RBB | Economics

[redacted]

The effect of congestion at Heathrow Airport

Comments on Frontier and FTI reports

RBB Economics, 11 January 2019

[redacted]

Contents

1 Introduction and Executive Summary

The present report, prepared by RBB Economics at the request of British Airways and International Airlines Group, contributes to the economic analysis on the effect of congestion at London Heathrow Airport ("**Heathrow**"), which is by now a well-established fact.

One issue that was brought to the forefront is the extent to which the shortage of slots at Heathrow gives airlines that use these slots the ability to raise ticket prices, thus enabling them to realise so-called *scarcity rents*. A recent report by Frontier Economics at the request of Heathrow Airport Limited – the "**Frontier report**" – argues that capacity constraints at airports leads directly to higher ticket prices. This argument is based on an "*economic theory of capacity constraints*" which focuses on "*price setting in terms of supply and demand for seats*" at Heathrow.[1] To empirically test its thesis, Frontier has undertaken a "*detailed econometric analysis to estimate the cost of the [scarcity rent] today*".[2] The Frontier report finds that, in 2016, "*ticket fares at Heathrow were on average 23.3% higher than at other London airports and 24.4% higher than at other European hub airports, due to the congestion premium*".[3] These results lead Frontier to conclude that congestion at Heathrow has given rise to a significant premium on ticket prices.

FTI Consulting, at the request of the Civil Aviation Authority ("**CAA**"), has developed a critique of the Frontier report – the "**FTI report**". The FTI report acknowledges the economic framework set out by Frontier, though it notes that this analysis does not address the fact that "*the direct constraint in the market is on the slots available to airlines*",[4] whilst "*air travel to and from Heathrow is* […] *more accurately represented as a collection of individual* [city-pair] *markets*".[5] Further, and importantly, the FTI report identifies several "*deficiencies [regarding the Frontier's econometric analysis] that cast doubt on the accuracy of the 23% scarcity rent premium*". Nevertheless, FTI considers that the "*Frontier report reflects a substantial econometric analysis*", thus lending some credibility to Frontier's analysis.[6] The FTI report also contains an assessment of IAG's response to the Frontier report. FTI "*does not agree with IAG's view that airlines currently compete with one another on a route-by-route basis, […], such that scarcity rents are competed away*".[7] FTI, therefore, appears to consider that airport congestion necessarily leads to some rents that may well accrue to airlines.

In this report, we show that the analysis presented in the Frontier report (and acknowledged to some extent by FTI) is flawed. In particular, the Frontier report does not establish a clear, robust link between the shortage of slots at Heathrow and passenger ticket prices, casting doubts on Frontier's prediction that ticket prices are higher at Heathrow because of congestion. Instead, we show that airport congestion is expected to raise airlines' cost of using slots. However, since this cost is mostly fixed (i.e. it does not vary with the number of passengers), any increase is unlikely to affect passenger ticket price.

---

1 Frontier report, p. 20.
2 Frontier report, p. 5.
3 Frontier report, p. 28, emphasis added.
4 FTI report, para 3.12
5 FTI report, paragraph 4.24
6 FTI report, para 1.18
7 FTI report, para 1.19

In summary, our report shows:

- The economic framework that underlies the econometric analysis presented in the Frontier report is ill suited to analyse how congestion may affect the airline markets (see *Section 2*). Frontier's analysis fails to account for the following facts:
  - Airlines compete to transport passengers on city-pair markets. This implies that the assessment of airport congestion should focus on airline competition on city-pair markets; and that such an assessment should take into account the extent of congestion in the greater London area (as airlines are able to add seats on city-pair markets at other London airports); and
  - Congestion at Heathrow manifests itself through a shortage of slots rather than through a shortage of seats to transport passengers on city-pairs. Frontier fails to establish how there could be a link between a shortage of slots and a limited supply of seats on city-pair markets.
- To assess the impact of congestion on airlines' operations, we develop a relevant economic framework, which takes into account the fact that slots are scarce inputs used by airlines to fly aircraft on city-pair markets. Our economic analysis shows that congestion is unlikely to cause ticket prices to go up as predicted by Frontier and FTI (see *Section 3*). This is because:
  - Even if Heathrow runs at full capacity in terms of slots, the supply of seats on any city-pair market can be expanded, by increasing the number of passengers per flight, by using or acquiring additional slots to serve that city-pair market or by using slots at different London airports.[8]
  - As congestion becomes more important, the cost of using slots to airlines increases. However, since this cost is fixed, any change is not expected to be passed on to passengers in the form of higher ticket prices.
  - Even if the cost of using slots to airlines has increased, there is no evidence that congestion at Heathrow airport led to an increase in market concentration. While an increase in the cost of slots could theoretically lead to more concentrated city-pair markets, by pushing out airlines from city-pair markets or raising barriers to entry, we do not find evidence of such a trend: concentration on city-pair markets has not increased between 2011 and 2018 (a period during which congestion at Heathrow further increased).
  - In fact, in response to congestion at Heathrow, the data shows that airlines have expanded seat capacity at other London airports. As the cost of slots increases at Heathrow, airlines have increasingly served city-pair markets from other airports, yet they have continued to compete effectively with airlines based at Heathrow.

[8] We have developed two econometric analyses that show that rival airlines that operate at other London airports exert a significant competitive constraint on British Airways' ticket prices at Heathrow. This clearly indicates that competition on city-pair markets involving London should include frequencies in all London airports, and not just Heathrow. (Appendix A and B present these analyses in further detail).

<u>Appendix D further contains a detailed assessment of Frontier's econometric analysis, finding that Frontier's econometric analysis is seriously flawed</u>: As acknowledged by FTI, the results of Frontier's econometric analysis are unreliable. Crucially, unlike FTI, we consider that the empirical approach that Frontier pursues cannot be improved to provide evidence for the existence of scarcity rents, even if it were to adopt the recommendations provided by FTI.

2 The analytical framework proposed by Frontier and FTI is flawed

Capacity constraints, and in our case congestion at airports due to the shortage of available slots, might give rise to scarcity rents. Scarcity rents are the returns earned in excess of normal profits from holding a scarce good (in this case, slots) and "*exist when capacity constraints prevent the market from achieving the equilibrium price and quantity*".[9]

The analytical framework developed by Frontier, and to some extent adopted by FTI, is not adequate to assess how airport congestion affects the ability of airlines to compete on city-pair markets, and in particular in markets involving London. Frontier's analysis fails to establish a clear link between the shortage of slots and passenger ticket prices. As a result, this casts serious doubts on Frontier's prediction that congestion at Heathrow would give rise to a premium on ticket prices.

This section discusses Frontier's analytical framework (Section 2.1) and its key limitations (Section 2.2).

2.1 Frontier's analytical framework

The Frontier report presents a simple analytical framework to assess the impact of airport congestion. According to this analysis, the shortage of slots would inevitably lead to a price increase, as airlines would be unable to expand the supply of seats to transport more passengers.

This framework can be summarized by means of the figure below, which is taken from the FTI report. The x-axis represents quantity, in this case the quantity of seats and the y-axis indicates the cost or price per seat. The demand curve (blue line) shows the relationship between ticket prices and passenger demand for seats. The supply curve (grey full line) is the relationship between seats offered by airlines and the cost of supplying these seats. Capacity is constrained at point C, which can be considered the point at which the cost of supply becomes infinite (as illustrated by the vertical portion of the supply curve, which represents the maximum capacity level).

- When passenger demand is below the maximum capacity level, then the supply curve is flat (in this particular setting). In that case, the "unconstrained" price will be equal to p*.
- When demand exceeds the maximum capacity level, the ticket price rises above the average cost of supply to P_1.

As explained in the FTI report, when demand at the "unconstrained" price exceeds capacity, this "***results in an excess profit, designated by Area A, that would not have been earned by producers in the absence of the capacity constraint. This Area A is the scarcity rent***".[10]

---

9 FTI report, paragraph 3.3

10 FTI report, paragraphs 3.8 and 3.9

Figure 1: Effect of capacity constraint on market equilibrium according to FTI

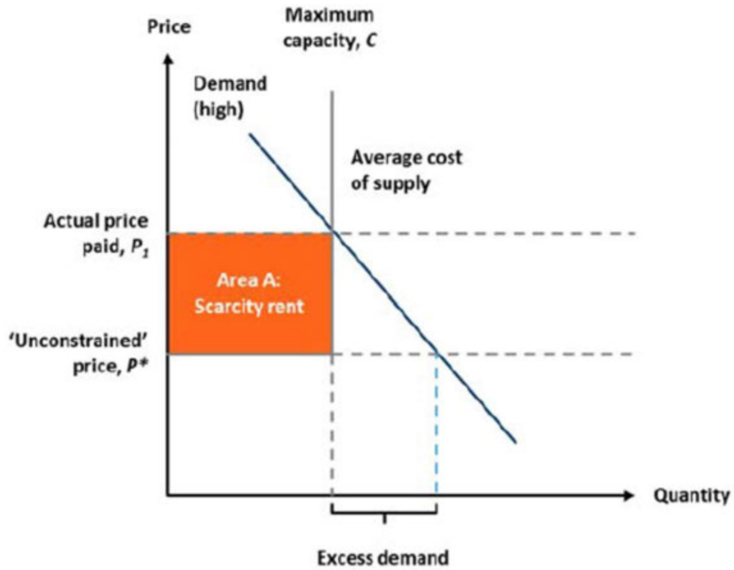

Source: FTI report, figure 3-3.

In summary, the capacity constraint on slots would lead directly to a ticket price increase, which implicitly would be evidence of scarcity rents accruing to airlines.

2.2 Key limitations of Frontier's analytical framework

While the FTI report adopts the economic framework set out by Frontier, it also identifies some of its key limitations, which cast serious doubts on Frontier's prediction that congestion would give rise to a premium on ticket prices at Heathrow.

First, airlines compete to transport passengers in city-pair markets. FTI notes that the "*Frontier report inherently considers Heathrow airport to be a single market*".[11] It further states that "*with numerous airlines operating out of Heathrow, and not all airlines serving the same routes, air travel to and from Heathrow is perhaps more accurately represented as a collection of individual markets*".[12]

This implies that the assessment of airport congestion should focus on airline competition on city-pair markets. This also implies that such an assessment should take into account the extent of congestion in the greater London area, as airlines are able to add seats on city-pair markets at other London airports.

Second, the shortage of available slots does not imply that airlines cannot offer more seats on city-pairs on which they compete. Indeed, the FTI report notes "*passengers purchase seats, whereas the direct constraint in the market is on the slots available to airlines*".[13] In its report,

---

[11] FTI report, paragraph 4.24
[12] Ibid.
[13] FTI report, paragraph 3.12

Frontier fails to establish a clear link between the shortage of slots and a shortage of seats on a given city-pair market, nor does it explain how such a link could work in theory.

As will be shown in Section 3, taking into account that the capacity constraint is on slots and not on seats has important implications on (i) the extent to which airlines' seat capacity on city-pair markets is constrained; and (ii) the extent to which airport congestion gives rise to higher ticket prices to passengers.

[redacted]

3 Heathrow's congestion is unlikely to cause ticket prices to increase

This section presents an alternative economic framework, which shows that the direct impact of congestion is to raise the cost of using slots to airlines.

That slots are a necessary input for airlines to serve passengers on city-pair markets is an obvious starting point. In particular, at slot-controlled airports, airlines have to obtain slots for landings and take-offs in advance. For instance, in city-pair markets involving London, airlines have to use a slot that they may request from slot coordinators at the various London airports.[14]

When slots are in shortage, this may affect airlines' operations; however, as we show below, this would not give rise to a ticket price increase, contrary to what Frontier, and to a lesser extent FTI, predict. There are several reasons why Heathrow congestion is unlikely to cause ticket prices to go up.

- Even if Heathrow runs at full slot capacity, the supply of seats on a given city-pair market can be increased (Section 3.1):
  - by increasing the number of passengers per flight: the evidence suggests a significant increase in the number of passengers per flight through the use of larger aircraft, improved yield management and seat densification of existing aircraft;
  - by using or acquiring additional slots to serve their preferred city-pair markets: airlines can increase seat capacity on a city-pair market by adding frequencies through re-directing a slot in its portfolio or by acquiring a slot from another airline, provided it is profitable to do so;
  - by using slots at different London airports: competition takes place on city-pair markets, which implies that airlines can transport passengers from other London airports. This is further evidenced by an econometric analysis, which shows that rival airlines at other London airports exert a significant competitive constraint on British Airways' ticket prices in Heathrow.
- The increased cost that airlines incur as a result of congestion is not expected to be passed on to passengers in the form of higher ticket prices (Section 3.2):
  - The direct impact of congestion is to raise the cost of using slots to airlines: if slots are widely available (and thus the supply of slots exceeds its demand), no airline will be willing to pay for them. However, as the shortage of slots becomes material (i.e. the supply of slots is not enough to satisfy demand), they become more valuable. As the cost of slots to airlines increases, airlines become more selective on the frequencies they use slots for, choosing only to operate the most profitable frequencies.

[14] Five London airports, City, Gatwick, Heathrow, Luton and Stansted are slot-coordinated airports. London Southend is not.

- – The increase in the cost of slots is unlikely to be passed on to ticket prices: As congestion raises the value of slots, this, in turn, increases the *opportunity cost* of operating existing frequencies. However, slots, if valuable to airlines, represent a fixed cost, and as such, any change in their value is unlikely to be passed on.

- Even if congestion raises the cost of using slots to airlines, there is no evidence that this has led to an increase in concentration on city-pair markets involving Heathrow (Section 3.3):

  – An increase in the cost of slots may push airlines out of some city-pair markets, or deter entry by potential entrants, thus leading to higher levels of market concentration (which may in turn lead to higher prices). However, we do not find evidence of this trend: concentration on city-pair markets has not increased between 2011 and 2018, a period during which congestion at Heathrow further increased.

  – Rather, the data shows that airlines have expanded seat capacity at other London airports. As the cost of slots increases at Heathrow, airlines may increasingly serve city-pair markets from other airports, yet they continue to compete effectively with airlines based at Heathrow.

3.1 Seat capacity may be increased even in the face of slot constraints at Heathrow

Even though Heathrow may be operating at (close to) full slot capacity,[15] seat capacity on individual city-pair markets can be expanded, either by increasing the number of passengers on a flight, by acquiring or re-directing additional slots to serve that city-pair market, or by flying from other airports.

3.1.1 Airlines can increase the number of seats on a flight

Airlines can alleviate the impact of congestion by increasing the number of seats on a flight. This increase is done by up-gauging, that is, by switching to larger aircraft, through improved yield management (e.g. by ensuring that aircraft seat capacity is better utilized), or by densifying seat configuration on existing aircraft. For example, British Airways is engaged in a strategy to replace its A319 with larger A320neos and A321neos. By 2022, the number of A319 aircraft will be reduced from 44 to 22.[16] Similarly, the number of (supersized) Airbus A380 flights being scheduled at Heathrow has increased significantly, from 2,000 in 2011 to 9,000 in 2017. [17]

Further, data from Airport Coordination Limited ("**ACL**") shows that the number of seats per air traffic movement ("**ATM**") has increased by 7% – see the table below.

---

15 The extent to which Heathrow is congested is measured by comparing the annual number of air traffic movements (the number of arrivals and departures) with the available number of air traffic movements (ATMs). For example, the Frontier report (section 3.2) identifies the capacity constraint at Heathrow as total ATMs being close to the available ATMs (in 2016, there were 476,000 ATMs at Heathrow on a total of 480,000 ATMs).

16 IAG Annual report and accounts 2017

17 https://blueswandaily.com/heathrow-hits-new-heights-as-airlines-work-to-maximise-slot-usage/

Table 1: Seats per ATM at Heathrow, Summer 2011 vs Summer 2018

	Summer 2011	Summer 2018	Change (in %)
ATMs	292,121	292,906	+0.3%
Total seats (in million)	58.7	63.2	+7.7%
Seats per ATM	200.9	215.8	+6.9%

Source: ACL; https://www.acl-uk.org/wp-content/uploads/pdf/LHR%20S11%20Start%20of%20Season%20report.pdf for summer 2011; https://www.acl-uk.org/wp-content/uploads/2018/03/LHR-S18-Start-of-season-report.pdf for summer 2018. Only ATMs for passenger flights were counted.

The evolution in number of seats per ATM could also be driven by a shift towards city-pair markets for which airlines already employed larger aircraft. To control for this, we also look at the evolution in the number of passengers per flight for scheduled flights that took place in all years. As shown in the figure below, the average number of passengers on these flights increased by 2-3% between 2011 and 2018.

Figure 2: Passengers per flight departing from Heathrow, 2011-2018

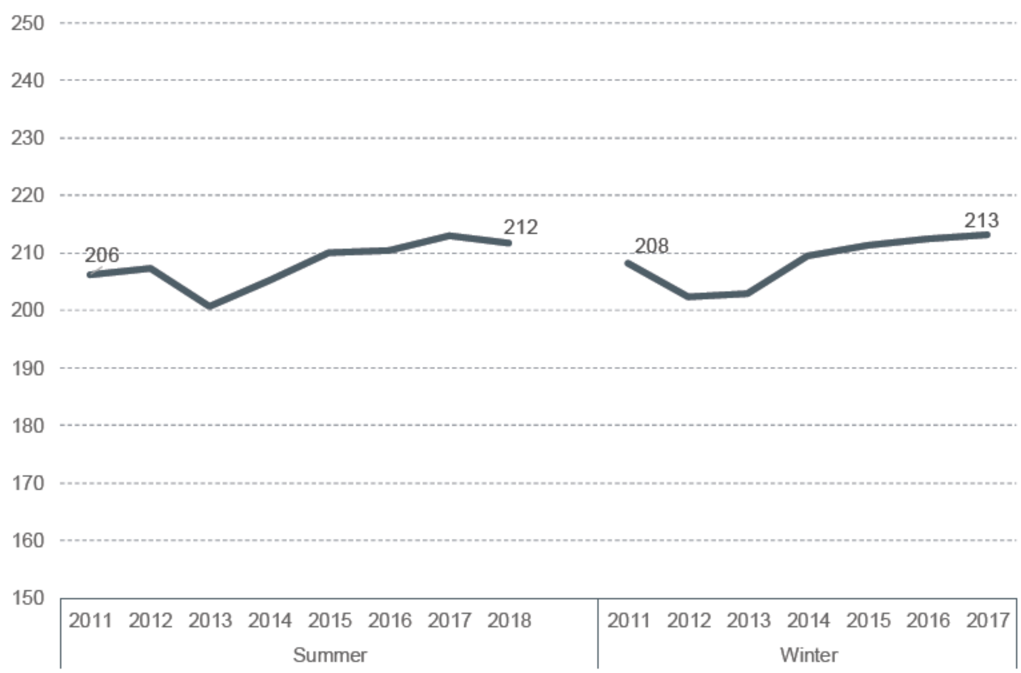

Source: OAG data; Only includes flights that have taken place in every year.

3.1.2 Seat capacity on a city-pair market can be added by re-directing or acquiring slots

Airlines can also increase capacity on a given market by either re-directing slots from another city-pair market served from the same airport, or by trading slots. Slot trading was approved by the UK High Court in a ruling over a slot deal between British Airways and KLM in 1999, and has since taken place on numerous occasions, in particular at Heathrow.[18] In each of the years 2014-2016, more than 200 slots were transferred at Heathrow.[19] British Airways in particular has bought c. 30% of the slots it held in November 2018 (18% if not accounting for the slots acquired in the acquisition of BMI assets).

3.1.3 Capacity can be increased by flying from other London airports

As set out in Section 2.2, airlines compete for passengers on city-pair markets. In the case of London, this means that airlines operating at Heathrow are competing with airlines that fly from other London airports on the same city-pair markets. Consequently, the supply of seats on city-pair markets involving London can be expanded by increasing seat capacity on all London airports, not just Heathrow.

That airlines that fly from different London airports compete with each other is evidenced by two sets of econometric analyses, which show that rival airlines operating from other London airports exert a significant competitive constraint on British Airways ("**BA**") at Heathrow.

<u>The first econometric analysis</u> examines the extent to which entry or exit on a city-pair at another London airport would significantly affect BA's ticket price on flights operated from Heathrow. Specifically, this analysis tests the following hypothesis:

- When competitors enter (exit) a city-pair at another London airport, everything else being equal, this may reduce (increase) BA's ticket price on flights operated from Heathrow.

We develop a *difference-in-difference* regression model that compares the evolution of BA ticket prices on city-pair markets where entry or exit occurred (the treatment markets) against the evolution of BA ticket prices on markets where neither entry nor exit took place (the control markets).

We find that entry (exit) on a city-pair by a competitor at another London airport leads to a statistically significant 7% reduction (increase) in BA's ticket price at Heathrow (corresponding to a £41 price drop (increase) on average).

<u>The second econometric analysis</u> examines how a change in the share of seat capacity held by competing airlines that operate from other London airports affects BA's ticket price at Heathrow on the same city-pair markets. Specifically, this analysis tests the following hypothesis:

- An increase (decrease) in the share of seats held by BA's competitors at other London airports, all else equal, may reduce (increase) BA's ticket price on flights operated from Heathrow.

We find that a 20%-point increase (decrease) in seat share in the hands of all competitors at other London airports leads to a statistically significant 4.5% reduction (increase) in BA's ticket price at Heathrow (corresponding to a £26 price drop (increase) on average);

---

18 House of Commons Airport slots briefing paper, page 6.
19 Ibid, page 7.

The results of the second econometric analysis corroborate the finding that airlines operating from other London airports exert a significant constraint on British Airways at Heathrow, thus suggesting that this finding is robust.

More details on the entry/exit analysis and the seat-capacity share analysis can be found in Appendix A and B respectively.

3.2 The cost of congestion is unlikely to be passed on to passengers in the form of higher ticket prices

In this section, we present a relevant economic framework to assess the impact of congestion on airlines. In particular, this framework shows that a shortage of slots raises the cost of using those for airlines, and that this cost increase is unlikely to be passed on to passengers in the form of higher ticket prices.

If slots are widely available, (and the supply of slots exceeds demand), the cost of slots would be zero since they have no alternative use. However, as slots become scarce (i.e. the supply of slots is not sufficient to satisfy demand), their value increases, as a slot used for one frequency cannot be used to operate another one. Slot scarcity therefore raises the opportunity cost of using them, and naturally, airlines become more selective, choosing to operate only frequencies that earn sufficient profits to justify the increased *opportunity cost*.

The increase in the cost of using slots, which accompanies congestion, is, however, unlikely to be passed on to passengers. This is because, slots, if valuable to airlines, are fixed cost investments. Consider that an airline raises (reduces) ticket price by a small but significant amount on a frequency, it may lose (gain) passengers, and thus save (incur) variable costs. However, the cost of using the slot in question remains the same. This implies that any change in the cost of using slots is not expected to affect how airlines set ticket prices.

3.2.1 Congestion leads to higher cost of using slots

Airport congestion is predicted to have a direct impact on the cost of slots.

- If airport slots are widely available (i.e. the demand for slots is less than the maximum number of slots available), the cost of an additional slot should be zero. This is because the slot that is demanded by an airline has no alternative use; hence, the value of the slot is zero.

- Instead, when there are no slots available at an airport, then in that case, the economic cost of a slot cannot be zero. Indeed, any airline wishing to add a frequency will have to obtain a slot that is already in use (irrespective of whether it re-directs a slot it already owns from another city-pair market or it acquires the slot from another airline). Engaging a slot to run a particular frequency has an opportunity cost, which is determined by the value of the best alternative foregone in making that decision. As a slot can always be sold, the opportunity cost would be at least equal to the expected return on the proceedings from the sale, which can be approximated by the *weighted average cost of capital* ("**WACC**") times the sales price.

It is important to note that airlines will have a strong incentive to use slots optimally. That is, as we show below, airlines may consider reshuffling slots they currently have to optimise their operations, so that they earn enough to compensate the opportunity cost of the slot (more specifically, the frequency would have to earn a return on the potential sales price that is above WACC to justify using the slot for that frequency). Alternatively, they may simply trade slots to other airlines, if these would use the slots more effectively.

To see this, consider the simple example set out in the table below. There are two airports (in different cities), A and B, each with eight slots.

- At Airport A, only five slots are being used, with different levels of profitability, as shown in the table below. For example, the airline using Slot 1 generates a profit of 100 on a slot; while the airline using Slot 5 records a profit of 150. Slots 6 to 8 are not used but available through the slot pool.
- At Airport B, all slots are used.

For the sake of simplicity, we make abstraction of timing differences in slots (e.g. a slot being more valuable because it is a morning slot).[20]

Table 2: Profitability of slots at airports with different degrees of scarcity

	Slot 1	Slot 2	Slot 3	Slot 4	Slot 5	Slot 6	Slot 7	Slot 8
Airport A	100	120	80	90	150	-	-	-
Airport B	100	120	80	90	150	70	110	140

Source: RBB example; it is assumed that Airport A and Airport B are located in different cities.

In this example, the difference in slot availability between the two airports will affect the cost of using slots to airlines.

- Airport A is not congested. There are slots available, which can be requested from the slot coordinator. If an airline would like to start operating in a new city-pair market (or add a frequency to an existing city-pair market), it would be able to do so at zero cost by getting a slot from the pool (i.e. slots 6 to 8). This is because these slots are assumed to have no alternative use. Using them does not deprive another airline of a profit opportunity.
- Airport B is congested as all slots are used. This implies that the use of slots has an opportunity cost. In this context, an airline wishing to add a frequency, would have to acquire a slot from an airline holding a slot (or from its own slot portfolio, which implies sacrificing an existing frequency). The airline would thus forego profits associated with the best alternative investment that could have been pursued instead of paying for the slot in question (or foregoing profits earned by using this slot on another frequency or city-pair market). In this simple example, the acquiring airline will have to pay at least 70 to compensate the airline holding that slot. This is given by the

[20] The analysis also holds when considering that there are slots that are more valuable because their time is more convenient for certain routes.

lowest profit generated by the pool of slots in use at Airport B (in this case Slot 6). The airline using Slot 6 would not sell that slot for less than the profit it generates using it. Hence, 70 is the lowest price at which any of the slot holders would be willing to sell.

Standard economic theory predicts that as the shortage of slots becomes more acute, the cost of using slots will increase. To see this, consider the extension of the simple example presented in the table below, where Slots 6 to 8 are removed from Airport B. Slots are removed in our example to mimic how an increase in scarcity (and hence an increase in excess demand) would affect the value of slots. Before slot trades, this leads to a situation where the previous holders of Slots 6 to 8 are looking to acquire a slot. Among these airlines:

- The airline holding Slot 6 will not attempt to acquire a slot as the maximum price it is willing to pay (70) is lower than the minimum price any of the slot holders is willing to accept.
- The previous holders of Slots 7 and 8 will be willing to acquire slots as the maximum price they are willing to pay is higher than the minimum price any of the slot holders is willing to accept. The airlines that were using Slot 7 and 8 can generate a higher profit than airlines holding Slots 3 and 4. A trade of Slots 3 and 4 will therefore take place, leading to a situation where the five remaining slots are used for the five most profitable flights.

The price from holding a slot (and hence the opportunity cost) in this situation has therefore increased to 100.

Table 3: Profitability of slots at airport B with increased scarcity

	Slot 1	Slot 2	Slot 3	Slot 4	Slot 5	Slot 6	Slot 7	Slot 8
Airport B	100	120	80	90	150	70	110	140
Airport B without Slots 6-8 – before trades	100	120	80	90	150			
Airport B without Slots 6-8 – after trades	100	120	110	140	150			

Source: RBB example

The above example illustrates the shift towards only the most profitable frequencies as the shortage of slots grows (raising the cost of slots). In this example, the elimination of Slots 6 to 8 pushes out the least profitable flights (with profits 70, 80 and 90), as the airlines operating these flights are not able to generate the profits that would justify holding on to the slots.

In reality, airlines operating a hub-and-spoke network may consider not just the profitability of operating a frequency on a city-pair market, but they also consider the profit that may be earned on connecting passengers (which are not part of the city-pair market in question). Indeed, adding or removing a frequency on a city-pair market affects not only local passengers but also transfer passengers. In this context, an airline will assess the opportunity cost of adding a frequency by considering the impact not only on the city-pair market in question but also on the connecting traffic.

Contrary to what the FTI report claims, the mechanism described above also applies when network effects are taken in consideration, as well as when there are restrictions on city-pair market access on the basis of Air Services Agreements ("**ASA**"). [21] When taking a decision on which frequency to allocate a slot to, airlines take into account the impact of that choice on overall profitability, i.e. considering the impact of operating a frequency on their performance on other (complementary) city-pair markets. For example, an airline with a hub at another airport may be able to profitably operate a frequency from Heathrow to their hub which offers good domestic or regional connections. Airlines will choose to operate the mix of frequencies, which yields the highest profit, taking into account potential ASA restrictions.[22]

Further, the mechanism by which slot congestion pushes airlines towards operating more profitable frequencies works under the assumption that the airport is slot congested. Contrary to what the FTI report claims,[23] there is no obvious connection between airport congestion and a restriction on the supply of seats on city-pair markets. As explained above, capacity on a city-pair market can always be increased by increasing the seats per flight; by acquiring or re-directing slots; or by using slots at other airports. Slot congestion does not lead to capacity constraints at the city-pair level; rather it raises the cost of using slots. Although this might lead airlines to abandon some frequencies (as the slots may be more profitably used elsewhere), and possibly withdrawn from some city-pair markets, the shortage of slots does not automatically reduce the supply of seats on city-pair markets.

3.2.2 The increased cost of using slots is not directly passed on to passengers

If the cost of an input (in this case the cost of slots) is increased, this cost increase could potentially be passed on to customers by the firm that experiences this change (in this case, airlines that could raise the ticket price to passengers).

In order to assess the extent to which an increase in the cost of slots is being passed on, we examine two scenarios.

First, consider a scenario in which an airline acquires a slot from another airline to operate a new frequency at a congested airport. To do so, the airline in question will have to compensate the airline holding the slot (as per our simple example above). This compensation represents the cost paid for the slot. That acquisition cost, however, constitutes an upfront, fixed investment cost. In other words, irrespective of the number of passengers that the airline in question will eventually transport on this new frequency, the acquisition cost of the slot will not change. As that expense will be incurred before operating the new frequency and will be the same irrespective of the number of passengers, this fixed cost is not expected to determine prices.[24] In principle, only changes in the cost that vary with the number of passengers

---

21 FTI report, paragraphs 6.10-6.13.

22 The FTI report makes the point that in some cases capacity constraints cannot be alleviated by other airlines stepping in, mentioning the example of the London to Cairo route. However, in this particular case, the capacity constraint does not emanate from slot congestion, but from regulation (only airlines designated by the UK can operate up to seven services per week on the London – Cairo route), which renders this example irrelevant for the purpose of the discussion on the nature of scarcity rents.

23 FTI report, paragraph 6.14 and following.

24 Clearly, airlines will pay this cost only if they expect that they would earn sufficient revenue to cover that expense (and earn a return above WACC). This means that they will ensure that sufficient passengers are transported using the slot.

transported are expected to affect prices.[25] This is also acknowledged by the CMA's latest report on airport slot allocation, stating that "*it is […] not clear that airlines would pass on the costs [*of higher prices for slots at constrained airports*], because […] if slot payments were required to be paid upfront, [they] would represent a fixed (sunk) cost rather than a variable cost*".[26] In summary, even if an airline must pay for slots, this should not directly affect ticket prices.

Second, consider a scenario in which the airline already holds and uses slots, but congestion at the airport in question has grown such that slots are no longer available for free. In this case, the use of slots gives rise to an opportunity cost. By using these slots whilst they could be employed elsewhere more profitably, the airline foregoes revenue either for its own use on another city-pair market, or from a potential trade. As congestion increases, the opportunity cost may increase. However, this too represents a fixed cost such that it is unlikely to alter an airline's pricing decision.

3.3 Slot scarcity does not significantly increase market concentration on city-pair markets

While airport congestion is unlikely to have a direct impact on ticket prices, it may be argued that it may have an indirect impact, through an increase in market concentration.

- First, the increase in the cost of using slot should lead airlines to be more selective and shift their operations to the most profitable frequencies. This move may push some airlines out of a city-pair market (as they re-direct a slot toward another frequency; or sell their slot to other airlines).
- Second, the increase in slot prices may also raise barriers to entry, as potential entrants on a city-pair market will face a higher cost of entry.

These two effects could lead to higher concentration levels, and hence possibly to higher prices (although the relationship between market concentration and price levels is empirical).

Consider our simple example (the table below reproduced the example), in which the airline holding Slot 1 and the airline holding Slot 4 at Airport B before elimination of Slots 6 to 8 are competing on the same city-pair market. The elimination of Slots 6 to 8 will push the airline previously holding Slot 4 out of the market; the slot price will be too high to justify holding on to the slot. This may enable the airline holding Slot 1 to increase prices on the city-pair market.

Table 4: Profitability of slots at airport B with increased scarcity

	Slot 1	Slot 2	Slot 3	Slot 4	Slot 5	Slot 6	Slot 7	Slot 8
Airport B	100	120	80	90	150	70	110	140

[25] This is because when a profit maximising firm raises prices, its output declines, yet this will not result in any fixed cost saving. Fixed costs do not vary with output. That implies that fixed costs are not taken into account by firms when they set prices to maximise profit. For more details, see RBB report for the Office of Fair Trading on Cost pass-through - https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/320912/Cost_Pass-Through_Report.pdf

[26] Advice for DfT on competition impacts of airport slot allocation, CMA, paragraph 5.

	Slot 1	Slot 2	Slot 3	Slot 4	Slot 5	Slot 6	Slot 7	Slot 8
Airport B without Slots 6-8 – before trades	100	120	80	90	150			
Airport B without Slots 6-8 – after trades	100	120	110	140	150			

Source: RBB example

However, this reasoning may be too simplistic. Barriers to entry on city-pair markets are likely to remain low, even in the face of slot congestion at Heathrow, because potential entrants can always choose to enter a city-pair market via another London airport.

As such, while airport congestion could theoretically increase market concentration, the evidence suggests that congestion at Heathrow has not led to a significant change in market concentration. Analysis of data covering 2011 and 2018, a period during which congestion at Heathrow increased, does not reveal a significant increase in market concentration (Section 3.3.1). Rather than pushing airlines out of a city-pair market, we observe that competition has shifted towards other airports (see Section 3.3.2). This suggests that the increased congestion did not lead to the reduced levels of competition that are a prerequisite to observe the congestion premium claimed by Frontier.

3.3.1 Heathrow airport congestion did not lead to greater market concentration

We show below that increased congestion at Heathrow has not led to greater market concentration.

We examine the link between congestion and market concentration, by evaluating notably the change in concentration on city-pair markets involving London airports between 2011 and 2018. If it were the case that airport congestion has pushed airlines out of city-pair markets, then we should observe a marked increase in concentration between 2011 and 2018, a period where congestion further increased as set out in the FTI report.

We measure concentration using two standard metrics:

- the number of competitors on a city-pair route; and
- the Herfindahl-Hirschman Index ("**HHI**").[27]

The analysis shows that market concentration has not increased on city-pair markets served from Heathrow.

First, we calculate the average number of competitors and HHI over time on all city-pair markets to and from London-based airports, distinguishing between winter and summer – shown in the figure below. We do not include new city-pair markets since 2011 or city-pair markets that were discontinued since 2011, to allow comparison over time for a stable set of city-pair markets.

---

[27] HHI in this case is calculated as the sum of squared seat shares of each competitor on a given route.

Figure 3: Evolution in average concentration metrics, 2011 to 2018

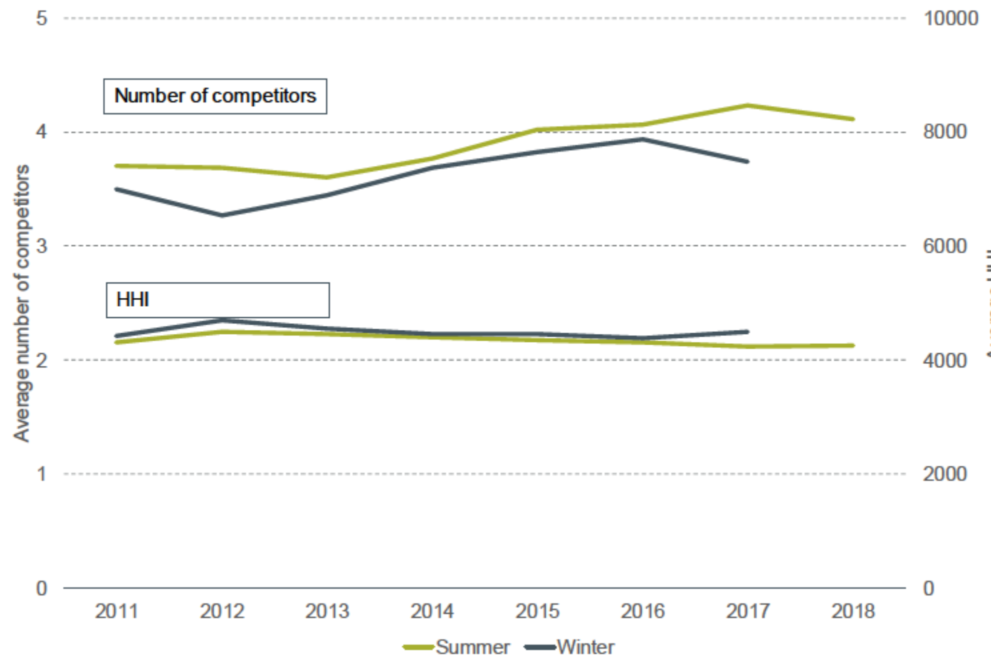

Source: OAG data; Mergers, joint ventures and joint business agreements on city-pair markets during 2011 and 2018 are assumed to have already taken place prior to 2011, to filter out their impact on HHI and the number of competitors.

This evidence indicates that, on average, market concentration has not increased between 2011 and 2018.

- The average number of competitors per city-pair market has gone up, from 3.7 in the summer of 2011 to 4.1 in the summer of 2018, suggesting that markets have become less concentrated.
- Similarly, the HHI has slightly gone down, from 4,320 in summer 2011 to 4,250 in summer 2018, suggesting again markets have become (marginally) less concentrated.

Further, we have also assessed the share of city-pair markets on which market concentration was likely to be significantly affected between 2011 and 2018, measured as the proportion of city-pair markets (weighted by number of seats) that evolved to a situation of duopoly or monopoly – shown in the figure below.

Figure 4: Proportion of city-pair markets (weighted by seats) by evolution in number of competitors per route, 2011 vs 2017/2018

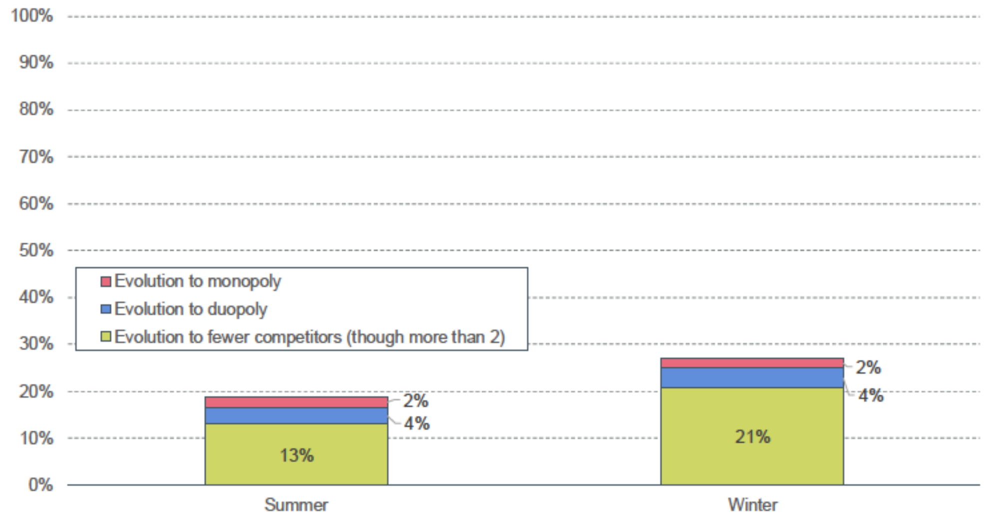

Source: OAG data

Only a minority of city-pair markets experienced a reduction in the number of competitors over time (i.e. for the vast majority of markets, the number of competitors increased or stayed constant). Further, only a very small share (c. 6%) of city-pair markets saw an evolution towards monopoly or duopoly, suggesting that the share of city-pair markets where market concentration was significantly affected over this period is minimal.

3.3.2 Airlines shifted flights towards other London airports in response to congestion at Heathrow

Rather than raising market concentration, the data show that increased congestion at Heathrow has in fact led to a shift of seat capacity to other London airports.

First, the average number of airlines operating from Heathrow per city-pair has fallen between 2011 and 2018, while the average number of airlines operating from other London airports per city-pair has gone up over the same period, as evidenced in the figure below. This suggests that on a given city-pair, airlines have tended to shift operations towards other London airports or that airlines operating from other London airports have replaced airlines operating from Heathrow.

Figure 5: Average number of operators per city-pair per airport, 2011-2018, summer routes only

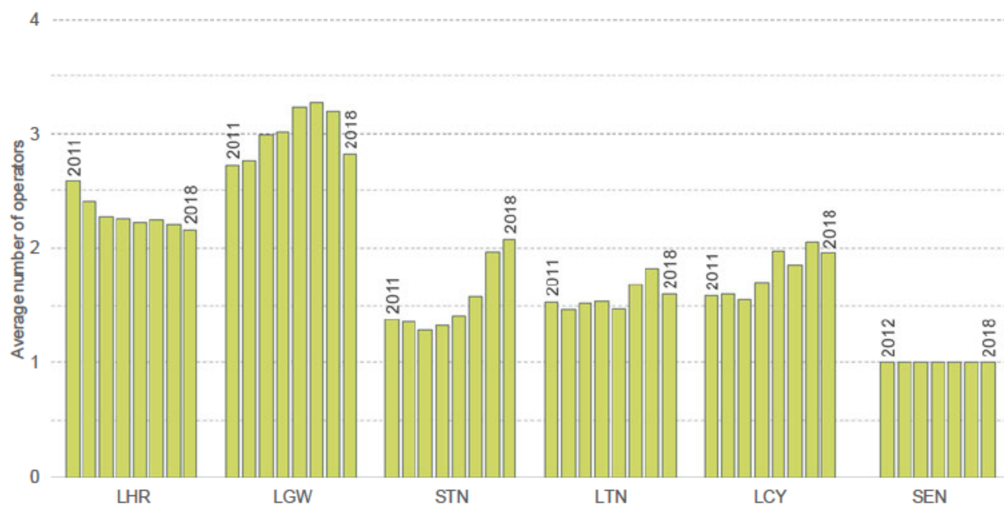

Source: OAG data; weighted by total seat capacity on a city-pair market. LHR = Heathrow; LGW = Gatwick; STN = Stansted; LTN = Luton; LCY = London City; SEN = London Southend.

Second, Heathrow's share in total seats offered by all London airports went down significantly as evidenced in the figure below – from 55% in 2011, to 45% in 2018.

Figure 6: Seat share by airport, Summer 2011 to Summer 2018

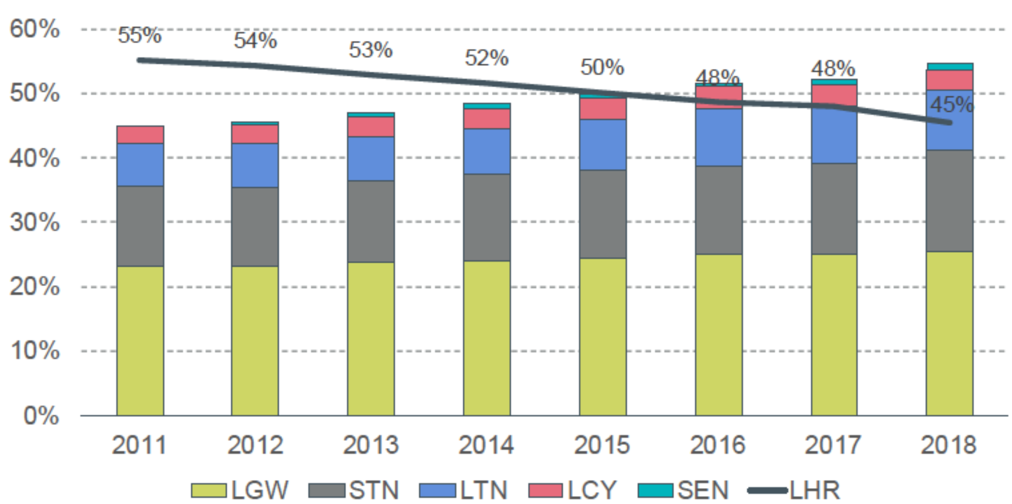

Source: OAG data. LHR = Heathrow; LGW = Gatwick; STN = Stansted; LTN = Luton; LCY = London City.

The decline in the share of seats held by Heathrow may be due to the fact that some city-pair markets predominantly operated from Heathrow have become less popular than city-pair markets operated from other airports. To control for this, we also assess the change in the proportion of seats held by Heathrow at the city-pair market level. The chart below indicates the proportion of city-pair markets for which an airport's share of seats has declined ("Lost share"), stayed the same ("Constant share"), or gone up ("Gained share"), as well as the proportion of city-pair markets that were eliminated at an airport ("Eliminated") or newly created ("New"). For example, the chart shows that on 61% of city-pair markets (weighted by seats) served from Heathrow, Heathrow's proportion of seats has gone down. The chart also shows that on 41% of the city-pair markets at Stansted, that airport's proportion of seats has gone up.

Figure 7: Proportion of city-pairs (seat-weighted) by evolution in airport seat share, 2011 vs 2018, Summer season

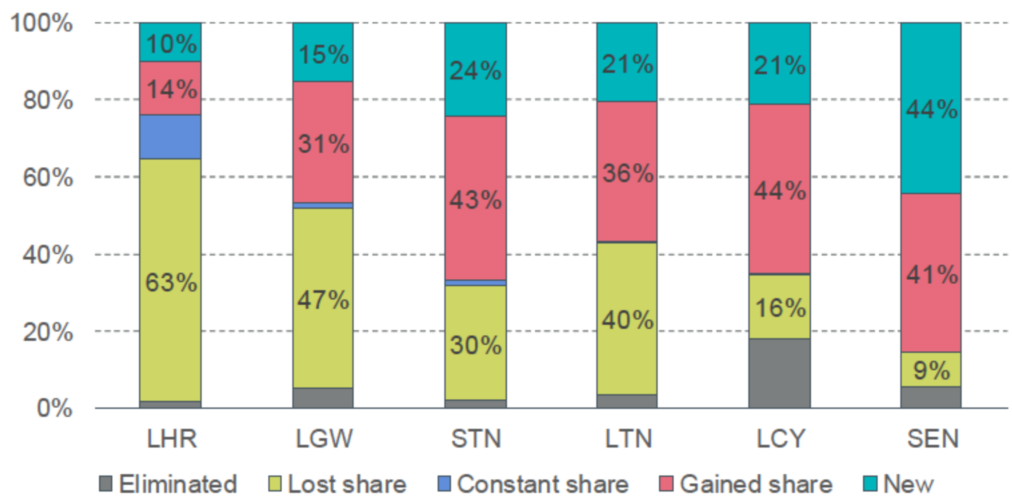

Source: OAG data. The weights are determined on the basis of the share of total London seats served by that route over the period 2011-2018. For Soutend, evolution between 2018 and 2012 is compared, as no OAG data are available for 2011 for Southend.

The figure shows that:

- Heathrow's proportion of "lost share" city-pairs is significantly higher than the corresponding proportion for the other airports;
- These other airports have gained share for a significantly higher proportion of city-pairs, effectively from Heathrow.

The figure thus confirms that airlines operating from other London airports have gained seat share at the expense of airlines operating from Heathrow, most likely in response to increased congestion at Heathrow.

Annexes

A Entry/Exit analysis

This econometric analysis quantifies the impact of entry or exit of competitor airlines at other London airports on British Airways plc's ("**BA**") ticket prices at Heathrow. The results of this analysis clearly show that airlines operating in the same city-pair markets but from other London airports apply a competitive pressure on ticket prices at Heathrow airport.

In the rest of this annex, we present in more details this analysis, which relies on BA data as well as capacity data from OAG.

A.1 The impact of entry/exit on British Airways' price

This analysis assesses empirically the impact of entry/exit on BA's ticket prices at London Heathrow. In particular, we examine the extent to which ticket prices have declined (increased) following entry (exit) of a competing airline at another London airport.

The figure below provides an example of the evolution of BA ticket prices on a selected city-pair market – in this case the city-pair market London-Luxemburg where Ryanair entered from Stansted in November 2016. There was a noticeable drop in average ticket price after Ryanair entered, for both ticket classes shown (Business High and Economy Medium). Our analysis aims to determine whether this price drop can be attributed to Ryanair's entry, or whether it is due to other factors that coincided with Ryanair's entry.

Figure 8: Example of BA ticket prices on city-pair market with entry served from Heathrow – London-Luxemburg

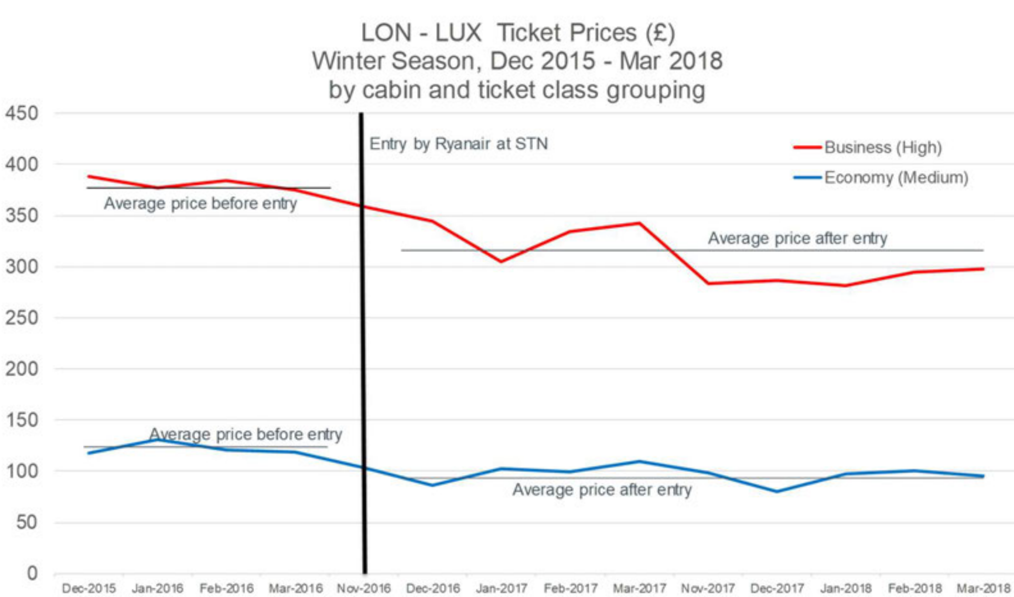

Source: British Airways ticket price database; OAG for the identification of entry and exit routes. Price given is the price for a return flight in the particular class type given.

A.2 The sample of relevant city-pairs

This analysis focuses on the impact of entry or exit of airlines on British Airways' ticket prices. To this end, we select a sample of 29 city-pairs from the OAG data in which British Airways operated from Heathrow every month between April 2011 and July 2018, but for which competitors entered or exited at other London airports between April 2011 and July 2018.

- To focus on important events of entry and exit, we consider only "permanent" entry or exit, which means that the airline that entered, stayed in the market until at least July 2018; and the airline that exited, stayed out of the market until at least July 2018.
- We also consider markets by season, i.e. winter or summer of the same city-pair constitute a different city-pair market.

The entry and exit events that are included in this study represent a significant share of seat capacity on the city-pair. An entrant at other London airports represented on average 18% of total seat capacity on a city-pair; an exit by an airline at other London airports accounted for 21% of total seat capacity on a city-pair. A list of all city-pairs considered is presented in Annex C below.

The ticket prices are based on revenue from point-to-point passengers (specifically passengers that have purchased a return ticket) on these city-pairs. This means that the price is based on so-called "local passengers", i.e. those passengers for which the origin and destination is the city-pair, and not "connecting passengers", i.e. those whose origin and destination do not correspond to the city-pair in question.

In addition, we are only using ticket prices of direct flights.

Further, we distinguish between different ticket class types, i.e. between Economy, Premium Economy, Business and First class, with further subdivisions by ticket class (High, Medium, Low and Groups).

Finally, price used is the net price paid by the passenger in GBP where the exchange rate has been kept fixed.

A.3 Estimation methodology

To quantify the impact of entry and exit at other London airports on BA's ticket price, it is important to control for the potential confounding influence of other factors that can affect ticket prices and that might coincide with the entry or exit event.

To "isolate" the impact of entry and exit, we apply a standard empirical approach – a so-called *differences-in-differences* analysis. Briefly, this approach compares the changes in BA's ticket prices following entry or exit (the *treatment markets*) against changes in BA's ticket prices on city-pairs where there was no entry or exit (the *control markets*).

Figure 9 provides an illustration of this approach, specifically for the case of exit (the approach is identical for entry). The approach consists of comparing the evolution of ticket prices on the treatment city-pair, i.e. one with exit by a competitor operating on a non-Heathrow airport, with that in the control city-pair, i.e., one without exit.

The advantage of this approach is that if ticket price exhibits a trend, this will not affect the analysis. Consider, for example, that the ticket prices in both groups rise over time, irrespective of whether exit took place on the city-pair. This complicates the analysis of the impact of exit on ticket prices for the city-pair: is any increase in ticket price due to the treatment or would this have occurred in any event? A differences-in-differences analysis can address this issue: if the difference in ticket price between the treatment group (i.e. city-pairs with exit) and the control group (i.e. city-pairs without exit) has increased, the exit is likely to have had an effect.

Figure 9: Illustration of difference-in-differences analysis

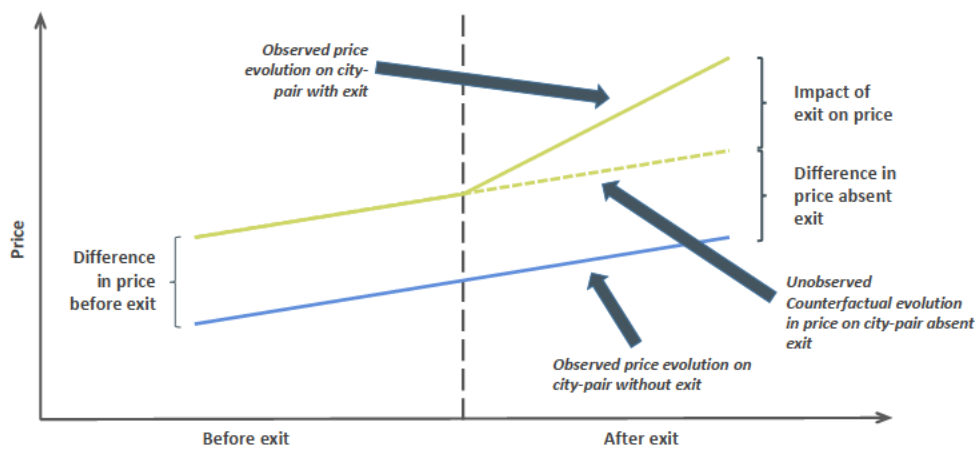

Source: Adapted from https://www.mailman.columbia.edu/research/population-health-methods/difference-difference-estimation

A differences-in-differences analysis assumes ticket prices in the treatment market would have evolved in the same way as prices in the control markets absent entry/exit (*parallel trend assumption*). This assumption, however, is not always realistic. It is, therefore, important to control for other factors that drive how ticket prices in the treatment and the control markets may be different.

The regression model is given below:

$$\ln(price) = [\alpha + \beta_1 Distance + \beta_2 FareClassDummies\ ] + [\beta_3 TimeDummies] + [\beta_4 EntryRoute + \beta_5 ExitRoute] + [\beta_6 EntryTimeDummy + \beta_7 ExitTimeDummy].$$

We can explain the regression model as follows:

- The part in the first set of square brackets controls for the impact of distance and fare class dummies, which may affect the price of the treatment and control city-pairs markets.
- In the second set of square brackets, the time dummies account for the change in ticket price over time in both the treatment and control city-pair markets.[28] This corresponds to the parallel trend assumption as the evolution of ticket price absent treatment is assumed to be the same after controlling for differences in distance and fare class (in Figure 9 this corresponds to the trend given by the blue line).
- The third set of square brackets gives the estimated average difference in ticket price between the treatment and control city-pair markets (separately for entry and exit). This is the difference between the blue and green line before exit in the above figure.
- The last set of square brackets gives the difference-in-differences estimate (separately for entry and exit). It indicates by how much entry (exit) has affected the difference in ticket price between the treatment and control city-pair markets. In Figure 9 above, this corresponds to the difference between the solid green and dashed green lines.

The above regression model may give rise to a potential endogeneity problem, as entry and exit may be correlated with the error term due to reverse causality. Indeed, it might be argued that high prices cause entry of new airlines and low prices causes exit. However, this endogeneity would lead to understate the true impact of entry and exit on British Airways ticket prices.

- Entry: Entry is expected to have a negative impact on prices, whilst higher prices are expected to encourage entry. Failing to account for the latter link, that is, that high prices may trigger entry, may understate the negative impact of entry on prices. This failure means that the regression results confound the two causal relation, which in this particular case may give rise to under-estimate significantly the negative effect of entry on prices.
- Exit: Contrary to entry, exit is expected to have a positive impact on prices, whilst low prices are expected to induce airlines to leave the market. As for entry, failing to account for this reverse causal link may lead to understate significantly the effect of exit on prices.

Since, in practice it is difficult to identify relevant instrumental variables, notably for dummy variables, we present results that may be biased in the sense that they would understate the effect of entry and exit on BA ticket prices. As the results indicate, even if downward biased, this analysis supports the thesis that airlines that operate in other London airport exert a significant competitive constraint on British Airways' operations at Heathrow. If anything, the actual impact might be larger than what we estimate.

---

28 In a standard difference-in-differences, one time dummy variable would be added (i.e. a variable with value 0 before the date of entry or exit and a value 1 after the date of entry or exit), which would pick up the change in ticket price after the date of entry or exit on the control routes. In this case, the various entry and exit episodes considered do not occur at the same time – they are *staggered*. Instead of adding a time dummy variable per episode, we therefore include time fixed-effects (i.e. a dummy for each month considered) which pick up variations in ticket prices in any given month.

A.4 Econometric results

The table below sets out the detailed regression results. The first column introduces the results of our base model, with separate dummies for entry and exit to account for a differential impact (in absolute terms) of these two events. The second column presents a restricted model where we assume the impact of entry and exit to be the same (in absolute terms). The next two columns estimate separate models for entry and exit respectively.

Table 5: Regression results for various models estimated

Variable	**(1)**	**(3)**	**(5)**	**(7)**
Model	*Base*	*Restricted*	*Only entry*	*Only exit*
Impact of entry on BA ticket price	**-5.4%**	**-6.7%**	**-6.6%**	
Impact of exit on BA ticket price	**10.5%**	**7.2%**		**12.7%**
Distance	0.666***	0.664***	0.662***	0.651***
EntryRoute	0.120***	0.123***	0.126***	
ExitRoute	-0.094***	-0.085***		-0.129***
EventTimeDummy		-0.070***		
EntryTimeDummy	-0.056***		-0.069***	
ExitTimeDummy	0.010***			0.120***
Observations	17553	17553	15943	13744
R-squared	0.951	0.951	0.956	0.946

Source: RBB analysis. Data based on BA ticket prices and OAG data for the city-pair markets. Only city-pair markets with fewer than five competitors have been included. Only route-fare class combinations have been included with observations across the full time period. Standard errors are clustered at the city-pair markets, date and ticket class level. Impact of entry/exit on BA ticket price has been calculated by calculating exp(X)-1 with X the coefficient for EntryTimeDummy and ExitTimeDummy respectively. Observations have been weighted by the number of passengers per city-pair markets per month per fare class.

The results of all regression models presented above find that entry/exit have a significant impact on ticket price, and in the expected direction (i.e. negative for entry; positive for exit). In particular, the "restricted" model shows that following an entry (exit) of a competitor airline at other London airports, BA's ticket prices at Heathrow on the same city-pair were reduced (increased) by at least 7%. The other regression models find that entry depresses ticket prices by 5-7%, whilst exit leads to higher prices by 10-13%.

Importantly, the difference in the impact of entry and exit on ticket prices should not be interpreted as indicating that entry has had less impact on the market than exit. This difference may simply stem from the fact that the entry and exit events considered in this analysis are not fully comparable.

- First, the exit events included in the sample are typically larger, and hence more impactful, than the entry events: airlines exiting the city-pairs had a seat share of 21% two months before exiting, whereas the airlines entering the market had a seat share of 18% two months after entry.

- Second, exit events took place on markets that were slightly less concentrated than markets on which entry events took place. City-pairs on which exit took place had on average 3.3 competitors before exit; city-pairs on which entry took place had on average 3.6 competitors after entry. We would expect the impact of exit to be larger therefore.

B Impact of change in share of seat capacity on BA ticket prices

In complement, we develop another econometric analysis to assess the extent to which competition from other London airports affect BA's ticket prices at Heathrow. This alternative approach examines empirically how a change in the share of seat capacity held by competing airlines that fly from other London airports on the same city-pairs affects BA's ticket price at Heathrow. The results of this econometric analysis corroborate the finding that airlines operating from other London airports exert a significant constraint on British Airways at Heathrow, thus suggesting that this finding is robust.

This analysis is based on a sample of city-pair markets, which has both BA operating from Heathrow as well as at least one competing airline operating from another London airport. We only consider city-pair markets in which fewer than five competitors were operating at some point in time.

This analysis seeks to quantify the impact of a change in the share of seats held by rival airlines at other London airports. Specifically, we test the following hypothesis: if competing airlines operating at other London airports exert a competitive constraint on British Airways, we would expect that an increase in the share of seat capacity held by competitors on the same city-pair would depress BA's ticket prices at Heathrow.

The econometric analysis is based on a *fixed-effects panel regression model*, where we control for price differences due to fare class, destination region of the city-pair market, and the month in which the flight took place (to adjust for seasonality). We estimate the following regression model:

$$\begin{aligned}\ln(price) = [\alpha + \beta_1 Distance + \beta_2 FareClassDummies + \beta_3 Month\&RegionDummies] \\ + \beta_4 SeatShareOfNonLHRAirports\end{aligned}$$

Our hypothesis can be tested by examining the coefficient estimate of the variable *SeatShareOfNonLHRAirports*, which is the share of seats held by competitors at other London airports.

Note that the share of seat capacity held by competitors at other London airports may be endogenous. This is because high BA prices may cause competing airlines to expand capacity, notably at other London airports. This reverse causality may lead to biased estimation results; however, in this case the effect would be understated.

To alleviate the endogeneity bias, we use an instrumental variable estimator. To this end, we use as instrumental variables the annual average seat share in the winter months for summer city-pair markets and the annual average seat share in the summer months for winter city-pair markets. The logic is that the annual average seat share in the winter season is likely correlated with the seat share of the same city-pair markets in summer, but it is not affected by price levels in summer (as these can be considered different markets).

The table below presents the two sets of results.

Table 6: Regression results for share seat capacity model

Variable	**(1)**	**(2)**
Model	*Base*	*IV*
Impact of 20 p.p. increase in non-LHR seat share	**-4.5%**	**-5.1%**
Distance	0.514***	0.509***
SeatShareOfNonLHRAirports	-0.002***	-0.003***
Observations	1,327	1,327
R-squared	0.96	0.96

Source: RBB analysis. Data based on BA ticket prices and OAG data for the routes. Standard errors are clustered at the route, month and ticket class level. Only includes city-pair markets with non-zero seat share at other airports. Only includes city-pair markets with fewer than 5 competitors. Observations have been weighted by the number of passengers per route per month per fare class.

We find that a change in seat share at other London airports has a significant negative impact on BA's ticket price: as competition from other London airports becomes more intense on a route, we observe that BA's ticket price on this route drops. A 20 percentage points change in the share of seat capacity held by all competitors at other London airports, would have a negative 4.5% effect on BA's Heathrow ticket price.

C List of city-pair markets with entry and exit

This annex presents the list of selected city-pair markets with episodes of permanent entry or exit. It should be clear the permanence of entry or exit is assessed over the sample period. For instance, entry is considered permanent, when a new airline starts operating on a city-pair and has not exited during the sample period.

Only city-pair markets with fewer than five competitors have been considered for the analysis.

Some city-pairs experienced multiple entry or exit episodes. We address this by splitting up the markets into several different sub-markets and estimating the impact of entry and exit on those. For example, the winter seasons on the London-Luxemburg city-pair markets had multiple episodes of permanent entry and exit – in November 2012, easyJet permanently entered the city-pair markets via Gatwick; in March 2014, Air France permanently exited the market via London City; and in November 2016, Ryanair permanently entered via Stansted. The London-Luxembourg winter city-pair market has therefore been split up into three different sub-markets. Each particular sub-market is related to an exit or entry event and is assigned the months closest to that event, as illustrated in the table below for the London-Luxembourg city-pair markets.

Table 7: An example of splitting the market that has multiple events – London-Luxembourg

Sub-route (event)	Event date	Time period of sub-route
LUX1 (entry by easyJet)	November 2012	November 2011 – November 2013
LUX2 (exit by Air France)	March 2014	December 2013 – November 2015
LUX3 (entry by Ryanair)	November 2016	December 2015 – March 2018

Source: RBB econometric analysis

The tables below present the city-pair markets with entry and exit which were used for the econometric analysis.

Table 8: City-pair markets with entry at non-Heathrow London airport

Route code	Route	Season	Airline	Date of entry
BCN1	Stansted-Barcelona	Summer	Ryanair	Apr-12
BCN1	Stansted-Barcelona	Winter	Ryanair	Dec-11
BOS1	Gatwick-Boston	Summer	Norwegian	Apr-16
BOS1	Gatwick-Boston	Winter	Norwegian	Mar-16
CHI1	Gatwick-Chicago	Summer	Norwegian	Apr-18
LAS1	Gatwick-Las Vegas	Winter	Norwegian	Nov-16
LED2	Gatwick-Saint Petersburg	Summer	Aeroflot	Apr-18
LUX1	Gatwick-Luxembourg	Summer	easyJet	Oct-12

Route code	Route	Season	Airline	Date of entry
LUX1	Gatwick-Luxembourg	Winter	easyJet	Nov-12
LUX3	Stansted-Luxembourg	Summer	Ryanair	Oct-16
LUX3	Stansted-Luxembourg	Winter	Ryanair	Nov-16
SEA2	Gatwick-Seattle	Summer	Norwegian	Sep-17
SEA2	Gatwick-Seattle	Winter	Norwegian	Nov-17
STR2	Gatwick-Stuttgart	Summer	easyJet	Apr-15
STR2	Gatwick-Stuttgart	Winter	easyJet	Mar-15
WAS1	Stansted-Washington	Summer	Primera Air	Aug-18
YYC1	Gatwick-Calgary	Winter	WestJet Airlines	Nov-16
YYC2	Gatwick-Calgary	Summer	WestJet Airlines	May-16

Source: RBB analysis; Shaded in green are those routes that have been sampled for the base regression analysis.

Table 9: City-pair markets with exit at non-Heathrow London airport

Route code	Route	Season	Airline	Date of exit
ACC1	Gatwick-Accra	Summer	Ghana International Airlines	Jun-12
ACC1	Gatwick-Accra	Winter	Ghana International Airlines	Mar-12
GIB1	Luton-Gibraltar	Summer	Monarch Airlines	Oct-17
GIB1	Luton-Gibraltar	Winter	Monarch Airlines	Mar-17
GOT1	Gatwick-Goteborg	Winter	easyJet	Jan-12
LED1	Gatwick-Saint Petersburg	Summer	Rossiya Airlines	Oct-12
LED1	Gatwick-Saint Petersburg	Winter	Rossiya Airlines	Mar-12
LUX2	City-Luxembourg	Summer	Air France	Oct-13
LUX2	City-Luxembourg	Winter	Air France	Mar-14
STR3	Stansted-Stuttgart	Winter	Germanwings	Oct-17
YYC1	Gatwick-Calgary	Winter	Thomas Cook	Oct-11

Source: RBB analysis; Shaded in green are those routes that have been sampled for the base regression analysis.

[redacted]

D Frontier's econometric analysis suffers from serious methodological flaws, rendering its conclusions unreliable

To estimate scarcity rent at London Heathrow, Frontier develops an econometric analysis. This analysis essentially compares average annual ticket prices (in 2016) for departing city-pair routes for Heathrow and other London-based airports, and for Heathrow and other European hub airports. The Frontier report finds that ticket prices at Heathrow were on average 23% higher than at other London airports and 24% higher than at other European hub airports. Frontier explicitly attributes the higher price at Heathrow to congestion: "*because Heathrow is capacity constrained [*as opposed to the other airports*], passengers face higher ticket prices compared to both London and other European hub airports*".[29]

At the outset we note that in Section 3 above we show that the theoretical foundation for scarcity rents being passed on to ticket prices is weak. In fact, as indicated above, Frontier has not explained how airport congestion would limit the supply of passenger seats on city-pair markets served by Heathrow and other London airports. This means that the ticket price premium at London Heathrow as quantified by Frontier's regression analysis may have little to do with congestion.

In any event, as we explain below, the results of the Frontier econometric analysis are not reliable. More importantly, and unlike FTI, we consider that the empirical approach that Frontier has pursued cannot provide evidence for the existence of scarcity rents, even if it were to adopt the recommendations provided by FTI.

D.1 The Frontier regression analysis

The Frontier report develops an econometric analysis to quantify the impact of congestion at London Heathrow on ticket prices. In a nutshell, the regression model compares average ticket prices in 2016 using a sample of flights from the different London airports, namely Heathrow, Gatwick, Luton, Stansted and City airports.[30]

Since prices may differ between flights because of different factors, the regression analysis controls for distance, whether this is a long-haul flight, the number of flights (frequency) and other variables. Importantly, the regression model includes a Heathrow dummy variable, which indicates by how much on average ticket prices of flights from Heathrow differ from ticket prices at other airports. For the comparison with other London airports, the coefficient estimate of this variable is positive at 0.233, and statistically significant. Using this result, Frontier concludes that its analysis shows a congestion premium of 23.3% (0.233 x 100%), which implies a one-way mark up of circa £59.[31]

---

29 Frontier report, page 28.

30 The Frontier report presents another, similar analysis that compares ticket prices at Heathrow with that at other major European airports.

31 Frontier's interpretation of the coefficient estimate is incorrect. The dependent variable is the logarithm of the ticket price of inbound flights. In this case, to estimate the impact of a binary variable, the percentage impact is 100 (exp (b) – 1). In this case, 100 (exp(0.233) -1) = 26%.

D.2 The FTI critique of the Frontier econometric analysis

The FTI report has undertaken an assessment of Frontier's econometric analysis to conclude that "*the analysis is not sufficiently robust to draw the conclusions that scarcity rents – the typical ticket price premium to other London airports due to capacity constraints – at Heathrow are 23%*".[32] It further indicates that "*some issues with the analysis and data suggest that scarcity rents may be larger or smaller than Frontier states.*"

In the rest of this section, we briefly introduce the main points raised by FTI about the Frontier's econometric analysis. Note that FTI report additional technical issues that affect the regression analysis that we do not mention here.[33]

D.2.1 A residual approach

The main critique offered by FTI is that the Frontier econometric analysis is based on a "residual approach" to estimate the congestion premium. In other words, the Frontier analysis does not estimate directly the effect of congestion on ticket prices at Heathrow, but instead attributes any price difference that is left unexplained by the regression to congestion.

The Frontier regression analysis does control for some of the factors that may explain the difference in ticket price between flights at Heathrow and other airports. For instance, it includes distance of the flight. This variable may approximate difference in the cost of operating flights, as the distance increases so are the costs of flights.[34]

However, the regression model leaves out some important controls that may affect difference in ticket prices.

- FTI notes that Frontier's estimation approach "*ignores the possibility that passengers may simply perceive specific airports in England to provide a better or worse service*" compared to others.
- Further, FTI notes that the regression model does not account for the effect of competition on the price levels in city-pair markets. As airlines compete to serve passengers on a city-pair market, the ticket price might be expected to depend on the degree of competition. Since competition might vary across markets, this factor is expected to determine the price level.
- In addition, as noted by the FTI report, the regression model does not account for (non-passenger-related) airport charges levied by airport operators on airlines, covering parking and landing charges. We understand that there are substantial differences in parking and landing airport charges between London airports and that they could explain a significant proportion of the estimated price premium at Heathrow. For example, as set out in the table below, airlines paid c. £660 million in landing and parking charges at Heathrow (c. £8.5 per passenger), compared to only c. £100 million at Gatwick airport (c. £2.2 per passenger).

---

32 FTI report, paragraph 5.2.
33 Notably the FTI report mentions the issue of aggregated data, function form and measurement errors.
34 Figure 44 in Annex A to the Frontier report lists all of the variables that are included in the regression model.

Table 10: Parking and landing charges at Heathrow and Gatwick, forecast 2018

	Heathrow	Gatwick
Passengers (in million)	77.0	45.7
Parking and landing charges (in million £)	657.6	102.8
Parking and landing charges (per passenger)	8.5	2.2

Source: Heathrow data based on 2018 forecast revenue (see https://www.heathrow.com/file_source/Company/Static/PDF/Partnersandsuppliers/airport-charges-consultation-document-2018.pdf); Gatwick data based on 2017/2018 accounts

As the Frontier regression model omits these (and other) important determinants of the ticket price level, this means that the Heathrow dummy variable, which represents the difference in ticket prices with other airports, also accounts for other effects than just congestion. Indeed, according to FTI, "*the Heathrow airport dummy variables may have captured multiple additional effects as well as scarcity rents*."[35] In other words, by failing to properly account for other important determinants, Frontier has mistakenly attributed their impact on price to the impact of congestion.

D.2.2 Omitted variable bias

Furthermore, and following the issue described just above, the Frontier regression suffers from an omitted variable bias, which means that the regression results are unreliable. The Frontier regression model does not include important factors that may explain the price difference between flights from Heathrow and other London airports. As we expect these factors to be correlated with the Heathrow dummy variable (e.g. passengers may value more access to London Heathrow or its amenities than those of other London airports), this omission implies that the error term of the regression model is correlated with the explanatory variable. This implies necessarily that the ordinary least square (OLS) estimator that Frontier applies to estimate the regression coefficients gives rise to inconsistent (and biased) results.[36]

D.2.3 An "ideal world" analysis

Importantly, the FTI report considers that the Frontier approach could be improved, and even proposes an "ideal world" econometric analysis that builds on Frontier's econometric analysis.

First, the FTI report considers that some confusion exists between slot congestion, which affects directly airlines, and the supply of seats, which are purchased by passengers. FTI considers that further research must be conducted to clear up the confusion, but that in any event prior to estimate the impact of congestion on ticket prices, "*an analysis of load factors on individual flights and routes should precede and inform the model specification*."[37]

Second, FTI proposes important adjustments to the Frontier analysis, notably by suggesting adding more controls, which would remedy the flaw associated with the "residual approach". For example, there might be Heathrow-specific factors that contribute to the price difference

---

35 FTI report, paragraph 5.11.

36 For more on omited variables, see A. Colin Cameron and Pravin K. Trivedi "Microeconometrics – Methods and Applications". 2005 Cambridge University Press.

37 FTI report, paragraph 5.52.

with other London airports. These may include for example the non-passenger-related airport charges, the ease of accessing Heathrow from London, the perceived quality of passenger service offered at Heathrow relative to that of other airports. If these factors push price up at Heathrow, it would be wrong to attribute these effects to congestion.

D.3 Our critique

We broadly agree with some of the issues raised by FTI about the limitations of Frontier's econometric analysis, as summarised just above. That said, we do not agree that the Frontier's econometric analysis can be improved to the point where their analysis yields reliable estimates on the congestion premium at Heathrow as FTI indicates in their report.

D.3.1 The econometric results of the Frontier analysis are unreliable

The Frontier regression analysis is based on a cross-section comparison of ticket prices in 2016. As indicated above, Frontier's residual approach attributes the effect of congestion to the Heathrow dummy variable.

In our view, applying OLS to estimate the Frontier regression model is likely to be biased and inconsistent for the following reasons:

- First, and in agreement with FTI, we consider that the Frontier regression analysis suffers from an omitted variable problem (see above). In addition to the factors listed by FTI that Frontier omits to include, it is also important to consider the airlines' cost of operating at Heathrow (beyond the airport charges already discussed above), which might be significantly higher than that at other London airports notably.
- Second, we consider that the regression also suffers from an endogeneity problem, which means that applying OLS leads to biased and inconsistent regression estimates. Specifically, the model includes as explanatory variable two variables related to frequency, namely "frequency_own" (number of flights to the same destination at the same airport) and "frequency_other" (number of flights to the same destination at other airports).[38] These frequency variables, whilst they may explain ticket price differences, may also be explained by ticket prices. As ticket price increases, supply (and thus frequency_own) is also expected to increase. In other words, there might be a two-way causal relationship between ticket price and frequency of flights. In technical terms, this means that there is another equation where ticket price explains frequency. It follows that since frequency is also caused by ticket prices, frequency is related to the error term in the ticket price regression. This violates the assumption that OLS is unbiased and consistent.

D.3.2 The Frontier regression model cannot be improved

FTI considers that the Frontier's regression model can be improved to measure the size of scarcity rent. We disagree for the following reasons:

---

38 See Figure 44 of Annex A of the Frontier report.

First, before measuring the size of scarcity rents, it is crucial to establish a relevant economic framework to determine the possible impact of slot congestion on airports, airlines and passengers. As we show above, congestion may raise airlines' opportunity cost of using slots, but that does not imply that airlines would be able to monetise congestion by raising ticket prices. It is, therefore, likely that the ticket price premium measured by Frontier reflects Heathrow-specific factors that are unrelated to congestion

Second, Frontier's approach, which consists of attributing the difference in ticket prices between Heathrow and other airports that is left unexplained to scarcity, is inherently flawed. This cross-section analysis cannot guarantee that all Heathrow-specific factors that may affect ticket prices can be included in the regression model. In fact, the only way to guarantee that all Heathrow-specific factors are taken into account is to introduce in the model a dummy variable for Heathrow, which is exactly what Frontier does, but it interprets this variable as capturing only congestion at Heathrow. Furthermore, and by the same token, the ticket price difference between Heathrow and other London airports, may also be due to specific elements that impact individually some of these airports. Again, it appears inconceivable that the regression model would be able to include all these airport-specific variables.