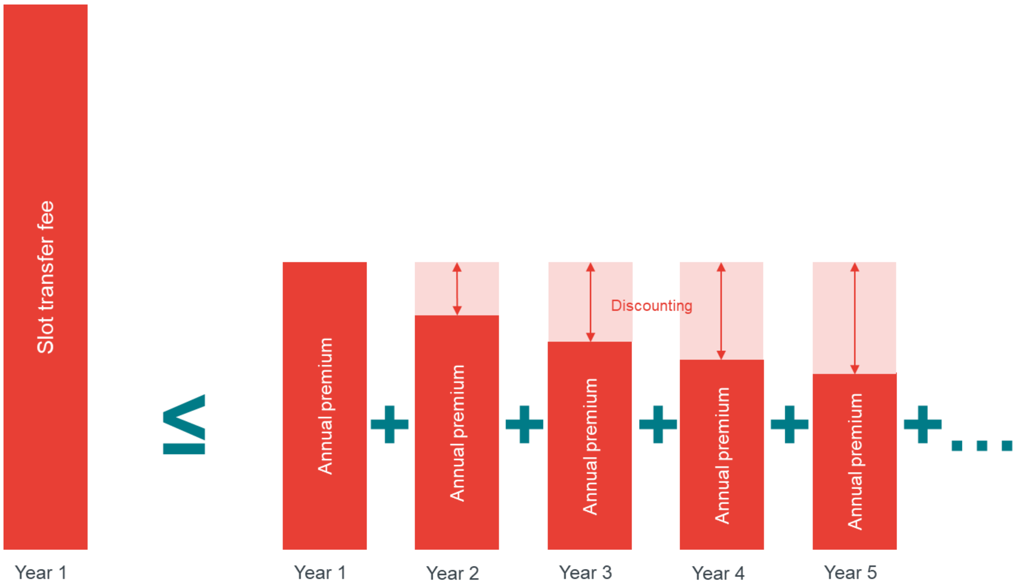

Source: Frontier

For the slot transfer fee we use the average (and, as a cross-check, the minimum) observed price paid per slot, calculated from the sample of slot transactions. This means that the annual premium is also calculated on a per-slot basis.

The calculation requires several assumptions:

- **Premium per year:** We assume the premium earned per slot per year to be constant over the asset lifetime, i.e. the length of time that the purchasing airline expects to continue using the slot. This is a simplifying assumption, because in reality we might expect the premium to generally increase over time as demand continues to grow.

  We note that Heathrow is set to expand with runway capacity increasing by more than 50% from around 2026-27. If the current slot allocation rules are used to allocate this extra capacity, this will result in many airlines receiving slots for free. As extra capacity comes into the market, we would expect the premium to reduce (and possibly disappear), hence reducing the net present value of slots today.

- **Asset lifetime:** When acquiring a slot, the purchasing airline is likely to make an assumption about the timeframe over which it expects to recover its initial outlay. As discussed above we expect the asset lifetime of a slot bought in the period 2013-2017 (the timeframe of our slot data sample) to be linked to capacity expansion which will make slots less valuable. If the airline expects to recover the transfer fee in 5 years, this implies a greater premium than if it expects to recover the premium in 15 years. We have estimated the premium using three different time horizons:
  - 5 years;
  - 10 years – this would be consistent with an airline buying a slot in 2017 (the last year of data in our sample) and expecting to recover the transfer free by the time the runway opens in around 2027; and
  - 15 years.

- **Discount rate:** The central assumption for the discount rate in our analysis is 7.5%. This is equal to the historical average Weighted Average Cost of Capital (WACC) in the global airline sector, based on an IATA report.[22] This is shown below.

Figure 36 Airline WACC

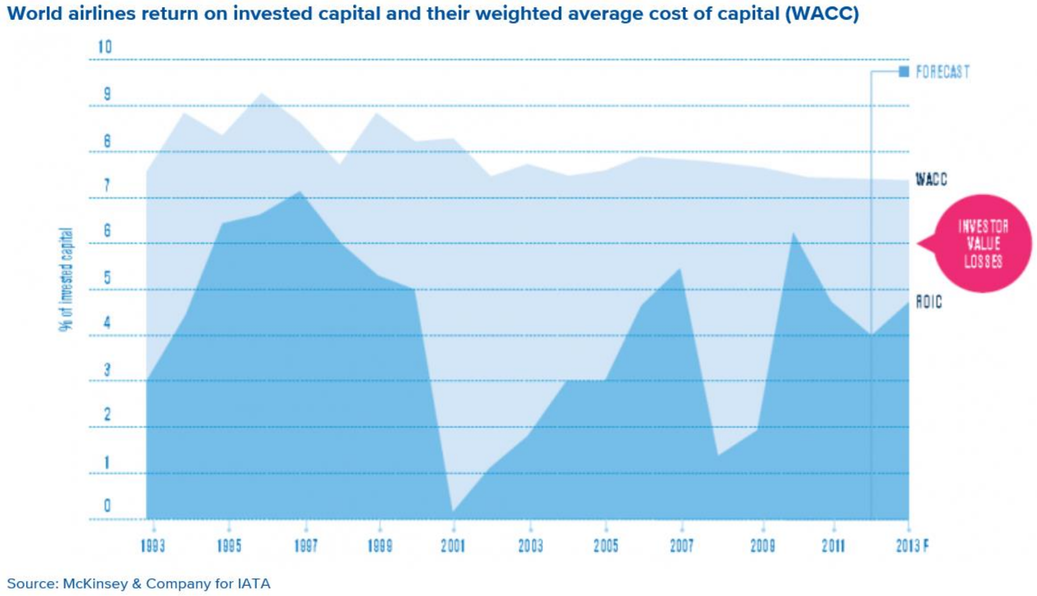

Source: IATA

Using a higher discount rate implies a larger premium per annum for the same slot transfer fee, as future expected cashflows are discounted more heavily. As a sensitivity, we have calculated the total premium using alternative assumptions of 7% and 8% for the discount rate. Based on the chart above, the WACC has fallen within this range since just after 2000. However, to broaden the range of outcomes modelled, we have also repeated the analysis using a WACC of 6% and 9%.

We consider this estimate to be conservative. The WACC figures above are based on actual outturn performance across the whole airline industry. However, buying slots could be considered more risky than business-as-usual conditions. First, the sector is already relatively risk as evidenced by many recent airline collapses, such as Flybmi and Monarch in the UK, and WOW air, Air Berlin, Alitalia and Germania amongst others in the rest of Europe. Buying slots adds to the airline cost base, and it is uncertain whether the airline will go on to recoup the value of the investment, especially given that expansion at Heathrow will create thousands of extra slots ‘for free’, based on the current WSG rules. Therefore, investors may possibly seek a greater return when investing in slots at Heathrow.

As discussed, 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. This gives us a lower bound

[22] https://centreforaviation.com/analysis/reports/airline-profitability-airlines-can-no-longer-afford-to-be-the-poor-relations-of-aviation-117521

estimate for the annual premium: in practice, the airline could expect to earn a greater annual premium on the slot being purchased.

3.2.3 Implied total premium

The calculations described above produce an estimate of the average premium per individual slot per annum. We have then scaled this up to the 480,000 flights / c18,500 slots at Heathrow per annum which gives us an estimate of the total premium across all slots at Heathrow.

The figure below sets out these calculations

- Average price per slot: we use the average observed price and also the minimum observed price as a cross-check;
- Discount rate: we use 7.5% as our central assumption. We also repeat the analysis using WACCs of 6%, 7%, 8% and 9% as sensitivities;
- For a given premium per slot, we then calculate the discounted cashflow earned in each year from that slot;
- For the assumed asset lifetime (we use 5, 10 and 15 years), we then add up the discounted cashflows in each year to calculate the Net Present Value (NPV);
- We set the implied NPV of the premium per slot equal to the initial average price paid – in principle it should be greater than or equal to the price paid;
- We then multiply the implied premium per slot per annum by the total number of slots at Heathrow to calculate the implied total premium per annum.

Figure 37 Slot premium analysis

Assumptions	
WACC	7.5%
Total annual movements	480000
Total annual slots	18462
Price per slot (£1000)	£812

The annual premium per slot is calculated as the amount that equalises the upfront price paid and the 10-year Net Present Value

Calculations										
Year	1	2	3	4	5	6	7	8	9	10
Discount factor	100.0%	93.0%	86.5%	80.5%	74.9%	69.7%	64.8%	60.3%	56.1%	52.2%
Annual premium per slot (£1000)	£110	£110	£110	£110	£110	£110	£110	£110	£110	£110
Discounted (£1000)	£110	£102	£95	£89	£82	£77	£71	£66	£62	£57
Cumulative cash flow (£1000)	£110	£212	£308	£396	£479	£555	£626	£693	£754	£812

Total premium	
Annual premium per slot (£1000)	£110
Total annual slots	18462
Total premium per annum (£m)	£2,031

Premium per passenger movement	
Total number return trips (m)	22
Annual premium per return trip (£)	£92

Source: Frontier

3.3 Results

The table below sets out our results. We estimate that if we take the average price paid per slot in the sample above (excluding Oman Air), which was £812,000, and assume an asset lifetime of 10 years and a WACC of 7.5%, the annualised value is equal to **£2.0 billion**. The implied premium would therefore be c. £2.0 billion. This is our central estimate, however we still believe it to be conservative due to the WACC and asset lifetime assumptions.

If we take the lower price per slot (the minimum in the sample), which was £364,000, then assuming a 10 year time horizon and a WACC of 7.5%, the premium per annum would be at least **£910 million**.

Figure 38 Slot transfer prices: Implied total premium per year (£m)

		Price per slot									
		Average observed					Minimum observed				
	WACC	*6.0%*	*7.0%*	*7.5%*	*8.0%*	*9.0%*	*6.0%*	*7.0%*	*7.5%*	*8.0%*	*9.0%*
Asset lifetime	5 year	£3,357	£3,416	£3,446	£3,476	£3,535	£1,504	£1,530	£1,544	£1,557	£1,584
	10 year	£1,921	£1,994	£2,031	£2,068	£2,143	£861	£893	£910	£926	£960
	15 year	£1,456	£1,538	£1,580	£1,621	£1,706	£652	£689	£708	£726	£764

Source: Heathrow slot price data, ACL completed slot trades database, Frontier calculations.

As previously discussed, we believe that the premium is paid only by point-to-point passengers. Given that there were around 22 million point-to-point return trips made at Heathrow in 2018, this implies an average premium per return trip of approximately:

- **£91** per passenger making a return trip (using our central estimate); or
- **£42** per passenger making a return trip (using our minimum estimate).

4 ECONOMETRIC ANALYSIS

4.1 Introduction

In this section, we describe the econometric analysis we have carried out to estimate the congestion premium at Heathrow. This involves comparing ticket prices at Heathrow with those at other large airports in Europe, and quantitatively explaining the drivers of any differences.

The scatter chart below shows that ticket prices for point-to-point passengers are generally higher at Heathrow compared to those at other large European airports. It also shows that there is a positive correlation between fares and distance travelled, but that there is still large variation in fares once those factors are taken into account.

Figure 39 2018 average fares by distance – point to point passengers only

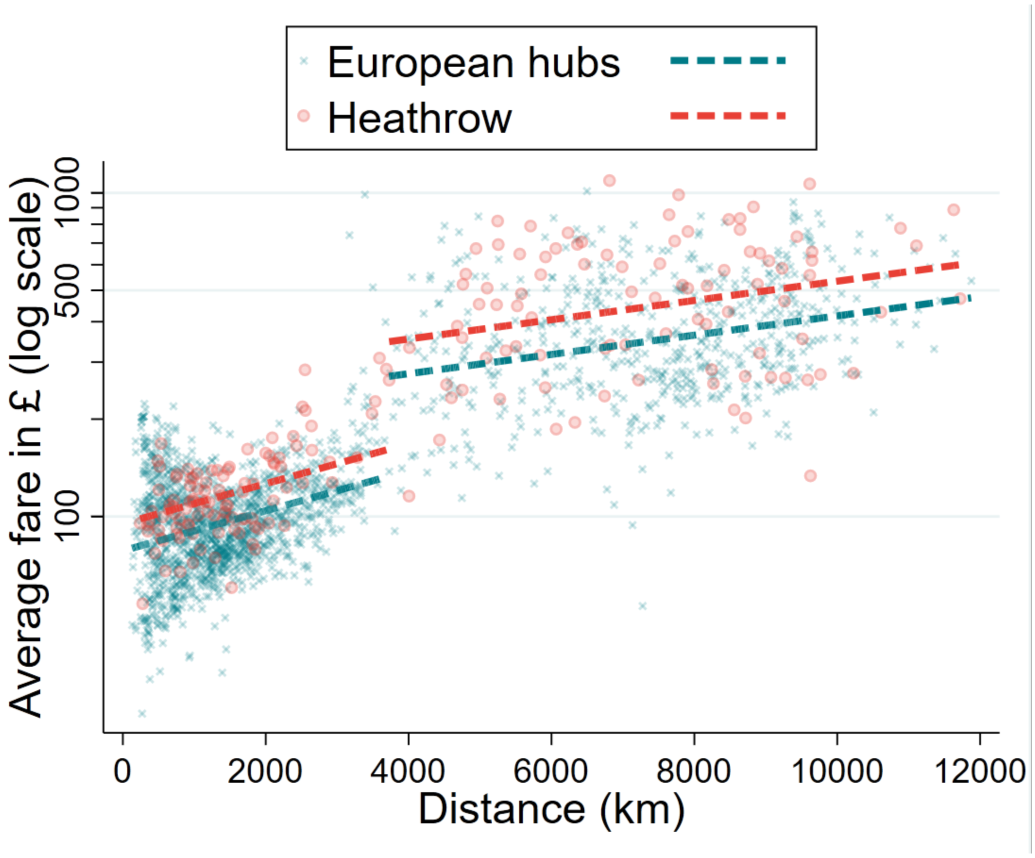

Source: Frontier analysis of IATA data

Note: Each point represents a route from either Heathrow or a European long haul hub in terms of distance and average fare. The dashed lines represent a linear fit of log prices with respect to distance, fitted separately for short and long haul flights.

On average, 2018 fares on short haul and long haul routes from Heathrow are 23% and 27% higher respectively than from other routes in our sample. Controlling for the distance of those routes gives a short haul premium of 21% and a long haul premium of 28% - the gap between the dashed lines in the above chart. Clearly, there are many factors that impact on price, and it would not be accurate to argue that all of this difference is attributable to congestion.

Figure 40 Many factors impact on prices

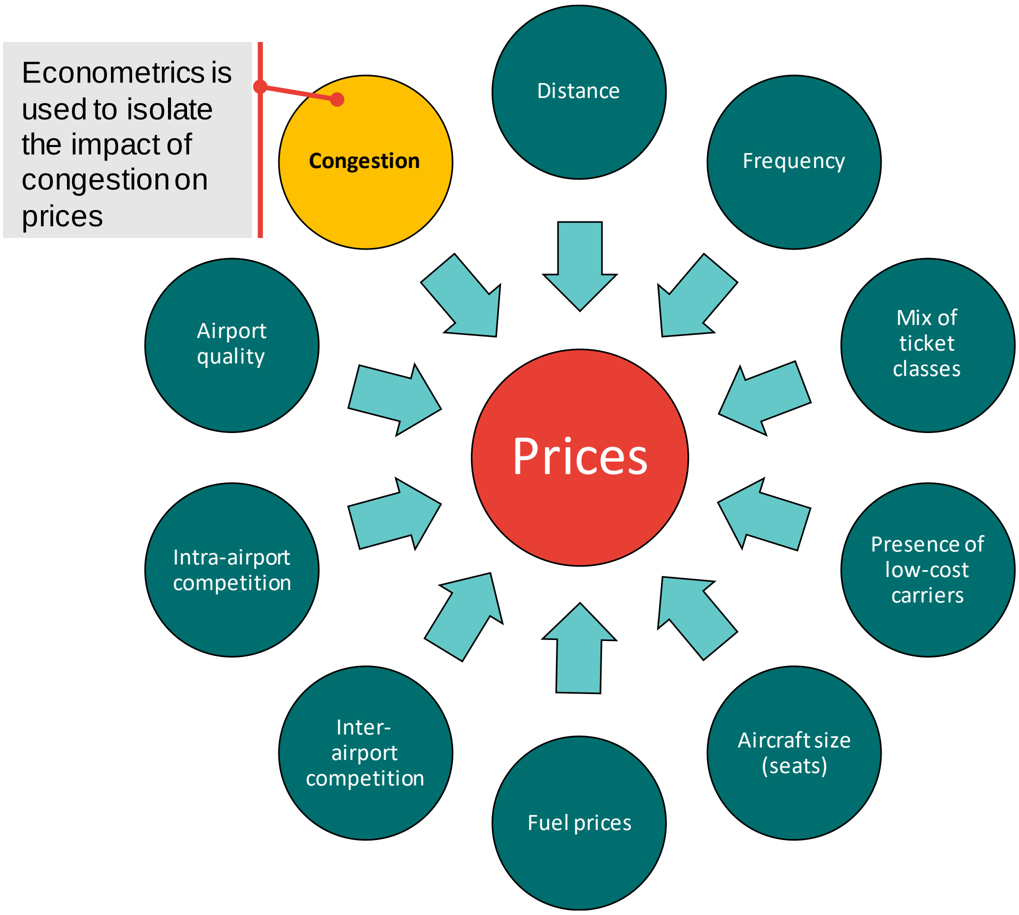

Using econometric analysis we aim to explain ticket prices as a function of these various factors and isolate the impact of congestion. This is our third report in which we estimate the premium using econometric analysis. This version represents our most comprehensive yet, as it contains:

- More years: In previous versions, we estimated the premium using just one single year of data. In this updated analysis we estimate the premium using a longer time series of data over the period 2011-2018. This covers all of the years for which we have consistent data on ticket prices.
- More comparator airports: This analysis involves comparing ticket prices at Heathrow with those at other relevant airports. Our analysis now includes more comparator airports. Our sample of airports covers the largest 10 airports in Europe, measured in terms of the number of long haul departures in 2018.
- More control variables: Econometric models can be accused of 'omitted variable bias' (OVB). Given that there is no one "correct" model for fares it is always possible to consider more and more possible explanatory variables. Hence any model can be accused of having "omitted variables" But OVB only occurs if the omitted variable is both an important explanatory driver of price *and* also if the omitted variable is correlated with the congestion we are trying to measure. By way of illustration, the simple average ticket fare at Heathrow is greater than that at Gatwick. This is partly explained by the fact that Heathrow

has many more long haul flights, so the average flight distance is longer than at Gatwick. But Heathrow is also more congested. So a model which does not include distance would not be very robust and could be subject to OVB. By contrast, comparing fares over time requires an understanding of the movement in fuel costs. So a model omitting fuel costs may not be a good predictor of fares. But fuel cost variation over time is not correlated with congestion at a given airport, so its omission should not lead to bias in the estimate of congestion premium. Nevertheless, in this analysis, we have done our best to include all relevant explanatory factors of fares and to ensure nothing that might be correlated with congestion has been omitted.

- More approaches to capturing congestion: In our previous models, we captured congestion at Heathrow by using a Heathrow 'dummy' variable. This argues that if we have included all of the relevant factors which impact on price, then any remaining difference in prices at Heathrow can be attributed entirely to congestion. In this analysis, we have expanded our approach to include 4 separate approaches to capturing congestion, including explicit measures of congestion. We describe these points in more detail below. However, as discussed below in more detail, the dummy variable approach remains our preferred approach.

We note that econometrics is not a perfect science because there will never be one single model which can perfectly explain ticket prices. However, the purpose of this exercise is to demonstrate that we have estimated the premium using many models which vary in terms of the number of variables included and also in terms of how we capture congestion, and all of them point to there being a considerable premium at Heathrow. We have produced a range of estimates which consistently point to a congestion premium of around £2.4 billion in 2018, equivalent to a premium of around 17% for short haul routes and 25% for long haul routes. This equates to around:

- £34 per passenger for a return trip to short haul destinations (flying point-to-point); and
- £217 per passenger for a return trip to long haul destinations (flying point-to-point).

The rest of this section is structured as follows:

- Data: We describe the data that we have used;
- Approaches to capturing congestion: We describe our different approaches to isolating the impact of congestion at Heathrow on prices;
- Model specification: We describe our approach including how we progressively add in more explanatory variables, split out short haul and long haul, and pool the data across multiple years.
- Results: We summarise our results; and
- Conclusions: We then present our overall conclusions.

In **Annex A** we set out more details underlying our analysis, including the full tables of coefficient estimates in each of our models.

4.2 Data

Our source of data on ticket prices is IATA's Airport Intelligence Service (AirportIS). A full description of the data can be found in Annex A. We have price data reported on individual routes, split out by point-to-point, beyond and transfer passengers[23], for each year in the period 2011-2018. The prices are reported on an average calendar year basis.

We compare prices at Heathrow with those at 9 other airports. Collectively, Heathrow and these other 9 airports make up the 10 largest airports in Europe by the number of long haul departures in 2018. This ensures that we are comparing prices at Heathrow with those at other airports which are broadly comparable in size (noting that Heathrow is the largest airport in the sample) and airports that are configured like Heathrow to offer long haul services

Figure 41 Our sample of airports

Code	Airport	Passengers (m)	Movements ('000s)	Connections*
LHR	Heathrow	80	476	176
CDG	Paris	72	481	239
AMS	Amsterdam	71	499	236
FRA	Frankfurt	70	512	242
MAD	Madrid	58	410	180
MUC	Munich	46	413	194
LGW	Gatwick	46	280	175
FCO	Rome	43	308	149
DUB	Dublin	31	226	142
ZRH	Zurich	31	278	143

Source: CAA; Royal Schiphol Group; Groupe ADP; IAA; Assaeoporti; Fraport AG; Aena; ADV; Zurich Airport; Frontier analysis of OAG Schedules Analyser
Connections is defined here as the number of destinations which were served with at least 100 departures in 2018 based on OAG Schedules Analyser

Each row in the data constitutes an airport-to-airport pair in a year between 2011 and 2018, along with the number of passengers and the average fare paid. We have supplemented this data with various explanatory variables. This includes the following:

[23] The distinction between point-to-point, beyond and transfer passengers is explained in Section 2.2.5

Variable	Description	Source
Distance	We have included a control for distance in kilometres of the route flown. This is because there is a clear relationship between distance and prices (as shown in Figure 39 above);	
Frequency (own)	Expressed as the average number of departures per day in each year and included in regressions in log form. Own frequency refers to the number of departures from the airport to the particular destination in question. This variable is intended to capture the supply of seats on a given route, and the established relationship between demand and frequency in transport economics.	OAG Schedules Analyser
Frequency (other)	Expressed as the average number of departures per day in each year and included in regressions in log form. "Other" frequency measures the number of flights *to* that particular destination *from all other* airports in the sample, reflecting its desirability and popularity as a destination beyond what is indicated by the own-route frequency. Since this variable is included log form, it implies that all destinations must be served by at least one other airport in the sample. In other words it requires that a route is an overlapping destination with another origin.	OAG Schedules Analyser
Transfer share	Transfer share measures the proportion of passengers on a given route who are transfer passengers. This is intended to proxy for demand on a route beyond the number of passengers travelling directly.	AirportIS
LCC share	LCC share refers to the proportion of seats on a route flown by Low Cost Carriers (LCCs). While the fares paid by LCC passengers themselves do not appear in the data due to limitations of the source dataset, the presence of LCCs on a route gives mainline customers the option of a lower price, so would be expected to drive down prices.	OAG Schedules Analyser
Business class share	Business class share measures the proportion of passengers on a flight who are estimated to have flown business class or above. This is because the mix of passenger classes is assumed to influence overall fares, and potentially proxy for overall demand. Note that, for the purposes of this analysis, premium economy has been treated as business class.	OAG Traffic Analyser
European jet fuel price	European jet fuel price is a common measure of fuel prices (in logs) across all airports in a given year, and is a key source of costs for airlines. Since we are already including year dummies, we would expect this only to decompose some of the already controlled-for sample-wide trends.	Bloomberg (JET1NEFB)

Variable	Description	Source
Average seats (plane size)	Average seats measures the average number of seats per departure on a given route, as a measure of the supply available to be filled, although this is itself arguably a potential response to the congestion premium.	OAG Schedules Analyser
Airport competition	Airport competition is a measure of the proportion of seats between two cities that is served by the airport-to-airport pair. For example, it reflects the proportion of seats between London and New York covered by the Heathrow-JFK route.	Frontier analysis of OAG Schedules Analyser
Route competition	Route-level competition measures the extent to which airlines on a route dominate or compete with other airlines. There are two variables we have used: firstly, we include a dummy for routes served by only a single carrier; secondly, we include a Herfindahl-Hirschman Index (HHI)[24] score. Together, these two variables imply an overall relationship between carrier market shares at a route level and fares, while allowing for the single carrier market to be a special case.	Frontier analysis of OAG Schedules Analyser
Skytrax rank (airport quality)	Skytrax rank is a ranking of airport service quality across the world. The lower the score, the better the airport's quality in a given year and vice versa. Since Skytrax only reports the top 100 airports, we have at times assumed airports that were unlisted to be ranked 100.	Skytrax

4.3 Approaches to capturing congestion

As described above, in our previous analyses we estimated the congestion premium by using a Heathrow dummy variable. This assumes that all statistical differences between Heathrow and comparable airports, after controlling for other drivers of fares in the model, are attributable to capacity constraints at Heathrow.

In FTI's critique of our approach it argued that:

> *"The 'residual approach'…ignores the possibility that passengers may simply perceive specific airports in England to provide a superior service, be in a more convenient location, etc. compared to others."*[25]

The suggestion here is that part of the perceived premium at Heathrow may be driven not by the fact that Heathrow is congested but rather by the fact that it

[24] This is a commonly used statistical measure of competition in the market, which takes the sum of the squared market shares of each of the individual participants. As a smaller number of firms account for more of the market, the value of the index increases.

[25] http://publicapps.caa.co.uk/docs/33/CAP1722b%20FTI%20Consulting%20-%20Report%20to%20CAA%20on%20Scarcity%20Rents%20-%20Final%20-%2030%20October%202018.pdf

provides a superior service and is more convenient to passengers than the other airports in the sample.

We discuss at length in Section 1 the fact that differentiation of Heathrow as an airport is a *necessary* but not *sufficient* condition for a congestion premium. It is nevertheless worth reiterating here that it is the capacity constraint that prevents competition from eroding a premium for service, location and other aspects of differentiation. Our core model therefore retains the “residual approach” noted here by FTI, as it is the best empirical measure of the level of the congestion premium that Heathrow’s level of excess demand and differentiation have led to.

However, in addition to this, and to demonstrate that our results are not sensitive to the use of the residual approach, we have estimated the premium in two additional ways:

- Including as a regressor our top-down measure of unserved demand (as calculated in Section 1); and
- An ‘out of sample’ approach.

One other possible approach to estimating the congestion premium is using the variation in runway utilisation across the sample, as it is directly observable and measures the level of capacity used at each airport in each year. In our view, runway utilisation cannot capture the nature of the congestion premium that we are seeking to measure. We have nonetheless tested this approach, and discuss the results and shortcomings in further detail in Annex A.

4.3.1 Unserved demand as an independent variable

The congestion premium is really driven by the volume of unserved demand, and this is not directly observable at an airport. In the section on our top down analysis we estimated the level of unserved demand at Heathrow (see Figure 23).

The way that we estimated this variable is such that all of the other airports in the sample have unserved demand of 0%. This is because Heathrow was assumed to be constrained in the year 2000, and the unserved demand was assumed to have been growing since then. However, none of the other airports in our sample are as constrained, even now, as Heathrow was back then. Therefore, following the same approach implies no unserved demand for the other airports. As mentioned, we believe there may be a premium at other airports at least during some parts of the day. Therefore, the fact that we are comparing prices at Heathrow to the average prices at the other airports (which may also have a premium in reality) means that our estimate at Heathrow is therefore likely to be an underestimate.

We can include this variable as an additional explicit measure of congestion at Heathrow rather than needing to rely on a residual / dummy approach.

Similar to the above, under this approach we can estimate a general relationship between unserved demand and prices. To turn this into an estimate of the congestion premium at Heathrow, we estimate what prices would be expected to be if its unserved demand becomes equal to 0%. The difference in prices is the congestion premium.

4.3.2 Out-of-sample estimation

The level of congestion at Heathrow is clearly unique. Therefore, one potential critique of our models is that they may be biased / skewed precisely because they are based in part on Heathrow data. To test that this proposition, we have separately re-run many of our models while excluding Heathrow from the data. At one level this is helpful to check that our coefficients are consistent whether we include the Heathrow data or not. It also provides us with an opportunity to produce an ‘out of sample’ estimate of prices at Heathrow, which we can then compare to the actual prices.

In particular, we can take the coefficients estimated in the models excluding Heathrow, and then estimate what prices might be expected to be at Heathrow, based on these coefficients. The difference between actual fares at Heathrow and the out-of-sample estimate would, on average, indicate the size of the congestion premium. We therefore use this as a further, independent approach to estimating the premium.

We have illustrated below the justification for testing an out-of-sample approach by positing a situation in which congestion at Heathrow results in not just a different level of price but in a different relationship between price and independent variables.

4.4 Model specification

In this section we describe our model specification in more detail. A complete discussion of each of the methodological decisions taken is available in Annex A. However there are a number of common characteristics of our modelling that should be noted:

- We have modelled **short haul and long haul routes separately**. This is to reflect the different underlying economics of operation and the potential for different trends in the congestion premium. Modelling the two together presents similar aggregate results.
- We have **pooled together our data for 2011 to 2018** in our models, rather than estimating each year separately. To account for common yearly trends, we have included a dummy for each year in the sample. To allow the congestion premium to vary over time, we have included interactions between those yearly dummies and the Heathrow dummy. We have done this because, from estimating yearly models separately, it is clear that the coefficients on control variables (such as those on distance or frequency) are stable over time, and that pooling the observations therefore results in more precise estimates of the variables of interest.
- We have chosen to model **fares in log form**, to allow for a clear interpretation of coefficients as proportional effects, and to better approximate the relationships between key variables of interest and fares. Modelling fares in non-log form produces similar results.
- For the reasons described in section 2.2.5 we have focussed our analysis on **point-to-point passengers only**. We nonetheless control for the proportion of passengers on a route who are transfer passengers in many of our models.

- Each model has been run twice to ensure that **outliers are excluded** from the final estimation, defined as observations with residuals more than 3 standard deviations from zero. This means that the number of observations will vary slightly between models, as the proportion of observations excluded ranges between 0.3% and 1.0% in our core specifications. The justification for and implications of this decision are discussed in detail in Annex A.
- Wherever it is displayed, the congestion premium is expressed as a **percentage increase on counterfactual fares** rather than as a percentage of actual fares. For example, a 25% congestion premium is equivalent to saying that 20% of an actual fare is due to the congestion premium.

To ensure that our estimates of the congestion premium are robust, and not a quirk of the models selected, we have presented a range of different specifications controlling for a number of potential drivers of fares. In all of our models, we include dummies for each year (to control for any unobserved time trends) and a Heathrow dummy interacted with the yearly dummies, to measure the annual premium.

In the illustration below, we describe how we have progressively added more control variables to the models. As the models move from narrow (few control variables added) to wide (many variables added), the models become more fitted. The purpose of this approach is to show that the congestion premium estimates are robust to the addition of extra control variables – i.e. the estimate of the premium neither disappears nor changes significantly when certain variables are added.

Figure 42 Structure and presentation of core models

Models are more determined/fitted

More controls are added

Variables	Narrow	1	2	3	4	Wide
Distance (000s KM)	✗	✗	✗	✗	✗	✗
Heathrow dummies	✗	✗	✗	✗	✗	✗
Year dummies	✗	✗	✗	✗	✗	✗
Frequency (own)		✗	✗	✗	✗	✗
Frequency (other)		✗	✗	✗	✗	✗
Transfer share		✗	✗	✗	✗	✗
LCC share			✗	✗	✗	✗
Business class share			✗	✗	✗	✗
European jet fuel price				✗	✗	✗
Avg. seats					✗	✗
Airport competition					✗	✗
Route competition					✗	✗
Skytrax rank						✗

Source: Frontier Economics

TECHNICAL SPECIFICATION

Our wide core model is an estimate of the following equation:

$$\ln(\text{AvgFare}_{r,y}) = \hat{\boldsymbol{\lambda}}_y + \widehat{\mathbf{CP}}_{r,y} + \hat{\gamma}_1 \cdot \text{Dist}_r + \cdots + \hat{\gamma}_{12} \cdot \text{SkytraxRank}_{r,y} + \hat{u}_{r,y}$$

Where r and y denote the route and year of an observation. $\hat{\lambda}_y$ is a set of yearly constant terms and $\hat{u}_{r,y}$ is a route and year specific residual. The parameter of interest – the congestion premium – is the coefficient on a dummy variable equal to 1 if the route originates from Heathrow in a given year, and is shown in the equation as $\widehat{\mathbf{CP}}_{r,y}$. Each of the controls is shown with an estimated coefficient $\hat{\gamma}_n$. The narrower models are identical with the appropriate control terms removed.

Each model is estimated using OLS with robust standard errors.

Further technical details are provided in Annex A

4.5 Results

In this section we report and comment on the main results. This covers:

- The 'core' models – where we use a Heathrow dummy variable; and
- The alternative models which do not use a Heathrow dummy variable:
  - Unserved demand; and
  - Out of sample approach

The full coefficient tables for each model are available in Annex A.

4.5.1 Core models

Short haul

The table below shows the coefficients on the Heathrow year dummy variables estimated in our core models on short haul flights in our European sample. Each model estimates log fares for all short haul routes included in our sample, controlling for each of the independent variables as laid out in Figure 42. (For example, core model 1 includes all of the control variables listed in the table above).

The summary statistics refer to the number of observations used to estimate the model, the percentage premium over the entire 2011 to 2018 period and the implied average premium per passenger for a return trip.

The coefficients associated with the Heathrow dummies provide the estimates of the congestion premium in each year. A positive coefficient indicates that fares from Heathrow are higher than fares at other airports in the sample. The stars indicate that the associated increase in fares is statistically significant. In the table below , we see that at least one of the congestion premium estimates in each year in each model is significant at the 1% level, and that this is true for the majority of estimates.

Figure 43 Core models – Short haul

		Narrow	1	2	3	4	Wide
Heathrow dummies	2011	0.19***	0.21***	0.17***	0.17***	0.19***	0.15***
	2012	0.28***	0.28***	0.24***	0.24***	0.27***	0.22***
	2013	0.13***	0.14***	0.10***	0.10***	0.11***	0.07**
	2014	0.14***	0.14***	0.10***	0.10***	0.11***	0.09***
	2015	0.29***	0.29***	0.23***	0.23***	0.24***	0.23***
	2016	0.19***	0.19***	0.12***	0.12***	0.15***	0.14***
	2017	0.12***	0.12***	0.05*	0.05*	0.07**	0.07**
	2018	0.19***	0.21***	0.13***	0.13***	0.16***	0.14***
Observations		10,427	10,396	10,395	10,395	10,381	10,387
Average premium, 2011-2018 (%)		20.8%	21.7%	15.0%	15.0%	17.2%	14.4%
Est. premium per return flight 2018 (£)		£39	£43	£28	£28	£33	£29

Source: Frontier Economics analysis of IATA data
All models estimate average log fares for routes from our European hub sample between 2011 and 2018 using OLS
** denotes statistical significance at the 10% level; ** denotes statistical significance at the 5% level; *** denotes statistical significance at the 1% level*

The estimated premium across the entire period varies from 14.3% to 21.8%. The largest shift in our results occurs between models 1 and 2, when we add business class and low cost carrier shares at the route level, reducing our estimated congestion premium from around 20% to around 15%. This suggests that part of the congestion premium is due to the mix of carriers on a route, not just higher prices from comparable carriers. Arguably, the lack of low cost carriers at Heathrow is due to capacity constraints since their branding prevent higher prices at individual airports. However, to be conservative we retain it as a legitimate driver of fares and as a control in the majority of our models.

Within years, each of the coefficient estimates are roughly consistent across models – the patterns between years appear to be an accurate reflection of the data and not random noise produced by the modelling process.

The chart below plots the coefficients for each year. This chart shows that the range of estimates is uneven with rises and falls between years.

Figure 44 Estimated short haul congestion premium, 2011-2018

Source: Frontier Economics

Note: The range represents the range of central model estimates of the congestion premium, and not the uncertainty associated with those estimates within a given model.

The shape of the results is driven by either changes in prices at Heathrow or changes in prices at the other airports in the sample – or both, i.e. the premium is really measured as the difference in average prices, after controlling for the other variables.

As described in the section on the top down analysis, in principle, the premium will decrease in a given year if Heathrow experienced a very large growth in that year compared to the unconstrained growth rate. For example, in 2013 (when the premium decreased) we note that Heathrow actually experienced a 3.4% increase in passenger volumes, which exceeded our estimate of unconstrained demand growth in that year. Also, we note that the premium compares fares at Heathrow with those at the other airports in the sample. Therefore, it follows that if the other airports are becoming increasingly constrained, the coefficient on Heathrow would actually be an underestimate because it would represent a premium on top of airports which are themselves experiencing a congestion premium.

Below, we repeat the analysis for long haul. The interpretation of coefficients and summary statistics is identical to that for short haul.

Long haul

The results below show that the premium for long haul over the whole period was around 25%. This is somewhat higher than the figure for short haul, although similar in the years with a higher short haul premium. However, we note that because long haul fares are higher, the premium in absolute terms is much larger. The results suggest that whereas the combined short haul premium in 2018 was around £300 million to £500 million, the estimates for long haul are around £1.8 billion to £2.2 billion, or around £200 per passenger for a return trip.

Figure 45 Core models – Long haul

		Narrow	1	2	3	4	Wide
Heathrow dummies	2011	0.19***	0.15***	0.14***	0.14***	0.15***	0.14***
	2012	0.23***	0.19***	0.20***	0.20***	0.21***	0.20***
	2013	0.21***	0.18***	0.18***	0.18***	0.18***	0.17***
	2014	0.26***	0.25***	0.24***	0.24***	0.25***	0.23***
	2015	0.28***	0.28***	0.27***	0.27***	0.25***	0.24***
	2016	0.30***	0.27***	0.25***	0.25***	0.23***	0.22***
	2017	0.31***	0.28***	0.28***	0.28***	0.25***	0.25***
	2018	0.26***	0.27***	0.25***	0.25***	0.22***	0.22***
Observations		4,454	4,444	4,444	4,444	4,436	4,437
Average premium, 2011-2018 (%)		28.8%	26.3%	25.3%	25.3%	24.1%	22.9%
Est. premium per return flight 2018 (£)		£205	£247	£231	£231	£201	£195

Source: Frontier Economics analysis of IATA data
All models estimate average log fares for routes from our European hub sample between 2011 and 2018 using OLS
** denotes statistical significance at the 10% level; ** denotes statistical significance at the 5% level; *** denotes statistical significance at the 1% level*

Each Heathrow dummy coefficient is statistically significant (at above the 1% level). More so than in the short haul models, the yearly trend in the long haul premium is smooth.

Figure 46 Estimated long haul congestion premium, 2011-2018

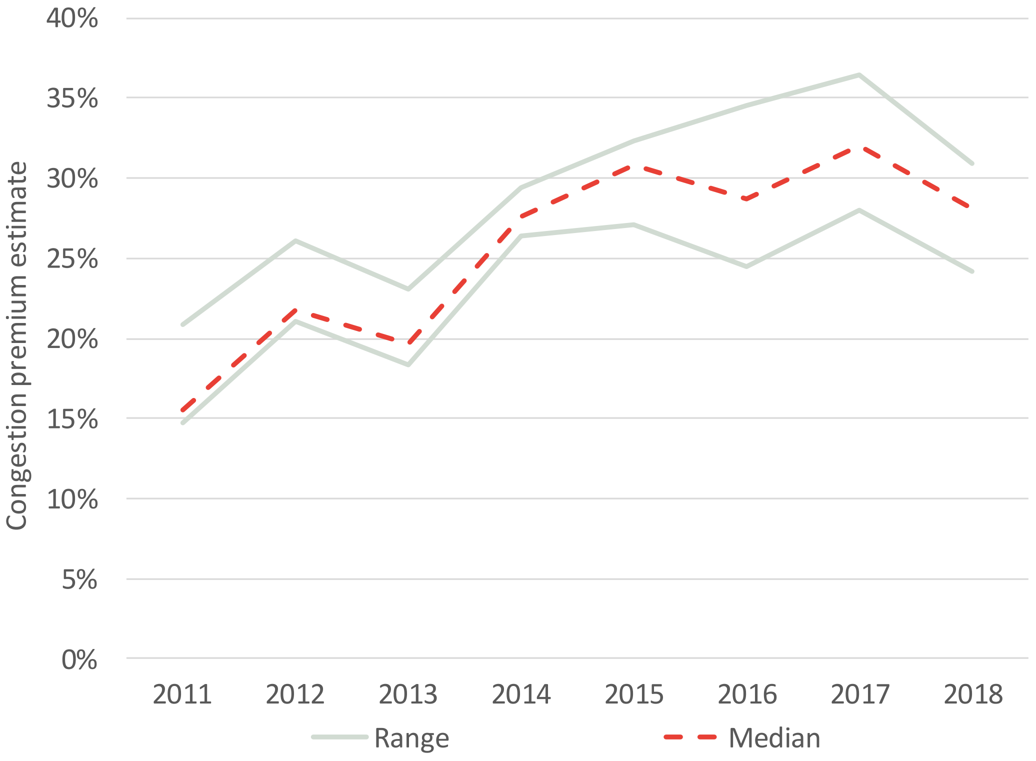

Source: Frontier Economics

Note: The range represents the range of central model estimates of the congestion premium, and not the uncertainty associated with those estimates within a given model.

As mentioned, the trend of increased congestion elsewhere in Europe may lead to an underestimate of the congestion premium, as in effect it will be a premium on top of airports which are themselves becoming increasingly constrained.

Our results show that the premium is greater in absolute terms for long haul than for short haul. The share of long haul departures at Heathrow has been increasing over time, which is consistent with airlines wishing to enter routes with the largest premia.

Figure 47 The share of long haul departures at Heathrow has been increasing

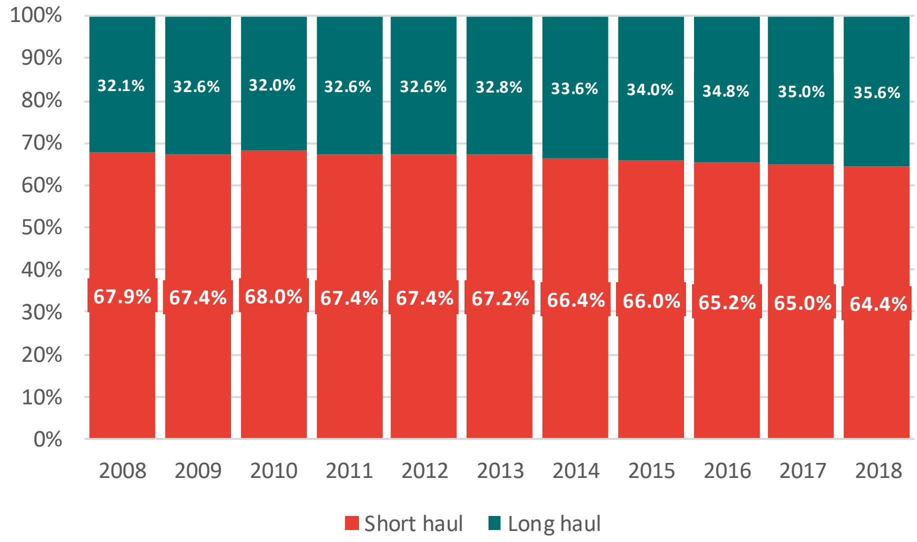

Source: Frontier analysis of OAG data

However, we note that Heathrow is hub, and network carriers need to be mindful to strike a balance between short haul and long haul flights – e.g. long haul flights carry many transfer passengers who fly in to / out of Heathrow on short haul flights. This limits the extent to which airlines can replace short haul flights with long haul flights.

4.5.2 Unserved demand approach

Using the estimates of unserved demand from the top down analysis we can estimate the congestion premium directly as a function of unserved demand. Under this interpretation, a percentage point increase in unserved demand results in a percentage increase in fares equal to the estimated coefficient. In each of these models, the following additional control variables are included: frequency, transfer share, LCC share, ticket class mix and yearly dummy variables. These controls are identical to those used in one of our core models, specifically core model 2.

Our results are shown in the table below.

Figure 48 Unserved demand models

	Short haul	Long haul
Unserved demand (%)	0.91***	1.60***
Observations	10,390	4,445
Average premium 2011-2018 (%)	12.0%	22.0%
Average premium per return flight 2011 (%)	6.0%	11.0%
Average premium per return flight 2018 (%)	19.0%	37.0%

Source: Frontier analysis
All models estimate average log fares for routes from our European hub sample between 2011 and

2018 using OLS
** denotes statistical significance at the 10% level; ** denotes statistical significance at the 5% level; *** denotes statistical significance at the 1% level*

In both models including unserved demand, the coefficient is significant at the 1% level. For short haul routes, the coefficient is estimated at 0.92% while for long haul it is almost twice as high at around 1.6%. These estimates imply that the congestion premium impacts more heavily on long haul routes, both in absolute and percentage terms.

For long haul, this means that as demand increases by a percentage point, with supply constrained, fares will rise by 1.6%. Interestingly, this can also be interpreted as the inverse of an elasticity between fares and demand. This implies a price-elasticity of demand of around -0.62, which is broadly in line with other estimates of the price-elasticity of demand for airfares, e.g. please see the DfT's estimates in Figure 24 in the section on the top down analysis. This approach therefore supports the rationale underlying the top down analysis.

The unserved demand models imply a congestion premium of roughly the same level as our core models. However, since the premium in these models is a function of the unserved demand estimates, it increases over time directly in proportion to the trends in unserved demand. Because of this, the estimates of the 2018 congestion premium are much higher than our core models.

4.5.3 Out-of-sample estimation

As described above, we have also estimated the congestion premium at Heathrow using an 'out of sample' approach. The coefficients below compare our estimates from one of our core models including Heathrow to a version of the model without Heathrow. This approach allows us to test that our models are not overly sensitive to the inclusion of Heathrow in the sample, and it also allows us to produce a separate estimate of congestion premium.

Figure 49 Out-of-sample estimation models

	Short haul		Long haul	
	Core model 2	*Exc.Heathrow*	*Core model 2*	*Exc. Heathrow*
Distance (000s KM)	0.00025***	0.00025***	0.00005***	0.00006***
Frequency (own)	-0.007**	-0.013***	0.067***	0.071***
Frequency (other)	0.018***	0.020***	0.054***	0.053***
Transfer share	0.45***	0.45***	0.45***	0.44***
LCC share	-0.28***	-0.30***	-0.37***	-0.37***
Business class share	0.00	-0.04	0.54***	0.46***
Observations	10,395	9,659	4,444	3,814
Average premium 2011-2018 (%)	15.0%	22.9%	25.3%	26.6%
Average premium 2018 (%)	14.0%	20.4%	28.1%	31.9%
Average premium per return flight (£)	£28	£40	£231	£263

Source: Frontier Economics
All models estimate average log fares for routes from our European hub sample between 2011 and 2018 using OLS
** denotes statistical significance at the 10% level; ** denotes statistical significance at the 5% level; *** denotes statistical significance at the 1% level*

A full discussion of each of the coefficients in the core models is available in the technical appendix, however the results show that all of our models are broadly consistent with and without the inclusion of Heathrow in our sample.

Comparing estimated prices at Heathrow with actual prices at Heathrow also shows a significant difference. What this says is that if the drivers of prices elsewhere in the sample applied at Heathrow, these are the prices that Heathrow passengers would be paying. The out of sample estimation method indicates a premium very similar to that indicated by the core models, although the yearly pattern is slightly different, as shown in the table below.

Figure 50 Out-of-sample congestion premium estimates by year

Estimated premium	Short haul	Long haul
2011	25.6%	16.4%
2012	34.9%	20.0%
2013	16.9%	19.2%
2014	18.0%	27.1%
2015	34.8%	35.6%
2016	23.7%	32.4%
2017	12.4%	31.9%
2018	20.4%	31.9%

Source: Frontier analysis of IATA data

4.6 Conclusions

We have estimated the congestion premium using a wide variety of econometric approaches. We have produced many models and all of them point to there being a significant premium at Heathrow:

Figure 51 Summary of congestion premium modelling results

	Short haul	Long haul
Core models	14-22%	22-29%
Unserved demand models	12%	22%
Out-of-sample estimation	23%	27%
	£433m/year	£1.93bn/year
	£34/return flight	£217/return flight

Source: Frontier Economics

Notes: Per-year and per-passenger numbers are calculated using 2018 fare and passenger statistics, and assume a 17% premium for short haul and a 25% premium for long haul

In almost all of our models the short haul congestion premium is estimated around 17% of fares, while for long haul routes it is close to 25% of fares. These estimates mean a total congestion premium equal to over £2 billion a year in 2018, and over £30 per passenger making a return trip to short haul destinations and £200 per passenger making a return trip to long haul destinations. These estimates are more detailed, more robust and more consistent than our previous estimates that focused only on a smaller sample of airports and a shorter period of time. More detail on the robustness checks, diagnostic tests, model selection and implications are drawn out in the detailed technical appendix at the end of this report.

Airport charges

We note that we have not included airport charges as a control variable in our econometric models. Airport charges are higher at Heathrow than at the other airports in our sample, on average. If prices in a competitive market reflect LRMC, then differences in airport charges could be seen as a legitimate reason for differences in prices at different airports.

The charts below show the departing passenger charges (DPCs) at each of the airports in our sample. (This does not include other non-passenger charges such as landing charges. However, DPCs tend to represent the majority of airport charges).

Figure 52 Airport charges at the sample of airports

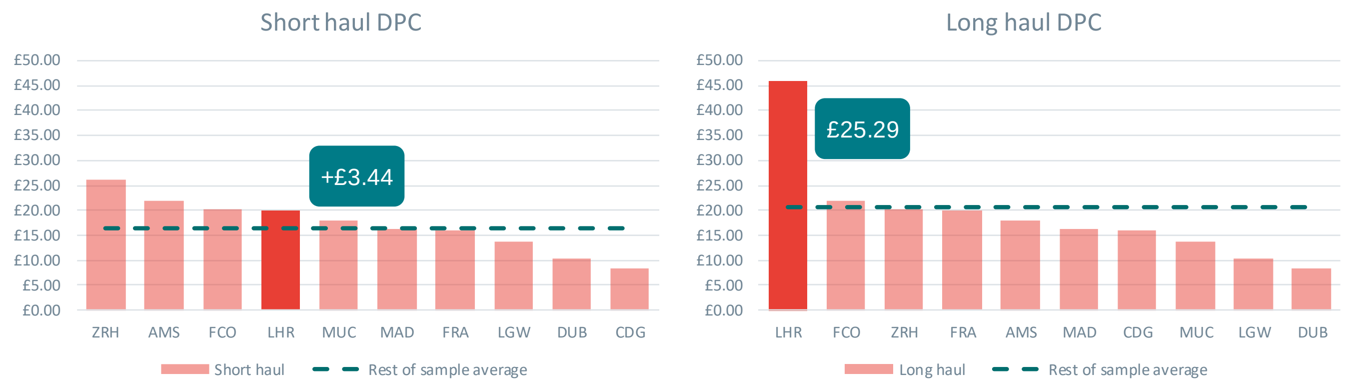

Source: Frontier analysis based on the most recent published airport charges at each airport

Note: For comparability, the short haul DPCs are based on assuming a destination in the Schengen area, and the long haul DPCs are based on assuming a destination outside of Europe / EU – e.g. the US.

The analysis shows that:

- For short haul, DPCs are about £3 greater at Heathrow than the weighted sample average; and
- For long haul, DPCs are about £25 greater at Heathrow than the weighted sample average.

To control for this difference, one approach could be to subtract this difference from our estimates of the congestion premium. This would produce the results below:

Figure 53. Per-passenger premium estimates for short haul, 2018

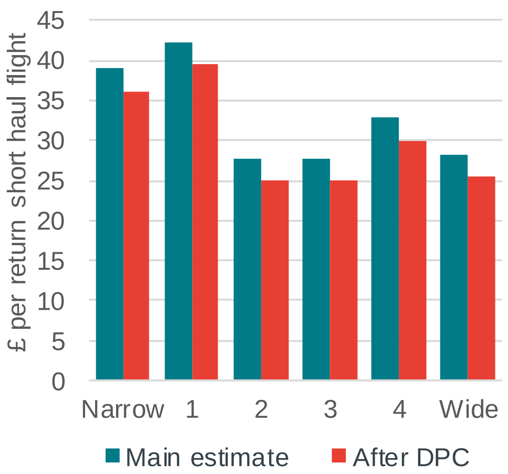

Source: Frontier Economics

Note: After DPC refers to the per-return-flight congestion premium estimate after the Heathrow DPC has been set to the weighted sample average

Figure 54. Per-passenger premium estimates for long haul, 2018

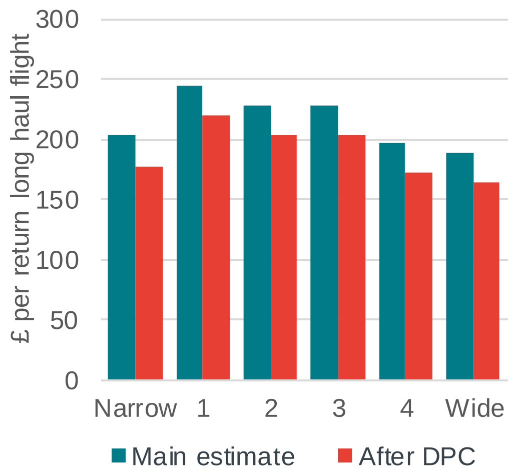

Source: Frontier Economics

Note: After DPC refers to the per-return-flight congestion premium estimate after the Heathrow DPC has been set to the weighted sample average

However, it is important to note that it would be misleading to suggest that if airport charges were lower at Heathrow then ticket fares would also fall by an equivalent amount. This is because, if ticket prices were to fall by an equivalent amount, then demand would increase. But we have already established that because Heathrow is constrained, there is no capacity to accommodate this extra demand. And prices would need to rise back up again to the original price, which is the price which ensures that the excess demand is priced out of the market. This dynamic is illustrated below.

Figure 55 If airport charges were to fall, prices would not necessarily fall by an equivalent amount

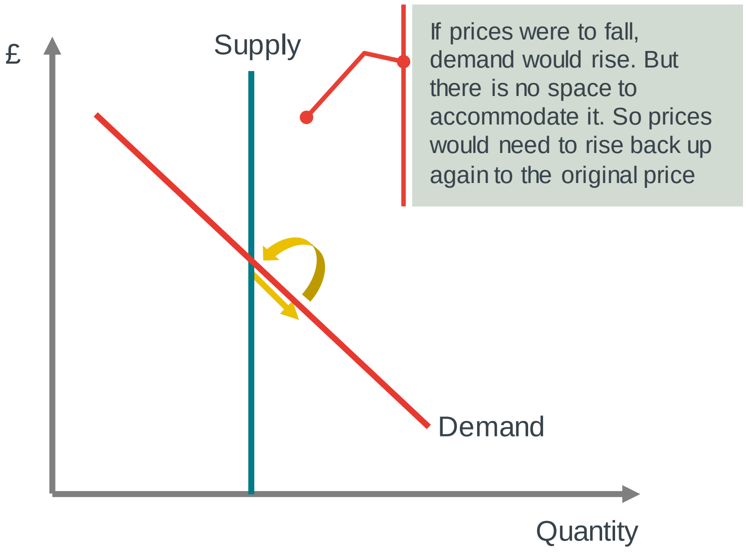

This pattern could vary from route to route depending on whether the individual route is constrained or on whether airlines could increase load factors. However, the general point remains. It is for this reason that we believe it is not appropriate to subtract the difference in airport charges from the premium. In any case, even if we did, this does not explain away the congestion premium and it remains very significant.

5 CONCLUSION

Heathrow has been constrained for many years. But in this time, demand has continued to grow. Where demand exceeds supply, prices must rise to price the excess demand out of the market. We have estimated the congestion premium at Heathrow using three separate approaches:

- **Top down analysis:** We have estimated that if Heathrow had been unconstrained, then in 2018 it would have handled around 97 million passengers. But because it was constrained, it was only able to handle 80 million. It follows that to have priced 17 million passengers out of the market (or 17.5% of total unconstrained demand), prices must have been higher than under an unconstrained scenario. Using a price elasticity of demand of -0.7 (the most conservative figure used in the DfT's demand modelling), this implies a congestion premium of 25% in 2018. This is equal to **£2.6 billion** across all short haul and long haul passengers.

  The chart below shows how the premium has evolved over time:

Figure 56 The congestion premium over time estimated on a top down basis

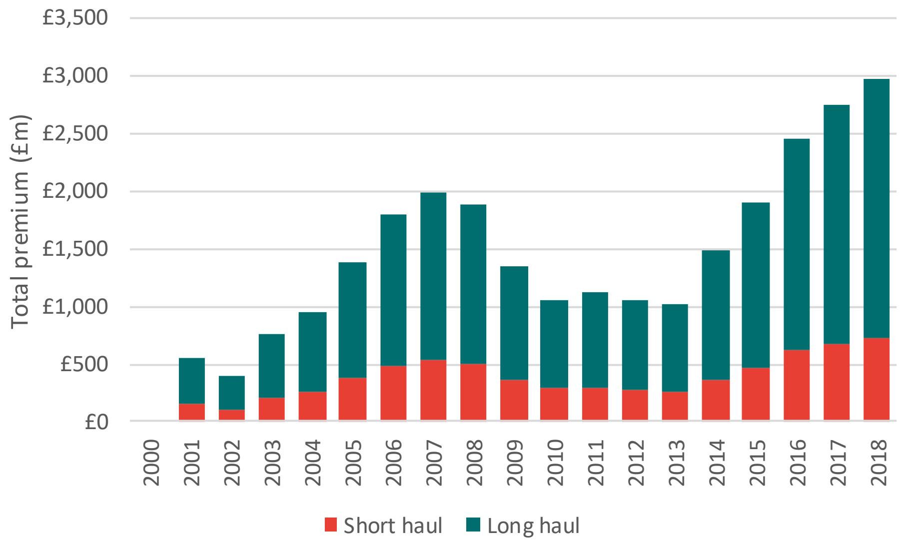

Source: Frontier analysis

- **Analysis of slot transfer prices:** In a competitive market with spare capacity, we would expect average prices on a route to reflect the long run marginal cost of operating the route. Why then do airlines pay millions to acquire slots at Heathrow, especially when slots at many other airports are free? This must suggest that at Heathrow incumbent airlines are able to earn a premium above this long run marginal cost. If there were no premium, then buying a slot for large sums of money would simply be loss-making.

  By analysing recent slot transfers at Heathrow we have conservatively estimated that the average slot at Heathrow is worth around £800,000. Across all slots at Heathrow, this implies a total value of all slots at Heathrow at around £15 billion in total.

Airlines buy slots so that they can then earn the congestion premium in future. Assuming an investment time horizon of 10 years, and a cost of capital of 7.5%, this implies that the value of slots of an annualised basis is £2 billion per annum. The premium must therefore be at least **£2 billion per annum**, otherwise buying slots would be loss-making. The table below shows how the results change if we vary these input assumptions:

Figure 57 Slot transfer prices: Implied total premium per year (£m)

		Price per slot									
		Average observed					Minimum observed				
	WACC	*6.0%*	*7.0%*	*7.5%*	*8.0%*	*9.0%*	*6.0%*	*7.0%*	*7.5%*	*8.0%*	*9.0%*
Asset lifetime	5 year	£3,357	£3,416	£3,446	£3,476	£3,535	£1,504	£1,530	£1,544	£1,557	£1,584
	10 year	£1,921	£1,994	£2,031	£2,068	£2,143	£861	£893	£910	£926	£960
	15 year	£1,456	£1,538	£1,580	£1,621	£1,706	£652	£689	£708	£726	£764

Source: Heathrow slot price data, ACL completed slot trades database, Frontier calculations.

- **Econometric analysis**: We have also carried out an extensive piece of econometric analysis to estimate the premium. Comparing ticket prices at Heathrow with those at other large airports in Europe shows that prices at Heathrow are generally higher. Using econometric analysis, we can isolate the proportion of this difference which is attributable to congestion.

Figure 58 Summary of congestion premium modelling results

	Short haul	Long haul
Core models	14-22%	22-29%
Unserved demand models	12%	22%
Out-of-sample estimation	23%	27%
	£433m/year	£1.93bn/year
	£34/return flight	£217/return flight

Source: Frontier Economics

Notes: Per-year and per-passenger numbers are calculated using 2018 fare and passenger statistics, and assume a 17% premium for short haul and a 25% premium for long haul

We note that econometrics is not a perfect science because there will never be one single model which can perfectly explain ticket prices. However, the purpose of this exercise is to demonstrate that we have estimated the premium using many models which vary in terms of the number of variables included and also in terms of how we capture congestion, and all of them point to there being a considerable premium at Heathrow. We have produced a range of estimates which consistently point to a congestion premium of around **£2.4 billion** in 2018, equivalent to a premium of around 25%.

There are caveats and limitations with each of these approaches when taken in isolation. For example, it is always possible to challenge the robustness of econometric analysis and technical issues associated with benchmarking prices at different airports. But this is not an issue at all with the other two approaches. We note that the size of the premium is strikingly consistent across all three, and we believe that this demonstrates that the premium is not a quirk of econometric analysis, but a clear pattern which can be observed across a variety of approaches.

Figure 59 All three approaches point to a consistent premium per annum

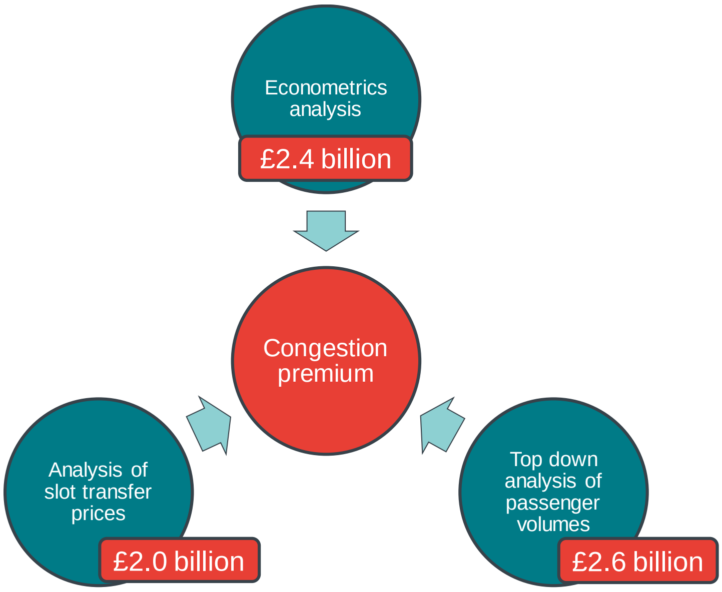

Source: Frontier analysis

We believe that there is clear evidence to suggest that the premium in 2018 was in the order of £2 billion. This equates to a premium of 25% or an average of £34 per passenger for a short haul return trip and over £200 per passenger making a long haul return trip, flying point-to-point.

ANNEX A TECHNICAL ANNEX ON ECONOMETRICS ANALYSIS

In this Annex, we outline in detail the econometric approaches we have taken to estimating the congestion premium, including more detailed tables of coefficients, diagnostic testing and robustness checks.

The Annex is divided into the following sections:

- **Data** – a more detailed summary of the data we have used in our analysis, including summary statistics and a description of the outliers excluded
- **Modelling fares** – an explanation of the key decisions made in estimating fares, including the reasoning of the relationship with distance, pooling our yearly observations and an interpretation of the other coefficients in our models
- **Modelling congestion** – a discussion of, and further tests of, our core approach to modelling congestion, along with details of alternative approaches
- **Diagnostic tests** – how we can know that our models are performing well, and accurately measuring the scale of the congestion premium
- **Coefficient tables** – the full results of our core models for short and long haul, and our alternative congestion models
- **Comparison to previous estimates** – a discussion of how our econometric approaches and results have changed since we estimated the congestion premium in past reports

A.1 Data

Fare 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).

26 https://www.iata.org/services/statistics/intelligence/Pages/market-intelligence.aspx

The fares listed exclude taxes but include airport charges. The consequences of including airport charges are discussed in detail in Section 4. The exclusion of taxes mean that our estimate of the difference in prices paid by passengers is an underestimate, as air passenger duties (APD) are higher in the UK than in the rest of Europe. It does not, however, affect our estimate of the congestion premium, as we are estimating the difference in prices relative to a world without capacity constraints, not relative to a world of equal taxes.

The data provided is for one-way fares, and we model our estimates on that basis (through the departures from sample airports), but calculate indicative return fares using our congestion premium estimates and one-way fares.

A.2 Modelling fares

In this section, we detail the various methodological decisions we have made in modelling fares, the background controls against which we model the congestion premium.

Fares and distance

To ensure that our estimates of the congestion premium are correct, it is important to accurately model the variation in fares across our reference sample. As Figure 60 shows below, the distribution of fares in 2018 ranges from near zero to above £1000, with a long tail and a clear peak around £100. Visually, the distribution of route-average fares is much closer to exponential than normal, which informs our choice to model fares in log form.

Figure 60 Distribution of route-level fares in European sample, 2018

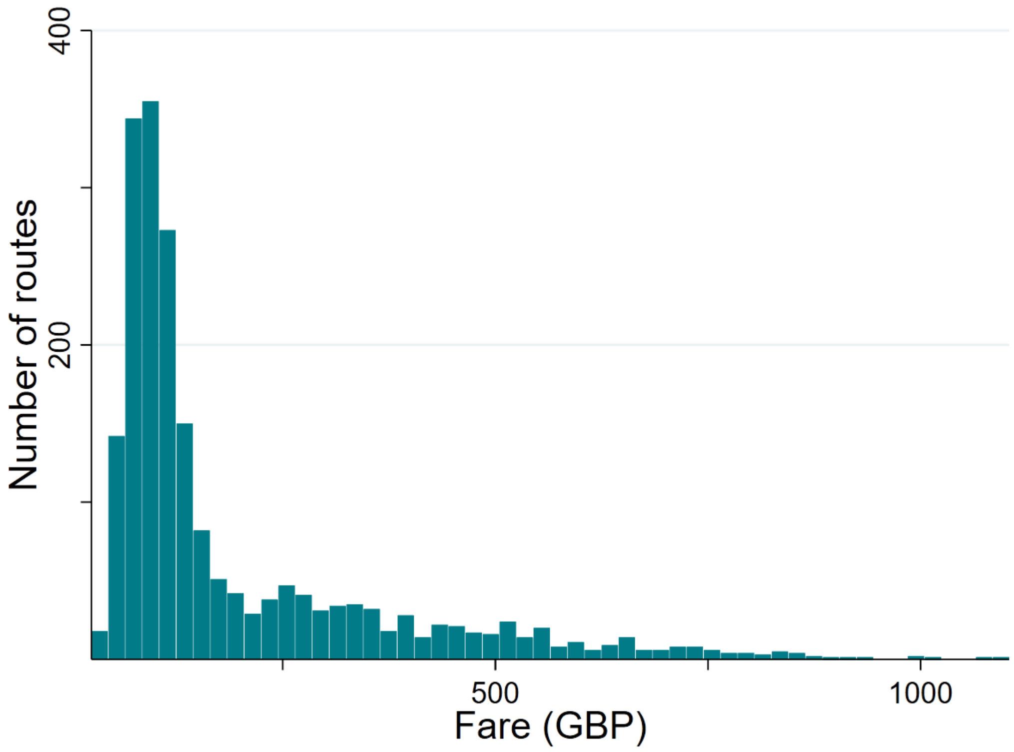

Source: Frontier analysis of IATA data

Taking the log of fares produces a more tractable set of data, as shown in the histogram in route Figure 61. However log fares do not remain smoothly distributed, with two peaks in the distribution of fares. Superimposing the average distance associated with fares reveals the effect of merging of two separate distributions – short haul and long haul flights. This is a distinction used widely in aviation, and the industry traditionally uses a cut-off of 2,000 nautical miles (roughly 3,700km) for short haul. Short haul and long haul flights reflect different economic models of operation, and therefore reflect different fares and different relationships with control variables. Imposing this distinction in our data neatly divides the two separate distributions of fares for different types of flights.

Figure 61 Route-level fares and distance in European sample, 2018

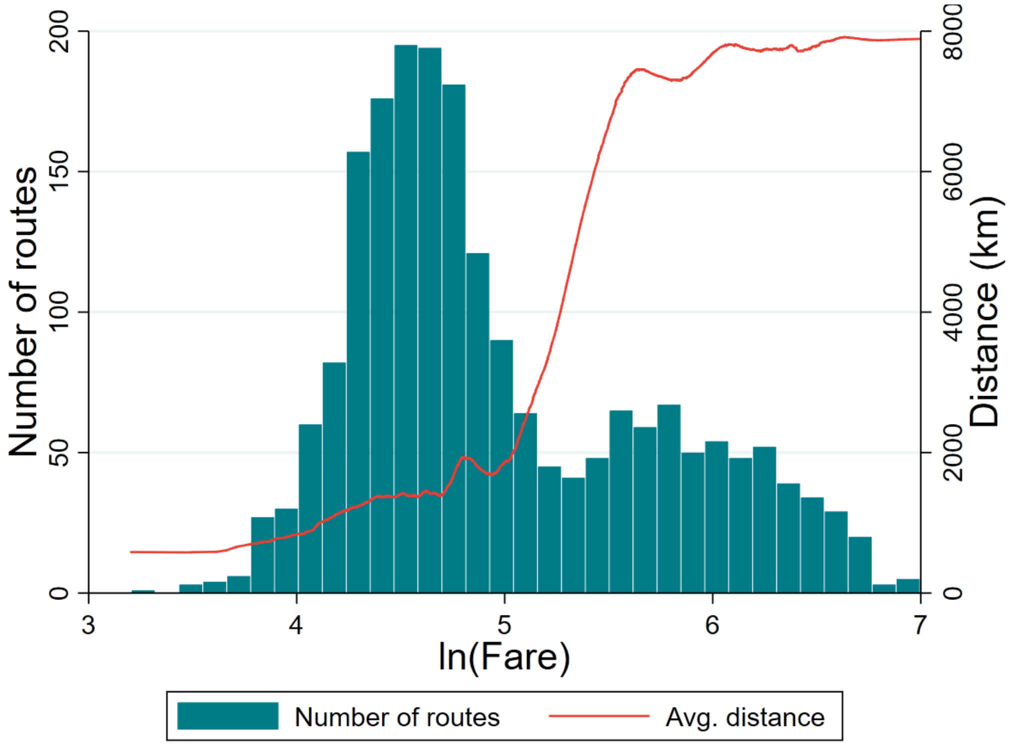

Source: Frontier analysis of IATA data

However, the separation of long and short haul flights is not crucial to our results – modelling routes in a combined form leads to a very similar estimate of the overall congestion premium. But it does produce less easily interpretable estimates of other coefficients in our models, and it covers up important variation in the different estimates of short and long haul congestion premia. As shown in the main report, the short haul congestion premium estimate seems to vary from year to year, while the long haul premium exhibits a reasonably stable trend. When combining short and long haul estimates, the overall premium is a weighted average of the two models, masking the difference in the volatility of the two segments.

Incorporating additional years

An important difference between this report and our previous work in this area is that we have used multiple years of data. Previously, we have used identical sources of data for one year only. Multiple years of data improves our analysis in two ways:

- **Greater power**: increasing the number of observations, including variation *across years* rather than just *between airports* improves the power of our models in estimating the parameters we are interested in; and
- **Variation over time**: rather than assuming that parameters remain fixed over time, we can allow them vary by year, giving a fuller picture of the way the relationship between fares and our control variables shifts from period to period.

To make the greatest use of this improvement in our data, we have fixed some relationships over time, but not others. At one extreme, we could run each model on each year of data we have available, resulting in somewhat different coefficients in each year. At another extreme, we could pool together our models and produce average estimates across the whole time period. Since we are interested primarily in estimating the congestion premium, we have pooled together our data but specifically allowed the congestion premium variable to vary using yearly interacted dummy variables. To control for other common trends that would bias our estimate of the congestion premium, we have included yearly dummies intended to capture all other sample-wide drivers of price over time (such as inflation). Taken together, this dramatically improves the power of our estimates of the congestion premium while effectively managing the existence of variation over time.

However, to test that we have appropriately managed the time-series nature of our modelling, we have also run our models separately over individual years. While the congestion premium and constant term vary as expected (that is, they are not materially different from our core models), the estimated coefficients on our additional control variables are reliably stable from year to year. The table below provides an example of this phenomenon, showing the estimated coefficients from one of our long haul models (specifically model 2) estimated separately across three years within our sample.

Figure 62 Sample of long haul coefficients estimated separately by year

	2013	2015	2017
Distance (000s KM)	0.00006***	0.00005***	0.00005***
Frequency (own)	0.07114***	0.08125***	0.06027***
Frequency (other)	0.04168***	0.05853***	0.07080***
Transfer share	0.41388***	0.44165***	0.52565***
LCC share	-0.65641***	-0.43326***	-0.35743***
Business class share	0.58472***	0.48680***	0.52920***
LHR dummy	0.16736***	0.26059***	0.29749***

Source: Frontier analysis of IATA data
All models estimate average log fares for routes from our European hub sample for the listed year using OLS
** denotes statistical significance at the 10% level; ** denotes statistical significance at the 5% level; *** denotes statistical significance at the 1% level*

The stability of these coefficients does not mean the fares themselves are stable, as our yearly dummies vary substantially, as does the congestion premium, however the proportional drivers of prices are stable. For example, an additional kilometre in distance will affect the average fare paid by a consistent percentage amount, regardless of in which year it has occurred.

Managing outliers

There are a number of routes in our data that are persistent outliers in ways for which our models cannot reasonably be adjusted to control without overfitting. This discrepancy occurs in both directions, and seems to reflect idiosyncrasies and non-market factors on particular routes. For example, the short haul route between Dublin and Kerry has substantially lower fares than our model would predict. To avoid the inclusion of these observations biasing our main estimates, we exclude them from our final results, meaning our models are estimated in two stages: first with, then without outlier observations.

We have defined outlier observations as those with an estimated residual greater than three standard deviations from zero. In our core models, this always means the removal of fewer than 1% of observations from the sample. Rather than trimming our sample generally, this method primarily results in the removal of a few esoteric routes, with roughly half of exclusions accounted for by 10 routes and over two-thirds of outlier exclusions accounted for by 25 routes out of a total of more than 700.

Figure 63 Impact of outlier exclusion on congestion premium estimates

	Short haul			Long haul		
	Est. premium	Outliers	Est. with outliers	Est. premium	Outliers	Est. with outliers
MIN	21%	64	20%	29%	13	28%
1	22%	95	21%	26%	23	27%
2	15%	96	15%	25%	23	25%
3	15%	96	15%	25%	23	25%
4	17%	110	17%	24%	31	25%
MAX	14%	104	14%	23%	30	24%

Source: Frontier analysis of IATA data

As the table above shows, this decision does not fundamentally alter our results. In all cases, the models have been run without excluding outliers and produce results broadly consist with post-exclusion models. The purpose of excluding outliers is simply to instil confidence that our results reflect true and robust statistical relationships rather than quirks of data or the results of model selection, and to improve the precision of our estimates of the model parameters.

Controlling for additional drivers

The full list of control variables is shown in Section 4.2, including the source. The estimated coefficients associated with those controls in our core models are shown in Figure 70 and Figure 71.

A.3 Modelling congestion

In this section we describe in more detail our approaches to measuring congestion in our econometric models. In addition to the three approaches taken in the main body of the report, we also discuss a further approach using runway utilisation which we have tested, although we do not endorse it as an appropriate way to pick up the effect we seeking to measure.

All of the estimated coefficients are shown in Figure 70, Figure 71 and Figure 72.

Our core method of estimating the size of the congestion premium is to include a dummy variable in our sample for routes originating from Heathrow in each year. Since we are modelling fares in logs, the coefficient on the Heathrow dummies is close to, but not identical to, the estimate of the congestion premium. To transform the coefficient into a congestion premium we take the exponent and subtract 1.

We believe this is the most appropriate option for modelling the congestion premium, as it does not pre-suppose the way in which differentiation, excess demand and fares interact. While the theoretical basis of the congestion premium is explored at length in this report, we have aimed to keep the empirical specification as simple as possible, and use alternative approaches to reinforce our core findings.

To confirm our core estimates, the main body of the report also uses two alternative approaches to estimate the counterfactual level of fares at Heathrow:

- **Unserved demand**: this model uses the estimated unserved demand variable from our top down analysis as an input to the econometric model,
- **Out-of-sample estimation**: this model estimates coefficients on the sample excluding Heathrow, and applies those parameters to flights from Heathrow.

These methods corroborate the findings from our core models and confirm our estimate of the congestion premium. In addition to these, we discuss below the use of runway utilisation as a measure of the congestion premium, why we believe it is not appropriate, and what the results of its use can tell us.

Unserved demand

As discussed at length in Sections 1 and 2, the key condition that drives the congestion premium at Heathrow is the presence of demand that the current capacity cannot serve, and so prices must rise to choke this off.

The purpose of our approach in Section 2 was to provide further evidence that our estimates of the congestion premium were driven by unserved demand and not some other quirk of our econometric approach. However, it also provides the opportunity to directly test the relationship between two variables of interest – unserved demand and the congestion premium.

Testing the unserved demand variable in our models is equivalent to fixing a relationship between the yearly Heathrow dummies we are currently using. This is because a) only Heathrow has unserved demand according to our model and b) we only estimate one coefficient for the unserved demand variable.

Out-of-sample estimation

The out-of-sample estimation approach is entirely separate from the other approaches in this report and from our previous approaches to estimating the congestion premium. In our core models, we are comparing the difference in conditional means of routes that do and do not originate from Heathrow. In our out-of-sample approach, actual fares at Heathrow routes do not affect our modelling.

To run this procedure, we take three steps:

- **Estimate the drivers of fares on a sample of comparable airports**: using our European sample (minus Heathrow) we estimate the coefficients on other variables we believe to drive fares. For example, we estimate that kilometre in distance adds X to the fare.
- **Predict fares for Heathrow routes using those coefficients**: since we know the other characteristics of the routes from Heathrow, we can use the estimated model to predict the fares “out-of-sample”
- **Compare out-of-sample predictions to actual fares**: the average difference between the actual fares and the predicted fares is assumed to be the congestion premium (although we would not expect that difference to be uniform)

We do not consider this to be a core approach for two reasons:

- Heathrow is a comparable airport, and should be included in our sample when estimating the drivers of fares
- Our modelling approach is not intended to measure the drivers of fares exactly, but rather to control for them while comparing Heathrow and non-Heathrow routes. For the same reasons we are sure that our estimates of the congestion premium are unbiased, even while our estimates of other coefficients may not be, the out-of-sample estimates are likely to be worse at isolating the congestion premium than our core models.

Runway utilisation

Additionally, in this section we consider one further proxy of observable congestion: runway utilisation. We have excluded this analysis from the main body of the report as runway utilisation does not capture the nature of the congestion premium that we are seeking to measure. To reiterate, the premium we are trying to isolate is driven by the volume of *unserved demand*: the extent to which the unconstrained demand to fly from Heathrow at normal market fares exceeds the airport’s capacity to accommodate that traffic. But this measure is not directly observable. An airport with significant excess demand is likely to have very high runway utilisation. But for instance, a measure of runway congestion cannot tell the difference between an airport which is operating at 99.5% has been doing so for a number of years and has considerable unserved demand and an airport which is operating at 99.5% capacity but has little or no excess demand, perhaps because it has only just reached that point. . Observed runway congestion is the same in both cases, but we would expect the premium to be much larger under the first case.

As noted, an airport with significant unserved demand is likely to have very high runway utilisation, so we would expect to see a *partial* correlation between runway utilisation and the congestion premium. But we would not expect this to be systematic.

Indeed it is possible, although clearly a matter for empirical analysis, that one could observe a positive relationship between fares and runway utilisation with no unserved demand being present. Such a result could imply that, for various reasons, airlines find it more expensive to operate at an airport operating close to maximum utilisation (e.g. increased delays on the ground and holding on landing, loss of operational flexibility). Provided, as discussed previously, there is some degree of airport differentiation, these costs could be expected to be reflected in higher fares at the busy airport in the long run, through normal competitive pressures.

But as airports with significant unserved demand will also have high runway utilisation, there is a significant degree of overlap between these two variables, which means they cannot both be included in a model without them interfering with each other and producing biased estimates. Including only the runway utilisation variable will pick up some, but not all effects of unserved demand on price. In our view the same risks do not apply if runway utilisation is omitted. In this case estimates of the congestion premium will include that part of the premium that arises from higher airline costs. But this is still a legitimate element in the premium as those increased airline operating costs could also be expected to be reduced if the capacity constraint were alleviated.

However, leaving these concerns aside, for the benefit of completeness we have performed a sensitivity around runway utilisation.

The chart below shows the runway utilisation at Heathrow and the other airports in the sample over the period for which we have data on ticket fares (2011-2018).

Figure 64 Runway utilisation in the European Sample

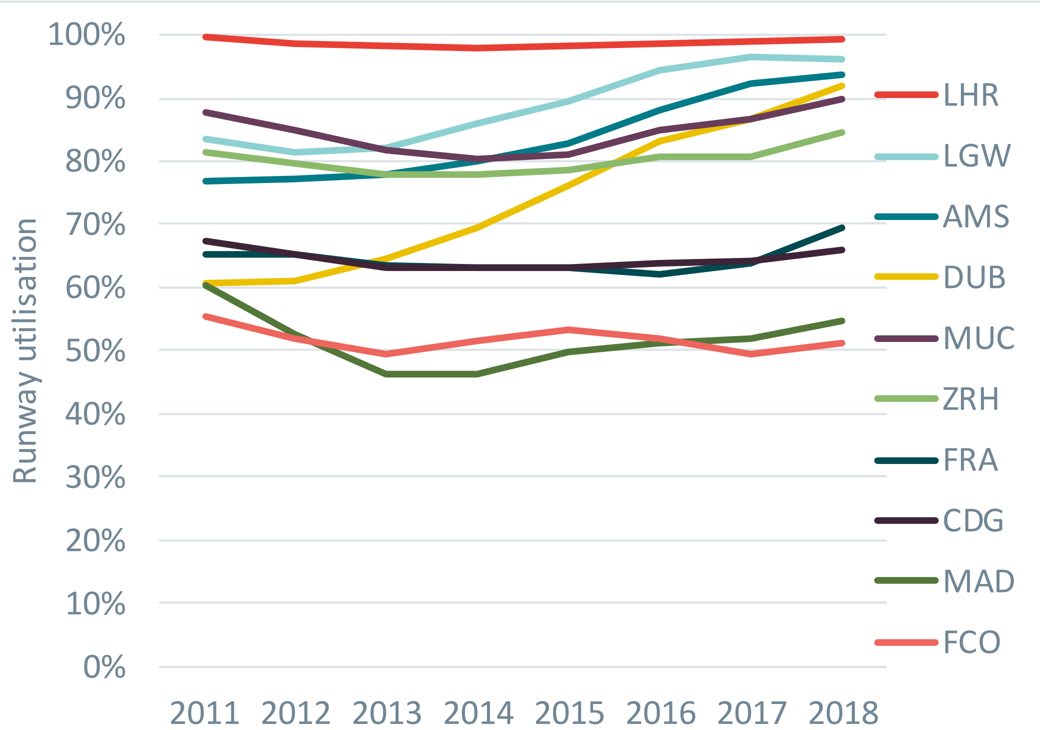

Source: Frontier Economics analysis of data form Eurocontrol and OAG Schedules Analyser

Under this approach we can estimate the relationship between runway utilisation in general and prices. But to produce a concrete estimate of the *premium* at Heathrow under this approach we also need to construct a 'counterfactual' scenario of Heathrow's utilisation rate if it were not constrained, and therefore, by how much prices would be lower under this scenario?

The approach we have taken is as follows. For 2018, we estimated in the top down analysis that unserved demand at Heathrow was 17.5%. Following expansion, Heathrow's runway capacity is estimated to increase by 54%. This means that if Heathrow had already expanded today, then its 2018 runway utilisation could be 117.5%/154% = 76%.

An alternative approach would be to take the average utilisation rate amongst the other airports in our sample. We estimate this to be 77%. Given the closeness of these two numbers, we have assumed that the counterfactual is that Heathrow would have a runway utilisation of 76%. We can then estimate the total reduction in prices under this scenario as an estimate of the congestion premium. We note that this may still underestimate the true size of the premium, as in principle Heathrow could still be constrained at many times of the day under this counterfactual.

The additional controls used are identical to those in one of our core models (specifically model 2), but these models have been tested with other sets of controls, and the relationship between runway capacity and fares remains consistent.

The results of the runway utilisation approach are shown in the table below, where the models produce similar results to the core models, but do not explain away the Heathrow-specific measures of congestion.

Figure 65 Log runway utilisation models of congestion

	Short haul	Long haul
Log transformed runway utilisation coefficient	0.03901***	0.04750***
Percentage congestion premium	14%	17%
Adj. R-squared	0.42347	0.48562
RMSE	0.27957	0.31881

Source: Frontier analysis of IATA data
Both models estimate average log fares for routes from our European hub sample for the years 2011 to 2018 using OLS
** denotes statistical significance at the 10% level; ** denotes statistical significance at the 5% level; *** denotes statistical significance at the 1% level*

A.4 Assessing models

In this section, we discuss the measures we have taken to ensure our models are accurately measuring the congestion premium. Specifically, we outline diagnostic testing undertaken and some key modelling cross-checks.

A.4.1 Diagnostic tests

Diagnostic tests were conducted for normality of the residuals, heteroskedasticity and Reset tests for non-linearity. In this section we will address the purpose of these tests and present the results for each.

The purpose of the diagnostic tests, respectively, are as follows:

- **RESET test** is a test for non-linearities in the chosen independent variables, a form of omitted variable bias. It uses a regression of the residuals on the squared fitted values– if there is a relationship between the two then some non-linear transformation of one or more independent variable is needed.

- **Heteroscedasticity test** estimates whether the variance of the residual is constant across the sample. If it is not, then the standard errors are not correct, although the estimated coefficients remain unbiased.
- **Residual normality** tests for whether the higher moments of the distribution of the residual are consistent with a normal distribution. That is, while the mean and standard deviation of the residuals are consistent with a normal distribution by definition, the skewness and kurtosis may not be. While the OLS estimator remains the best linear unbiased estimator in this case, it may indicate that the model is mis-specified.

For these tests, it is important to first note the null hypothesis is that the assumption holds, the residuals are normally distributed with constant variance and not related to powers of the fitted values. Large samples make the tests very sensitive, so there is a high probability of rejecting the null hypothesis even when there is only a very small and effectively irrelevant deviation from it.

The results for the RESET, heteroskedasticity and residual normality tests for each of our core models is shown in Figure 66 in the form of p-values for the acceptance of the null hypothesis. In all tests and all models for our short haul models, the null hypothesis is rejected at the 1% level. The results for long haul are less clear, but nonetheless largely show the rejection of the null hypothesis.

Figure 66 Diagnostic test p-values for core models

	MIN	1	2	3	4	MAX
Short haul						
RESET test	<0.001	<0.001	<0.001	<0.001	<0.001	<0.001
Heteroskedasticity test	<0.001	<0.001	<0.001	<0.001	<0.001	<0.001
Residual normality test	<0.001	<0.001	<0.001	<0.001	<0.001	<0.001
Long haul						
RESET test	<0.001	0.119	0.001	0.001	<0.001	<0.001
Heteroskedasticity test	<0.001	<0.001	<0.001	<0.001	<0.001	<0.001
Residual normality test	<0.001	0.007	0.004	0.004	0.006	0.009

Source: Frontier analysis

We address our response to each of these tests in turn.

Residual normality

The Shapiro-Wilk test for residual normality is particularly prone to problems of large sample size, and the statistics package used recommends not using it for sample sizes greater than 2000 observations. To judge how serious the problem was in cases of rejection, we plotted the histogram of the residuals and compared it to a normal distribution. As Figure 67 shows the residuals are very close to normal. Similarly, in cases where there was a failure on the other tests the seriousness of the failure was investigated graphically to see if there was serious misspecification. In many cases failure on the tests seemed to be caused by outliers. We note that normality is not required for the best linear unbiased estimator properties of the estimator.

Figure 67 Distribution of estimated residuals from SH Core Model 2

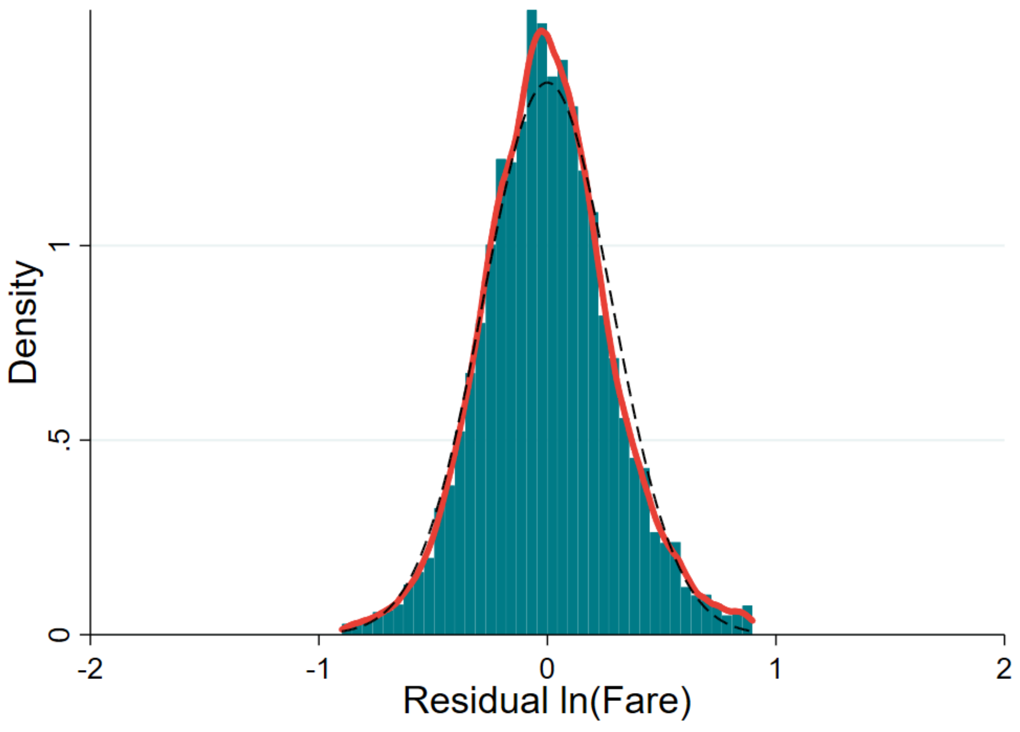

Source: Frontier analysis of IATA data

Heteroskedasticity

Failure of the homoscedasticity assumption leaves the estimator unbiased but not minimum variance and so the conventional standard errors are biased. Most importantly for our analysis, this might mean we had incorrectly interpreted the significance of our results. However, to deal with this issue, we report all coefficients with heteroskedasticity robust standard errors therefore avoiding this issue.

RESET test

The RESET test estimates whether the polynomial terms of the *predicted* value have any predictive power in explaining the dependent variable. If the test rejects the null hypothesis, it means that the model is mis-specified and that some non-linear or polynomial form of an explanatory variable ought to be included. This is the most concerning failure in our diagnostic tests, since it could mean that our estimates of the congestion premium are biased if the omitted non-linear variables are correlated with the premium. This could not be fixed by a simple statistical procedure. In our view it is likely that, as with normality, a large part part of the explanation is the very high sensitivity of the test to reject the null hypothesis even when there are very small deviations from it. It is nonetheless important to understand exactly what that deviation is.

In the main body of the report, to ensure we are not mis-specifying the congestion premium, we show that the source of non-linearity is not correlated with the Heathrow variable and so must be attributable to something else. As suggested by scatterplots like Figure 39, a key potential source of non-linearity is the irregular relationship between distance travelled and fare paid. We have already attempted to control for this by separating short and long haul flights, however this is insufficient for passing the RESET test.

The simplest form of controlling for the presence of non-linearities uses polynomial terms of the variable of interest, which can approximate a more complex function. Including a squared term for distance would alter the fit of log fares shown in Figure 39, and instead produce that shown in Figure 68 below.

Figure 68 Quadratic fit of log fares with respect to distance, 2018

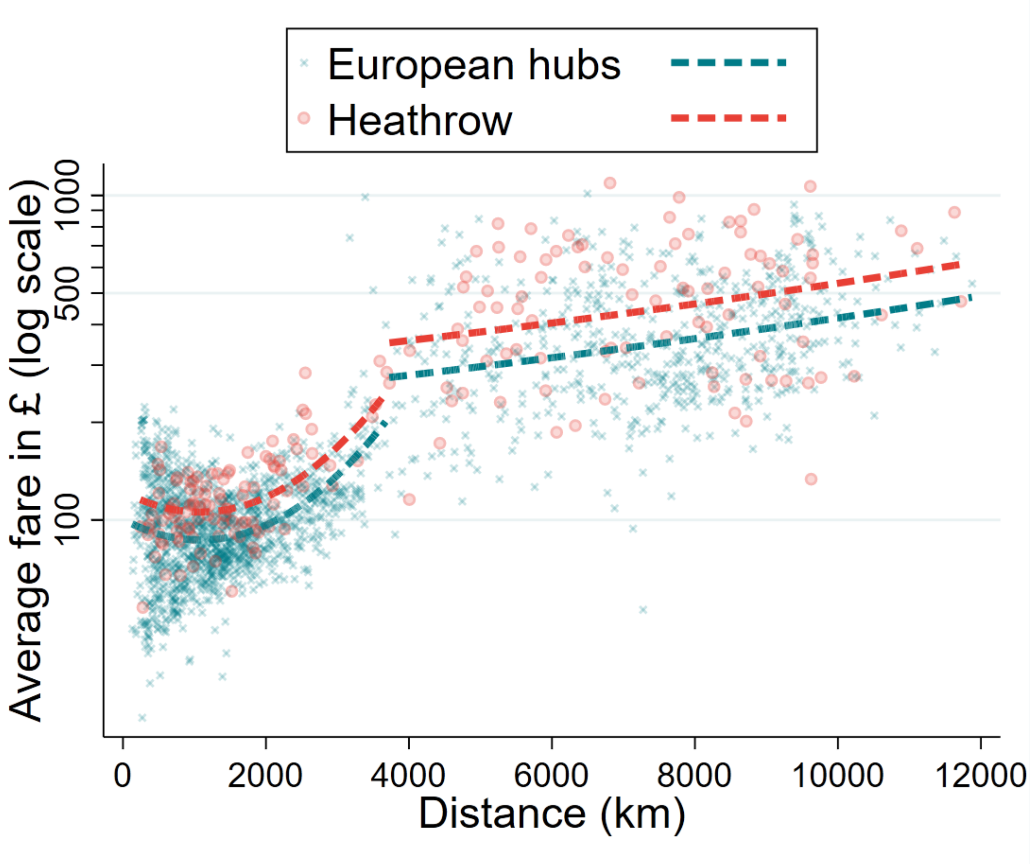

Source: Frontier analysis of IATA data

Note: Each point represents a route from either Heathrow or a European long haul hub in terms of distance and average fare. The dashed lines represent a linear fit of log prices with respect to distance and distance squared, fitted separately for short and long haul flights.

To test whether this could explain the RESET test results, we tested the inclusion of a squared term for distance in our simplest short haul model – the specification that failed most clearly. We found that the inclusion led to no longer rejecting the null hypothesis at the 5% level. This approach was confirmed by an even less parametric approach to controlling for distance, in which the relationship was estimated separately for each 1000km distance group. Using the distance buckets approach, the RESET test was not rejected at the 10% level. As well as allowing our model to pass the RESET test, we find that they do not impact the estimates of the congestion premium. As the table below shows, including the distance buckets as controls leaves the congestion premium estimates unchanged across all of our models.

Figure 69 Congestion premium estimates with alternative measures for distance

	Short haul		Long haul	
	CM	Alt. Distance	CM	Alt. Distance
MIN	21%	19%	29%	27%
1	22%	20%	26%	26%
2	15%	14%	25%	25%
3	15%	14%	25%	25%
4	17%	16%	24%	24%
MAX	14%	13%	23%	23%

Source: Frontier analysis of IATA data

These results, done on our largest sample and largest failure, suggest that the relationship between fares and distance is the source of omitted variable bias, and that this bias does not affect the variable of interest, the congestion premium. We still find that the RESET test fails on our more complex model even with the inclusion of non-linear distance terms. This likely reflects non-linear interactions with other controls, rather than any fundamental bias to the congestion premium. If there were a risk that the non-linearities were impacting the congestion premium, this would have been visible in the narrower models where the distance controls solve the failure of the RESET test.

A.5 Coefficient tables

This section contains the full tables of estimated coefficients from all of the models discussed in our report.

Figure 70 Core models – short haul (coefficient table)

Variable	MIN	1	2	3	4	MAX
Distance (000s KM)	0.00018***	0.00026***	0.00025***	0.00025***	0.00031***	0.00031***
	{0.00000}	{0.00000}	{0.00000}	{0.00000}	{0.00000}	{0.00000}
Frequency (own)		-0.01508***	-0.00671**	-0.00671**	0.02912***	0.03015***
		{0.00315}	{0.00305}	{0.00305}	{0.00338}	{0.00337}
Frequency (other)		0.02165***	0.01792***	0.01792***	0.02345***	0.02429***
		{0.00212}	{0.00208}	{0.00208}	{0.00201}	{0.00201}
Transfer share		0.62248***	0.44665***	0.44665***	0.32723***	0.30452***
		{0.01374}	{0.01506}	{0.01506}	{0.01513}	{0.01528}
LCC share			-0.28474***	-0.28474***	-0.20560***	-0.19799***
			{0.01028}	{0.01028}	{0.01055}	{0.01062}
Business class share			0.0001	0.0001	-0.00225	0.02794
			{0.02253}	{0.02253}	{0.02238}	{0.02296}
European jet fuel price (logs)				0.17962***	0.17222***	0.17408***
				{0.01089}	{0.01041}	{0.01038}
Avg. seats					-0.17233***	-0.17117***
					{0.01256}	{0.01249}
Airport competition (inverse)					0.07963***	0.06988***
					{0.01145}	{0.01155}
HHI (route level)					0.29036***	0.29158***
					{0.01757}	{0.01740}
Single carrier dummy					0.00119	0.00095
					{0.00924}	{0.00919}
Skytrax rank (100 if unlisted)						-0.00070***
						{0.00009}
Heathrow dummies	✓	✓	✓	✓	✓	✓
Year dummies	✓	✓	✓	✓	✓	✓
Constant	✓	✓	✓	✓	✓	✓
Observations	10,427	10,396	10,395	10,395	10,381	10,387
Total congestion premium (£bn)	£3.33	£3.45	£2.53	£2.53	£2.84	£2.45
Total congestion premium (%)	21%	22%	15%	15%	17%	14%
Adj. R-squared	0.21966	0.37636	0.41892	0.41892	0.47112	0.47235
RMSE	0.32744	0.29147	0.28140	0.28140	0.26839	0.26831

Source: Frontier analysis of IATA data
All models estimate average log fares for routes from our European hub sample for the years 2011 to 2018 using OLS, robust standard errors are reported in brackets
Each model is estimated with outliers (residual >3 sd) exlcuded
** denotes statistical significance at the 10% level; ** denotes statistical significance at the 5% level; *** denotes statistical significance at the 1% level*

Figure 71 Core models – long haul (coefficient table)

Variable	MIN	1	2	3	4	MAX
Distance (000s KM)	0.00008***	0.00006***	0.00005***	0.00005***	0.00007***	0.00007***
	{0.00000}	{0.00000}	{0.00000}	{0.00000}	{0.00000}	{0.00000}
Frequency (own)		0.07895***	0.06678***	0.06678***	0.10095***	0.10154***
		{0.00520}	{0.00509}	{0.00509}	{0.00608}	{0.00614}
Frequency (other)		0.05627***	0.05424***	0.05424***	0.05808***	0.05857***
		{0.00376}	{0.00375}	{0.00375}	{0.00360}	{0.00362}
Transfer share		0.51700***	0.45151***	0.45151***	0.40549***	0.39367***
		{0.02040}	{0.01994}	{0.01994}	{0.02000}	{0.02168}
LCC share			-0.37146***	-0.37146***	-0.26582***	-0.26407***
			{0.03395}	{0.03395}	{0.03311}	{0.03314}
Business class share			0.53955***	0.53955***	0.44506***	0.44246***
			{0.03726}	{0.03726}	{0.03538}	{0.03542}
European jet fuel price (logs)				0.15889***	0.15906***	0.16137***
				{0.01889}	{0.01823}	{0.01823}
Avg. seats					-0.26303***	-0.26141***
					{0.02591}	{0.02588}
Airport competition (inverse)					0.21921***	0.21824***
					{0.02925}	{0.02923}
HHI (route level)					0.29015***	0.29239***
					{0.03273}	{0.03275}
Single carrier dummy					-0.06247***	-0.06388***
					{0.01491}	{0.01490}
Skytrax rank (100 if unlisted)						-0.00034**
						{0.00016}
Heathrow dummies	✓	✓	✓	✓	✓	✓
Year dummies	✓	✓	✓	✓	✓	✓
Constant	✓	✓	✓	✓	✓	✓
Observations	4,454	4,444	4,444	4,444	4,436	4,437
Total congestion premium (£bn)	£13.10	£13.93	£13.50	£13.50	£12.38	£11.92
Total congestion premium (%)	29%	26%	25%	25%	24%	23%
Adj. R-squared	0.14967	0.45526	0.49006	0.49006	0.52499	0.52535
RMSE	0.41124	0.32797	0.31761	0.31761	0.30600	0.30619

Source: Frontier analysis of IATA data
All models estimate average log fares for routes from our European hub sample for the years 2011 to 2018 using OLS, robust standard errors are reported in brackets
Each model is estimated with outliers (residual >3 sd) exlcuded
** denotes statistical significance at the 10% level; ** denotes statistical significance at the 5% level; *** denotes statistical significance at the 1% level*

Figure 72 Alternative congestion models – coefficient table

Variable	Short haul - Unserved demand	Short haul - Out of Sample	Long haul - Unserved demand	Long haul - Out of Sample
Distance (000s KM)	0.00026***	0.00025***	0.00005***	0.00006***
	{0.00000}	{0.00000}	{0.00000}	{0.00000}
Frequency (own)	-0.00524*	-0.01316***	0.07104***	0.07068***
	{0.00304}	{0.00316}	{0.00498}	{0.00507}
Frequency (other)	0.01628***	0.02020***	0.05247***	0.05323***
	{0.00206}	{0.00214}	{0.00373}	{0.00385}
Transfer share	0.44806***	0.45299***	0.44683***	0.43680***
	{0.01511}	{0.01535}	{0.01990}	{0.02101}
LCC share	-0.27916***	-0.29583***	-0.36513***	-0.37342***
	{0.01033}	{0.01048}	{0.03389}	{0.03377}
Business class share	0.01272	-0.03762	0.54241***	0.46141***
	{0.02256}	{0.02336}	{0.03741}	{0.03644}
Heathrow dummies	✓	✓	✓	✓
Year dummies	✓	✓	✓	✓
Constant	✓	✓	✓	✓
Observations	10,390	9,659	4,445	3,814
Total congestion premium (£bn)	£2.02	£3.85	£11.79	£14.30
Total congestion premium (%)	12%	23%	22%	27%
Adj. R-squared	0.41848	0.40086	0.48918	0.48362
RMSE	0.28147	0.28328	0.31836	0.31298

Source: Frontier analysis of IATA data
All models estimate average log fares for routes from our European hub sample for the years 2011 to 2018 using OLS, robust standard errors are reported in brackets
Each model is estimated with outliers (residual >3 sd) exlcuded
** denotes statistical significance at the 10% level; ** denotes statistical significance at the 5% level; *** denotes statistical significance at the 1% level*

A.6 Comparison to previous estimates

This is our third report that estimates the congestion premium at Heathrow. Our past analyses were:

- Impact of airport expansion options on competition and choice (April 2014)[27]. This was based on 2012 data.
- Competition & Choice 2017 (December 2017)[28]. This was based on 2016 data.

The approach in this report builds directly on those past efforts, but also differs in a number of important ways:

- We have used an expanded, European sample of comparator airports instead of the sample of London airports that was the focus of previous work. Previous reports also included a European sample of airports, but this smaller selection consisted of Heathrow, Amsterdam Schipol, Paris Charles de Gaulle, Frankfurt and Madrid airports. The expanded sample includes Rome, Munich, Zurich and Dublin airports.
- Our previous reports focused on individual years of data, and the cross-sectional relationships within that year. This report includes data from multiple years.
- The list of controls used in our past work was more limited – we have expanded our array of models and controls in this report to instil further confidence in the estimates of the congestion premium.

Each of these changes has had an effect on our models, however we are in each case estimating the same economic parameter – the congestion premium at Heathrow.

The table below shows how our most recent estimates of the congestion premium compare to the estimates of our previous reports. We have specifically compared the range of our new estimates from the comparable years of data. For previous estimates, we have taken those provided by the European hub approaches, so as to be as comparable as possible to our new approach. These figures differ from the headline congestion premium figures offered in those reports, which were based on the London sample.

Figure 73 Comparison to previous congestion premium estimates

Year	Short haul		Long haul	
	Current	*Previous*	*Current*	*Previous*
2012	27% to 32%	18% to 23%	21% to 26%	19% to 22%
2016	12% to 21%	12% to 23%	25% to 35%	7% to 32%

Source: *The years 2012 and 2016 refer to the estimates of the congestion premium in that year, and not the year of publication of the report.*
The range of estimates refers to the short haul and long haul congestion premium estimates of using the European hub sample of those years.

[27] https://your.heathrow.com/takingbritainfurther/wp-content/uploads/2014/04/frontier-report-Impact-of-airport-expansion-options-on-competition-and-choice.pdf

[28] https://www.caa.co.uk/uploadedFiles/CAA/Content/Accordion/Standard_Content/Commercial/Airports/HAL%20-%20Frontier%20Competition%20and%20Choice.pdf

For the most part, our expanded sample has been effective at narrowing the range of estimates of the congestion premium by increasing the sample size. For example in 2016, our models produce less diverse estimates now than we did two years ago, despite providing twice as many models and a much wider range of controls. In 2012, the results have changed more significantly, especially in short haul where we know the year-to-year variation in the congestion premium is more pronounced.

ANNEX B RESPONSE TO FTI REPORT

In October 2018, a report by FTI Consulting was published on behalf of the Civil Aviation Authority (CAA) in response to the Frontier Economics Competition & Choice Report of 2017. In their critique, FTI described the Frontier Report as follows:

> *In our view, the Frontier Report reflects a substantial econometric analysis. This analysis, whilst sound in many respects, suffers from some deficiencies that cast doubt on the accuracy of the 23% scarcity rent premium identified.*

They list a number of shortcomings in Frontier's analysis:

- our "residual" approach to estimating the congestion premium may be attributing other Heathrow-specific elements to scarcity rents;
- we weight each route in our analysis equally and therefore overweight routes with a low number of passengers (and vice versa);
- our estimates may suffer from omitted variable bias due to the lack of controls for route-level competition;
- the coefficients on some of our control variables are unstable, suggesting misspecification; and
- we do not adequately justify our assumed relationship between the proposed expansion and the relieving of capacity constraints

They also provide a number of areas for potential consideration in future research on the topic:

- more granular data on fares at an individual, rather than route level;
- controlling for a wider range of variables, including passenger preferences and differences in within-airport route competition; and
- improved diagnostic testing of our results

We are grateful to FTI for their observations on our previous work, as they have helped to focus and clarify our analysis for this report. In this annex, we address each of the points raised in turn, noting where we have adjusted our approach as a result. Obviously, we have also endeavoured to reflect our response to all these comments in the preceding report. Overall, we believe that none of these objections undermine the core economic relationship between congestion and fares, nor do they change the estimated premium on fares paid.

B.1 "Residual" approach

> *The residual approach used assumes that all differences in price other than those reflected in the regression model can be attributed to scarcity rents. The 23% 'scarcity rent' premium may reflect Heathrow-specific elements (such as its more or less convenient travel links, actual and/or perceived differences in the quality of the airport, etc.)*

Before discussing this as a challenge to our empirical approach, it is important to reiterate the theoretical relationship between capacity and differentiation we have outlined in Section 1. It is not only the fact that Heathrow is full that leads to a congestion premium, but the fact that it is to some extent differentiated from other airports in or near London that may result in passengers paying higher fares. This differentiation does allow airlines to capture a small proportion of the difference between passenger willingness to pay and long run marginal cost.

But to explain systematic premium between Heathrow and other airports it would be necessary to argue that Heathrow is *more* differentiated from airports around it than the other airports in the sample. This is simply not plausible.

The most important differentiator is clearly location: passengers systematically show a preference for their local airport, all other things being equal. So to base an assumption that fares can capture a greater proportion of willingness to pay at Heathrow than, say, at say, Frankfurt, or Schiphol or Charles de Gaulle we have to suppose Heathrow has a greater surface access advantage over its catchment than those airports. No evidence we have seen leads us to think that is the case. Indeed given the number of independently owned airports in the London region and the quality of surface access this is not plausible.

The alternative interpretation would be that the quality of service at Heathrow is so exceptional, relative to other airports that passengers would not switch and so even the competitive level of fares, without the physical constraint would reflect this fact. Again, we see not plausible evidence that Heathrow is such an outlier. Heathrow is a good airport, and a much improved one in the last ten years to judge by passenger satisfaction measures, but this cannot reasonably be used to explain a premium of 20%+ on fares in an otherwise competitive market.

We can, in any event, capture variations in airport quality in our core models by the inclusion of an airport quality variable in our core models, which show no sign of materially impacting on the estimated fare premium..

We agree that it theoretically possible that there could be other Heathrow-specific characteristics that could be driving the difference in fares. It may be the case that, while Heathrow is uniquely constrained, there is some other unknown factor driving fares. The aim of this report has been to conclusively show this is not the case, and that congestion is *the* driver of higher fares. In particular:

- Our core econometric analysis is shown across a range of specifications with a number of different control variables. Whether controlling for capacity, distance, seat mix or other variables, our estimate of the congestion premium is consistent. The difference in fares at Heathrow is therefore not the result of an unusually high or low level of some common driver of fares. It is also not the case that a quirk of our model selection has driven our estimate.
- Top-down analysis of unserved demand – an explicit measure of congestion – produces similar levels of congestion premium to our econometric approach. This is also true, and statistically significant, when included as a control in our econometric analysis, as does use of out of sample estimation of the premium.
- Furthermore, prices paid by airlines for slots at Heathrow are consistent with the 25% premium under reasonable assumptions about cost of capital. Slot premia can only be explained by the presence of economic rents that require

access to the relevant slots. There would be no economic rationale for paying for them in the absence of a fare premium.

We are confident that our estimates of higher fares at Heathrow reflect one Heathrow-specific factor in particular: congestion. Our estimates of the scale of the congestion premium (roughly 25%) are consistent across a series of independent approaches we have followed. We do not believe it is possible to plausibly attribute this difference in fares to some other unknown factor.

B.2 Weighting

Aggregate data is used, which means that the fares on the least busy routes have a greater impact on the coefficient estimates than the busiest routes

The question being addressed by Frontier's econometric approach is whether fares are higher at Heathrow than if there were no (or fewer) constraints on airport capacity. To do this, we use variation in average fares paid across routes, controlling for known drivers of average fares on a route. FTI suggests that this over-weights the impact of fares paid by passengers on less popular routes.

The critique suggested by FTI contrasts our approach with one that either uses individual fares directly (and thereby implicitly weights passengers equally) or a weighted regression that gives greater weight to routes with more passengers. The use of individual fares is not feasible, since the data is not available, however it is possible to weight the route-level observations in accordance to the number of passengers.

The issue of weighting in econometrics is a complex one, and is therefore covered by a number of papers[29]. In our view the lack of weighting by passenger numbers could be a problem if we were estimating population statistics (e.g. the average fare paid by passengers at Heathrow). However, academic research such as the paper cited shows this is not a source of bias in estimating the causal impacts on fares using OLS.

There are situations in which it may be appropriate to use Weighted Least Squares (WLS), however best practice in these situations is to clearly identify the source of potential bias we are aiming to address with weighting. We have no evidence to show that the size of the congestion premium varies systematically with the size of a route, nor do we believe that the selection of routes is biased, as it covers the full population. Hence weighting is not supported and no reason to consider our results are biased.

Nonetheless, we have also produced WLS estimates of our models that are consistent with OLS, which demonstrates that not only is weighting not supported theoretically but it makes no difference to our findings empirically. In fact, as Figure 74 shows, using OLS rather than WLS reduces our estimate of the congestion premium by about 5 percentage points, although this is not consistent across models.

[29] For example, Solon, G., Haider, S. J., & Wooldridge, J. M. (2015). What are we weighting for?. Journal of Human resources, 50(2), 301-316.

Figure 74 Congestion premium estimates using OLS and WLS

	Short haul		Long haul	
	OLS	**WLS**	**OLS**	**WLS**
MIN	21%	34%	29%	39%
1	22%	31%	26%	34%
2	15%	20%	25%	29%
3	15%	20%	25%	29%
4	17%	15%	24%	23%
MAX	14%	12%	23%	21%
Average	**17%**	**22%**	**25%**	**29%**

Source: Frontier analysis of IATA data

B.3 Omitted variables

The analysis likely suffers from omitted variable bias, where relevant explanatory variables are missing from an econometric model. 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.

Omitted variable bias is a common source of error in empirical modelling, in which the absence of an important explanatory variable can bias the coefficient found on other estimated coefficients.

It would be correct to state that our analysis has likely omitted variables – there are a number of important drivers of variation in fares that we have not included, and even our best fitting model explains only 50% of the variation in log fares. However, this does not imply that our estimates of the congestion premium are *biased*, unless the omitted variable in question is disproportionately present on Heathrow routes. Since we are only interested in the scale of the congestion premium, the potential omitted variable would have to both be correlated with fares and be correlated with our Heathrow variables – this is the formal definition of omitted variable bias.

To ensure that omitted variable bias is not affecting our results, we have tested a wide range of controls in our core models, covering all reasonable drivers of fare differentials, to show that estimates of the congestion premium are robust to the inclusion of other variables in the model.

In addition to raising the point of omitted variables generally, FTI has pointed specifically to the lack of controls for intra-route competition between airlines. The impetus behind this suggestion is that the more airlines need to compete on a given route, the lower fares will be. This is a reasonable suggestion for a potential source of omitted variable bias, as it is likely to drive fares through the standard economic logic of competition, and could be correlated with our Heathrow variables.

The important point to raise here is that *if* routes at Heathrow had elevated fares because of less intra-route competition than at other airports you would have to ask: why specifically is there less competition at Heathrow. The answer will inevitably involve the lack of access to slots at Heathrow, which is another result of the capacity constraint. Hence fares may partially be raised because there are fewer competitors and partially raised for a given number of competitors due to reduced threat of competitive entry. But either way the net result is a premium due to congestion.

However, to see how route competition may affect fares, we have included two measures in our core models: a route-level Herfindahl-Hirschman Index (HHI) and a dummy for routes served by only one airline (see Figure 70 and Figure 71). As suggested by FTI and economic theory more generally, the coefficient on the HHI term is strongly positive, indicating higher fares as competition reduces. By contrast, the effect of having a single carrier on a route is roughly 6% lower than an HHI equal to 1 would suggest, perhaps reflecting the incentives of entry for new airlines. Through all of this however, the congestion premium estimates are extremely consistent. So while, lack of intra-route competition in fact could be expected to be part of the congestion premium, explicitly controlling for it makes no difference to our findings.

B.4 Unstable coefficients

Some coefficients appear unstable, indicating that the coefficient estimates could be biased in at least some of these model specifications.

A lack of stability in the coefficients of our model could indicate a broader failure to specify correctly the relationship between control variables and fares. However, the goal of Frontier's analysis is not to provide an accurate measurement of the impact of every included variable on the average fares paid. The inclusion of these variables in our models is intended to show that our estimate of the congestion premium is robust to the inclusion of other potential drivers of fares.

For example, we include two direct measures of supply in our models – frequency and average plane size. In models without average plane size, the effect of frequency on fares is more strongly negative. This does not mean that our estimate of congestion is biased, but that the two control variables are endogenous to fares or correlated with each other. The factors in question are controlled for, and our congestion premium estimates are robust to them, we simply are not able to use the coefficients on the control variables in the same predictive or causal way as the variables of interest.

B.5 Real-world application to expansion

The conclusion reflects an assumption that the expansion completely removes the capacity constraint

…

The Frontier Report does not discuss that the aero charges (per passenger) may change as a result of expansion … it is likely that this would have an impact on post-expansion ticket prices

Our top down analysis implies that in 2018, excess demand was 17%. The expansion at Heathrow will lead to 54% more movements which would therefore completely cover the current excess demand and lead the congestion premium to disappear. Arguably all busy airports face a congestion premium at some point during the day or year, as there will always be more demand for touching down at 8am (for example) and that the same would apply to a post-expansion Heathrow. That said, the scope for passengers and airlines to switch in response to congestion at peak times would be substantially increased if the capacity constraint were relieved *today.* Our estimate of the premium is thus an estimate of the current welfare loss resulting from congestion. We accept that what happens in 2026 depends on the extent of excess demand present by that point. But that in no way invalidates the observation that fares in 2018 are significantly elevated because of today's capacity constraint.

B.6 Granularity of fare data

Specifying the dependent variable at a more granular level, that is, at the level of individual flight tickets

Frontier's analysis of the congestion premium, both in this report and previous reports, is based on route-level averages of fares. The parameter of interest – the increase in fares driven by capacity constraints – is in fact paid on a per-passenger basis. Here, FTI is suggesting that analysis of fares at a or more granular level could result in a better estimate of the congestion premium and a worthwhile extension of our work.

While specifying the dependent variable at a more granular level may increase the power of our models, it would only lead to more precise estimates of the same parameters, and either way the relevant data is not available. The key source of variation is that some routes originate from Heathrow, while others do not, and that these two groups have a different conditional distribution of fares. Using the average fare paid on a route rather than the full distribution of fares on a route provides less information from which to estimate the congestion premium, but our current set of estimates have more than enough power to do so.

That said, more granular data would allow a number of interesting extensions to our work: for example we could isolate the within-day and within-year variation in scarcity rents (something we would expect to see). We might be able to identify the peak hours during which many other major airports are constrained. We could estimate the effect of individual passenger characteristics on their willingness-to-pay for fares at above-cost rates, and other behavioural responses of passengers to scarcity rents. However none of these worthwhile extensions are germane to the substantial, aggregate costs of limited runway capacity at Heathrow.

With the data currently available to Frontier (and other analysts using equivalent sources), the most detailed level of analysis possible is at monthly averages of routes rather than yearly. For the reasons discussed, the benefits of this granularity would be limited, and does not address the kind of individual-level variation that FTI is interested in.

B.7 Wider set of controls

Controlling for a wider range of explanatory variables, including (but not limited to) variables accounting for passenger preferences for flying from Heathrow, and differences in within-airport route competition

We agree that including a wider range of controls helps to show estimates of the congestion premium are robust, and not sensitive to the choice of controls. Hence the approach we have taken in our core models. The point regarding the inclusion of passenger preferences and differentiation is discussed in detail above, however we have included many of these variables in our regressions for completeness.

B.8 Further diagnostic tests

Performing further and more detailed diagnostic tests of the relevant regression specifications

Frontier believes that this is an appropriate extension on previous work, and is addressed in detail in Section A.4.1. It is nonetheless important to reiterate that the purpose of our modelling has at no point been to provide accurate predictions of fares, nor of the parameters driving fares beyond the congestion premium. We believe that we have clearly justified the validity of our estimates.

frontier economics

www.frontier-economics.com