

IEA Discussion Paper No.107

TERMINAL PROBLEM?

The case for a market-based airport slot
allocation system

Kristian Niemietz
May 2022



iea

Institute of
Economic Affairs

With some exceptions, such as with the publication of lectures, IEA Discussion Papers are blind peer-reviewed by at least one academic or researcher who is an expert in the field. As with all IEA publications, the views expressed in IEA Discussion Papers are those of the author and not those of the Institute (which has no corporate view), its managing trustees, Academic Advisory Council or senior staff.

Contents

About the author	4
Summary	6
Introduction: air travel in the UK – the big picture	8
The two bottlenecks: capacity and slot allocation	14
Mitigation measures and their limitations	20
A market-based alternative	26
Second-best solutions	33
References	36

Dr Kristian Niemietz is Head of Political Economy at the Institute of Economic Affairs. He studied Economics at the Humboldt University Berlin and the University of Salamanca, graduating in 2007 as *Diplom-Volkswirt* (≈MSc in Economics). He also studied Political Economy at King's College London, graduating in 2013 with a PhD.

Kristian previously worked at the Berlin-based Institute for Free Enterprise (IUF), and taught economics at King's College London. He is the author of the books *A New Understanding of Poverty* (2011), *Redefining the Poverty Debate* (2012), *Universal Healthcare Without the NHS* (2016) and *Socialism: The Failed Idea That Never Dies* (2019).

Summary

- Since the mid-1980s, the UK's aviation sector has undergone substantial liberalisation. This was a great success story as far as it went. But two major components have bucked that overall liberalisation trend: airport capacity, and the allocation of take-off and landing slots. These are the two remaining islands of central planning in an otherwise liberalised sector.
- At congested airports, airport operators have little control over the allocation of take-off and landing slots. If an airline has used a particular slot in the past, it has an automatic right (subject to some conditions) to carry on using it – indefinitely. The remainder of the slots are allocated by a quango, via an administrative process, under retained EU law (ultimately derived from the International Air Transport Association (IATA)).
- As a result of the 'grandfathering' of slots, most slots at large airports are simply held by the airline that has always held them. This distorts the market in favour of well-established incumbents. It is the main reason why large European airports continue to be dominated by national legacy carriers. Despite certain exceptions to the slot allocation rules, the market remains distorted. There is scope for major improvements.
- As in many other areas, Brexit is a double-edged sword when it comes to aviation. Brexit enables the UK to replace the EU/IATA slot allocation rules with an alternative system. This could mean replacing it with a more liberal and market-based alternative – but it could also mean replacing it with an even more dirigiste one, under which slots are allocated on the basis of political considerations.
- A more liberal alternative would mean introducing a primary market for airport slots. That could be done by a simple auction, or by congestion pricing, or it might be possible to allow airport operators to devise

whatever slot pricing mechanism they see fit, and rent out airport slots to whoever they choose. Grandfathered rights would be phased out, and the slot allocation rules would cease to apply.

- Taken to the fullest extent, this could mean different things in practice. Some airports might use periodic slot auctions, under which slots go to the highest bidder. Auction design and auction rules could then differ from airport to airport. Other airports might use a system of dynamic runway pricing, in which they set market-clearing prices for slots. It could also lead to the emergence an altogether different allocation mechanism.

Introduction: air travel in the UK – the big picture

At the time of writing, the aviation industry is climbing out of the deepest crisis of its history. In 2020, passenger numbers at UK airports dropped by an astonishing three quarters, from nearly 300 million in 2019 to less than 75 million (Department for Transport (DfT) 2021). Preliminary figures for 2021 do not suggest much of a recovery either (Eurocontrol 2022).

Like many other sectors, air travel virtually ground to a halt during the lockdowns (especially the first one in Spring 2020). But unlike many other sectors, it did not bounce back to anything like its pre-lockdown level afterwards. Ongoing travel restrictions, Covid-19 testing requirements and the associated costs, uncertainty about changes in travelling rules, and simply the hassle of the ‘Covid bureaucracy’ turned out to be a lasting drag on air travel. The sector’s perennial issue of insufficient airport capacity suddenly seemed like a luxury problem to have, as deserted airports became a more familiar sight than congested ones.

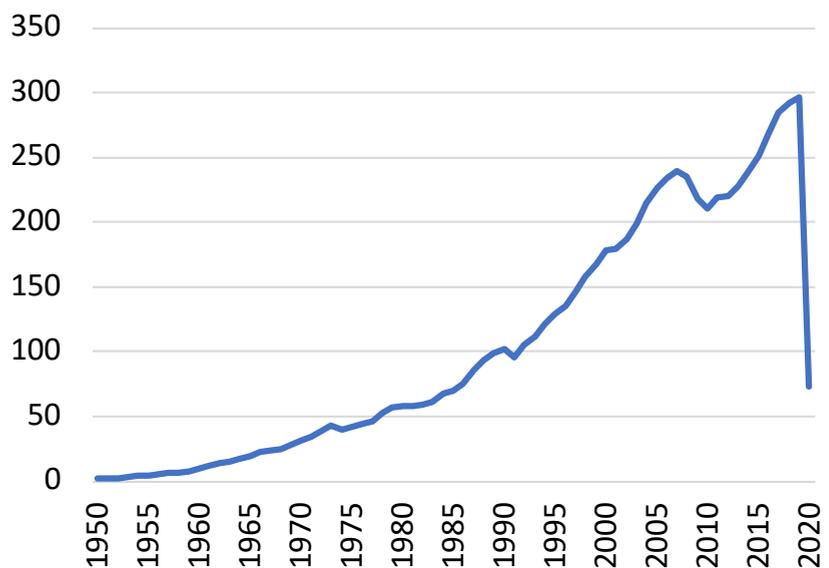
Nonetheless, we should not forget that this steep drop in passenger numbers follows decades of rapid growth. A three-quarters drop in a single year is undoubtedly dramatic, but it is worth remembering that until the mid-1980s, passenger numbers of the kind we saw in 2020 would have seemed impressive.

From the mid-1980s until 2019, however, the number of passengers passing through British airports had been growing by an average of almost 7 million per year. In 2019,

- Stansted alone handled more passengers than all UK airports taken together had done in 1969

- Gatwick alone handled more passengers than all UK airports taken together had done in 1977 and
- Heathrow alone handled more passengers than all UK airports taken together had done in 1986 (based on DfT 2021 and Civil Aviation Authority (CAA) 2021).

Figure 1: Terminal passengers at UK airports (in millions), 1950–2020



-based on DfT (2021)

This was not simply explained by a few frequent flyers, or by wealthy foreigners travelling to the UK. Before Covid, more than half of the UK adult population took at least one flight per year, and more than two thirds took at least one flight every four years (based on CAA 2019).

Table 1: Proportion of UK adults who have flown from a UK airport at least once, 2019

Time frame	% of UK adults who have taken ≥ 1 flight during that time
Within the last year	51%
Within the last 4 years	68%
Within the last 10 years	78%

-based on CAA & ComRes (2019)

These are remarkable figures if we bear in mind that there will always be some people who have no interest in travelling abroad, who prefer other means of transport or for whom flying is not a practical option (e.g. for health reasons, fear of flying, etc.).

There is an income gradient in travel behaviour, but flying is by no means reserved for the better-off. In 2018, at a typical UK airport, around one in five leisure travellers had an annual household income of less than £23,000.¹ Only London City Airport could fairly be labelled a 'rich people's airport'.

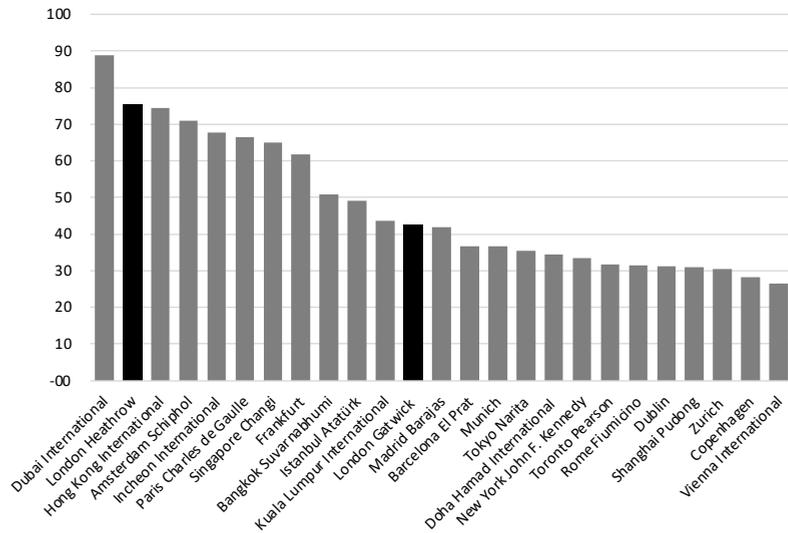
¹ Income figures are reported in bands, rather than as a full distribution. Hence the choice of £23,000 as a cut-off point (and not because there is anything special about that number).

Table 2: Proportion of UK leisure travellers with a household income of less than £23,000, 2018

Airport	% earning <£23,000
Heathrow	19.9%
Gatwick	15.3%
Manchester	21.7%
Stansted	20.3%
Luton	30.1%
Edinburgh	15.3%
Birmingham	22.8%
Glasgow	15.2%
Aberdeen	22.0%
East Midlands	17.1%
London City	11.2%
Inverness	27.0%

-based on CAA (2019)

In 'normal' times, London is home to two of the world's busiest airports. Heathrow is the busiest airport in Europe, ahead of Amsterdam, Paris and Frankfurt. Even so, Gatwick is still ahead of international top-tier airports such as Madrid, Barcelona, Munich, Dublin and Rome. The Greater London area, in particular, offers a level of choice between airports, and competition between airports, which is unparalleled in Europe. According to the 'Air Connectivity Index', compiled by the International Air Transport Association (IATA), London is the best-connected city in the world (Pickett & Hirst 2020: 23).

Figure 2: Number of passengers (in millions) by airport, 2018

-DfT (2019)

In short, UK aviation up until 2019 is a long-running economic success story and, in all likelihood, the pandemic will turn out to be no more than a dramatic one-off disruption of that.

It is a success story that did not happen by luck or accident. It happened as a result of specific policy choices.

Historically, in the UK and elsewhere, aviation used to be a heavily state-dominated, state-directed industry. In recent decades it has undergone substantial liberalisation, and it is no coincidence that the sector's growth has accelerated since then.

The OECD's International Transport Forum (ITF) explains:

The results of deregulation have been closely studied [...] providing a rich body of economic literature on the topic. Overall, liberalization [...] has driven down air fares, [...] improved connectivity and supported the growth of trade, tourism and the broader economy (ITF 2019: 17).

The International Civil Aviation Organization (ICAO) agrees:

Liberalization has led to substantial economic and traffic growth. [...] [I]ncreased competition in the aviation market [...] reduces price and stimulates traffic growth, producing efficiency gains [...]

The emergence and growth of low-cost carriers, made possible through liberalization, has [...] led to increased competition and reduced fares (ICAO 2016: 3).

This was an international policy trend, but the UK has been a forerunner in this regard: in the UK, the liberalisation process started in the mid-1980s, thus earlier than in most of Europe (albeit behind the US, where it had already started in the late 1970s). The UK also went further than most of its peers. It has been a game changer (see Lesh 2019: 6–8; Niemietz 2013: 12–14).

According to current forecasts, the European air travel industry is expected to return to its pre-pandemic level over the next two years (Eurocontrol 2022: 20). This is, of course, subject to major uncertainties, and it will mainly depend on factors about which the forecasters cannot know more than anyone else. It could happen years later, or it could already have happened by the time this paper is published. But the point is that even the more pessimistic forecast scenarios do not predict long-term scarring effects.

In the short to medium term, the sector can be expected to return to normality. As much as that is to be welcomed, it also means that the sector's long-standing chronic problems are going to reassert themselves once again. These may have seemed like luxury problems in the midst of a pandemic, but they did not seem so luxurious before Covid, and they will not seem so luxurious for long.

The two bottlenecks: capacity and slot allocation

The overall policy trend in aviation since the mid-1980s has been towards liberalisation. However, two key variables of the sector have so far bucked this trend: airport capacity investment, and the allocation of take-off and landing slots. These two variables have never really been subjected to market forces, and remain largely outside the scope of the market economy.

1. Airport capacity

Airport capacity is shaped by the political decision-making process controlling the building and expansion of airports. This makes it different from capacity investment decisions in most other sectors of the economy, where such decisions are, in the main, business decisions, not political ones (albeit subject to planning controls and other regulatory constraints). We have addressed this issue elsewhere (see Niemietz 2013).

2. Allocation of take-off and landing slots

In most other sectors of the economy, if a company wants to use an asset that it does not own itself, it has to purchase the right of use from the asset owner. It has to, for example, rent or lease the asset, and pay market rates for it. If demand for the asset outstrips supply – be it permanently or at particular peak times – various bidders will have to compete for the right to use the asset. The question of who gets to use the asset at what time is settled via the market process, not via political decisions.

Airlines do not own airport infrastructure. They have to effectively ‘rent’ take-off and landing slots at airports, that is, the right to use the airport’s runway at a particular time slot. However, this process of ‘renting’ take-off and landing rights looks nothing like a conventional rental market. Once

an airport is classified as ‘congested’ it becomes a ‘coordinated airport’. This means that slots are allocated via an administrative process under rules inherited from the EU (which, in turn, derived them from the guidelines of the IATA). This feature has been described as ‘a legacy of the days of state-run, flagship carriers that has not been addressed by privatisation or deregulation’ (Pheasant & Giles 2007: 31), and as ‘a relic of the heavily state-controlled system of the 20th century’ (Lesh 2019: 9).

In the UK, Heathrow, Gatwick, London City, Stansted, Luton, Manchester and Birmingham are classified as coordinated airports, while Bristol Airport is part-time coordinated. Thus, while the majority of airports are not coordinated, the coordinated ones account for the majority – about three quarters – of all passenger movements (based on CAA 2021). In this sense, coordinated status is the norm in the UK, not the exception.

At coordinated airports, slots are allocated in two ways.

1. If an airline uses a slot in the current winter/summer period, it has an automatic right (subject to some conditions) to use that slot again in the next equivalent period, and this right can be rolled over in perpetuity. This practice, known as ‘grandfathering’, means that the default position is that a slot is used by the airline that has always used it. It represents an implicit subsidy to long-established incumbents. Airlines pay service charges based on the airport’s costs, but they do not pay for the slot as such. At congested airports, where demand for slots exceeds supply, these regulated charges come nowhere near what the market value of the slot would be (Competition and Markets Authority (CMA) 2018: 6). We can see this by simply looking at the prices that slots fetch on the secondary market (more on which later). At Heathrow, a slot pair² is usually worth at least £5 million (in the less busy hours), and peak-hour slots can be worth many times more than that (Pickett & Hirst 2020: 26–27). Thus, the slot portfolio of a large incumbent can be worth billions of pounds.

Grandfathered slots are a bit like social housing, in the sense that social housing tenants also have an indefinite right to remain in their flat, and that they receive an implicit subsidy in the form of below-

2 In this paper, we will sometimes informally talk about ‘a slot’, although the technically correct term is a ‘slot pair’. Slots always come in pairs, since a take-off slot would be worthless if not bundled with a landing slot. What goes up must come down.

market rents. (The difference, of course, is that in the case of social housing, the implicit subsidy is intentional – it is the whole point.)

2. Slots that are not currently held by any airline are allocated by a quango, a designated slots coordinator, subject to the above-mentioned EU/IATA rules. In the UK, this role is performed by Airports Coordination Limited (ACL).

Thus, neither airport capacity planning, nor the allocation of airport slots, is governed by market processes. Capacity planning is subject to political determination, and slot allocation to a regulatory process. They are islands of central planning in an otherwise market-driven sector, and this gives rise to a series of problems.

There is certainly a problem of insufficient airport capacity. As early as 2003, a report by the DfT found:

[M]any airports in the UK are becoming increasingly congested as they attempt to cope with rising passenger numbers. In some cases, the capacity of terminals and runways is at, or near, saturation point. At Heathrow [...] the two runways are already full for virtually the whole day. The same is true at Gatwick, already the world's most intensively used single-runway airport. [...] Birmingham's runway is already close to its existing capacity during peak times [...] Edinburgh is approaching the limit of its existing terminal capacity and urgently needs further investment. The provision of some additional airport capacity will therefore be essential if we are to accommodate, even in part, the potential growth in demand. [...] Failure to provide additional capacity would become a barrier to future economic growth and competitiveness. Airports would become more congested; air fares would rise as slots became increasingly sought-after; and much of the future growth in air travel [...] could in due course migrate elsewhere (DfT 2003: 24–25).

In the meantime, numerous other reports have confirmed that same problem, or rather, an intensification of it (see Niemietz 2013: 14–17). Before Covid, lack of capacity was the main bottleneck holding back the sector. We have discussed solutions to this problem elsewhere (Niemietz 2013: 51–64).

When total capacity is insufficient, it becomes all the more important to use the existing capacity in the most efficient way possible. But this is precisely what the current system of slot allocation prevents.

It is not that the slot allocation rules leave no room at all for market forces. Subject to some constraints, airlines can swap slots between them or lease them (more on this later), which allows market forces in through the back door. Nonetheless, there are a number of problems with the slot allocation process, to which secondary slot trading has only been a partial solution.

1. Most obviously, the system hampers competition between airlines by distorting the market in favour of well-established incumbents (Pickett & Hirst 2020: 30–32). It is notable that, even today, decades after the beginning of the liberalisation process, Europe's large airports are often still dominated by national legacy carriers that established their position in the pre-liberalisation period. For example, at Heathrow, the International Airlines Group (IAG) – the parent company of British Airways – holds 55 per cent of all slots (WPI Economics 2019a: 13). A quarter of all long-haul routes, and over half of all short-haul routes, are only served by IAG (ibid.: 15–16).

Such figures are not unusual for busy airports under this system, and not just in the UK. At Frankfurt Airport, the Star Alliance – the parent company of Lufthansa – accounts for seven out of ten aircraft movements (Fraport AG 2022: 39).

2. It leads to an inefficient allocation of a scarce resource (Pickett & Hirst 2020: 29–30). Ideally, each slot would be held by the airline that can make the best use of it. In order to do this, the slot coordinator would have to know which airline that is, for each slot. It seems implausible that any organisation can possess that knowledge.

The basic problem here is that the current slot allocation process resembles a form of central planning, in which the slots coordinator assumes a role not unlike that of a planning board in a socialist economy. The CMA critiques this arrangement in almost 'Hayekian' terms:

In our view, it is infeasible for an administrator to make an ex-ante assessment of which airlines would be the most efficient user of the new capacity. [...] No administrator, however good, will have the knowledge to decide which slots, and how many

slots, should be allocated to optimise each airlines' network and business requirements. [...] There is no clear rationale for why these important commercial decisions should be made by a third-party administrator that is not privy to knowledge about each individual airlines' commercial and business interests (CMA 2019: 24).

One could add that even if the initial allocation of slots were optimal, it would be unlikely to remain so for long, as market conditions change.

3. The CMA points out that in addition to weakening competition between *airlines*, the system also, indirectly, weakens competition between *airports* (CMA 2018: 10). It does so in several ways.
 - 3a.) As mentioned above, when total capacity is insufficient, it becomes all the more important to use the existing capacity in the most efficient way possible. Inefficiencies in the use of runway capacity magnify the already existing capacity problems.

When two (or more) airports are running at full capacity most of the time, they no longer really compete with each other for new customers, because there is no point in trying to attract new customers if they could not accommodate them anyway. Whether they cannot accommodate them because the physical capacity is simply not there, or because something prevents them from using their existing capacity efficiently enough, makes no difference in the end.

- 3b.) Even when there is some spare capacity in the system, the slot allocation rules 'gum up' the market. Suppose an airline is (mainly or solely) based at Airport X, and ponders switching part of its fleet to Airport Y. In the current system, airlines will be reluctant to give up slots at Airport X that were allocated to them originally for free (in pre-congestion days), and they will also find it hard to acquire new slots at Airport Y. Thus, airlines' 'switching rates' between airports will be lower than they would be in a more market-based system.
 - 3c.) This then leads to greater market segmentation, such that consumers who have a preference for Airline A have to go to Airport X, ones who have a preference for Airline B have to go to Airport Y, and so on (CMA 2018: 10). This kind of airline–airport bundling has clearly

happened in the UK (e.g. Heathrow – British Airways, Gatwick – EasyJet, Stansted – Ryanair).

Of course, segmentation need not be a bad thing. It does give airports a degree of market power, but it can also be more efficient for airlines to bundle their activities around selected airports rather than spreading them out thinly, and the bundles may also align with consumer preferences. It *could* be a desirable outcome – we just cannot know whether it really is, until airports and airlines are also able to experiment with other business strategies.

4. As Gillen and Starkie (2016: 156–9) explain, the slot allocation system can contribute to the perpetuation of the capacity problem. This is because it also produces winners, namely, the incumbent airlines that benefit from the free past allocation of valuable slots, and from a less competitive market environment. Those incumbents are not a major political constituency in their own right, and they are not politically influential enough to block airport expansion on their own. But there is already a broad coalition of opponents to airport expansion, from local residents opposed to aircraft noise to environmentalist groups. In such a context, opposition from incumbent airlines can tip the scales.

Indeed, Gillen and Starkie (2016: 159–61) show that company leaders' reactions to expansion proposals have ranged from unenthusiastic to hostile.

5. Another unintended side effect of the system is that it fuels mergers and acquisitions in the airline industry (Gillen and Starkie 2016: 155–6). The reason is that if Airline X takes over Airline Y, it also takes over its slots. There need be nothing wrong with mergers and acquisitions, which can be efficiency-enhancing and beneficial for consumers. However, the slot allocation system creates motives for mergers that are not inherently efficiency-enhancing, but merely profitable because of the way landing slots are allocated. If mergers and acquisitions become a clumsy, roundabout way of acquiring take-off and landing slots, it distorts the industry structure.

Mitigation measures and their limitations

The current slot allocation system has a number of mechanisms to mitigate the problems described above. Unfortunately, those are either inadequate or they create further unintended and damaging effects of their own.

1. Use-it-or-lose-it rules

Airlines do not have an unconditional right to keep a slot forever. They have to actively use each slot at least 80 per cent of the time and, if they fail to do so, it is returned to the slot pool. This rule is supposed to prevent airlines from ‘hoarding’ slots they do not need.

But arguably, this solves the wrong problem. At a busy, congested airport, in normal circumstances it is quite unlikely that an airline will possess a slot for which it has no use whatsoever: it will always be able to find *some* use for it. But the question is whether it is an *efficient* use, and whether a different airline could have put it to *better* use.

There are some underused slots – and thus underused capacity – even at the busiest airports (Pickett & Hirst 2020: 29). What often happens in such cases is that airlines comply with the letter, but not the spirit, of the use-it-or-lose-it rules (Pickett & Hirst 2020: 32–33; Lesh 2019: 16). The CMA explains:

To comply with the ‘use it or lose it’ rule, many airlines resort to artifice – flying smaller planes than necessary in order to spread capacity across their slots, for example, and even running empty ‘ghost’ flights [...] [I]nstead of slots being recycled from established carriers to new ones, they are clung to (CMA 2018: 17).

Ghost flights became a particularly salient issue during the very unusual circumstances of the Covid-19 pandemic. As mentioned, in 2020, passenger numbers fell by three quarters. However, airlines cannot simply cancel three quarters of their flights. If they did, they would run afoul of the use-it-or-lose-it rules, and thereby risk losing valuable slots.

Although the use-it-or-lose-it threshold was temporarily cut from 80 per cent to 50 per cent (and even to effectively zero at one stage), this was not enough to accommodate the large and sustained drop in demand. Thus, some airlines flew empty planes in order to hold on to their slots. The Lufthansa Group, for example, flew about 18,000 empty flights into and out of British airports over the winter of 2021/22.³

This practice, quite understandably, enraged environmental campaigners. The policy director of Greenpeace UK, for example, said:

We know that the airline industry puts profit ahead of people and the planet but the absurdity of 'ghost flights' takes its recklessness to new heights. [...] These empty flights [...] must be grounded immediately.⁴

Greenpeace is, of course, right to highlight the absurdity of ghost flights. But by turning this into a story of moral failing rather than bad economics, it is missing the main point. Evidently, the flying of an empty flight is, in itself, a loss-making activity, so what needs explaining is why a profit-oriented company would ever engage in it. The answer is that it can become profitable once we take the knock-on effect on future slot allocation into account. That knock-on effect, however, only exists because of use-it-or-lose-it rules. Thus, airline companies respond rationally to the incentives provided by the system under which they operate. We do not need political action to order the grounding of those flights, but a rational slot allocation system under which airlines would simply have no reason to run such flights.

Others have responded by calling for use-it-or-lose-it rules to be scrapped. A petition submitted to the House of Commons (signed, at the time of writing, by over 6,000 people) reads:

3 Matthew Lesh: 'The EU rules creating an armada of empty "ghost flights"', *The Spectator*, 9 January 2022. (<https://www.spectator.co.uk/article/the-eu-rules-creating-an-armada-of-empty-ghost-flights->)

4 Samuel Webb: 'What are "ghost flights" and why are they so controversial?', *The Independent*, 12 January 2022. (<https://www.independent.co.uk/climate-change/infact/ghost-flights-eu-airlines-airports-climate-b1991531.html>)

UK regulation states that airlines must use their landing slots more than 80% of the time in order to keep them. [...]

We ask the Government to reduce this to 0% as a permanent measure.

Airlines have been flying planes empty to retain their landing slots. [...] [I]t beggars belief that planes fly empty. 'Ghost' flights are of no benefit to anyone. This is a needless, wasteful practice.⁵

This statement correctly identifies flawed regulations as the cause of the existence of ghost flights. However, merely calling for use-it-or-lose-it rules to be scrapped (which is what reducing the threshold permanently to 0 per cent amounts to) still falls well short of the ideal. Whatever their faults, use-it-or-lose-it rules exist for a reason, which the initiators of the petition fail to acknowledge. The current slot allocation system has an in-built bias towards slot-hoarding, and use-it-or-lose-it rules are supposed to counter that bias. So unless we are prepared to accept an increase in slot-hoarding, within the current system, use-it-or-lose-it rules cannot so easily be abolished without a replacement.

2. Reserved slots for new entrants

Under the current rules, when new capacity (and thus a set of new slots) is created, or when existing slots are returned to the pool and reallocated, at least half of them have to be reserved for new entrants with a small existing presence at the airport in question, or none at all (Pickett & Hirst 2020: 10–11).⁶ This is meant to mitigate the pro-incumbent, anti-competitive effect of the grandfathering system.

Unfortunately, in doing so, it also pushes the industry structure into a particular direction – and it is not necessarily an efficient one. It could well be that the most efficient way to tackle the market power of large incumbents is not to add lots of very small players to the market, but to allow some of the existing small-to-medium-sized ones to scale up.

5 'End "ghost" flights: reform historic rights to landing slots'. Petition. Available at <https://petition.parliament.uk/petitions/605749>

6 In this context, a 'new' entrant does not have to be a new company. It has to be an airline that has so far only had a minimal slot holding, or none, at a given airport (or its nearby competitors). Thus, a company like Lufthansa, which has existed for nearly a century (if one counts its predecessor), could, in this sense, be a 'new entrant' – 'new' as in 'new to the airport in question'.

This is what happened in, for example, the UK grocery retail sector, a formerly concentrated market that has become more competitive over time. Just over a decade ago, there were widespread concerns about the market power of Britain's 'Big Four' supermarket chains (Tesco, Sainsbury's, Asda and Morrisons), and in particular the market leader, Tesco. It has since become clear that those fears were unfounded. The combined market share of the Big Four has since fallen by about ten percentage points, and this is before we even take account of online delivery services.⁷ What is notable is that this development was *not* driven by new entrants or very small players. It was driven by already well-established minor actors, especially Aldi and Lidl, scaling up to become medium-sized ones. If we allocated retail capacity in the way we allocate runway capacity, the breakthrough of Lidl and Aldi might never have happened. The slot allocation system favours newcomers or actors with a very minor market presence, so Aldi and Lidl might not have qualified. The 'retail slots' would instead have been given to tiny corner shops. This would have created a market structure with a few very large and lots of very small players, but nothing in between. In all likelihood, this would have been a highly inefficient outcome.

And yet, in aviation, this is the type of markets that the slot allocation rules have created. The bulk of an airport's slots are often taken up by one single airline or alliance of airlines, with the remaining slots being thinly spread out among a large number of competitors. As mentioned, at Heathrow, IAG holds 55 per cent of all slots. At the opposite end of the spectrum, there are 62 airlines with a slot share of less than 1 per cent each. The closest to being medium-sized players are the Lufthansa Group, with a slot share of 8 per cent, and Virgin Atlantic, with a slot share of 7 per cent (WPI Economics 2019a: 9).

There is good reason to believe that that is not quite enough to mount an effective competitive challenge to IAG. The reason is that for airlines operating at hub airports, 'upscaling' does not necessarily just mean doing more of what they already do. It means acquiring features of a hub carrier. A hub carrier is an airline that does not just fly people from X to Y. Rather, it first brings passengers from A, B, C and D to X, pools them there, and *then* flies them to Y (ibid: 18–19). Thus, giving an airline a slot to fly from X to Y will not help it, if it cannot also bring in passengers from A, B, C and D.

7 See Kantar World Panel (<https://www.kantarworldpanel.com/en/grocery-market-share/great-britain>)

Hub carriers need not be gigantic, in terms of their slot holdings at any particular airport, but a typical hub carrier is considerably larger than any of IAG's rivals (ibid: 26–29). So here, and perhaps counterintuitively, the reserving of slots for altogether new entrants works against the creation of effective competition. As the CMA (2018: 9) points out, the system

risks creating unintended consequences including increasing the number of very small operators, whereas consumers may be better served by a smaller number of slightly larger operators.

3. Secondary slot trading

Airport slots cannot formally be bought and sold. Airlines can, however, swap slots bilaterally (Pickett & Hirst 2020: 14–15). Crucially, if Airline A holds a valuable slot, and Airline B a less valuable one, they can combine the swap with a side-payment from Airline B to Airline A. This was initially a legal grey area, but over the past decade or so it has become widely accepted practice. In this way, a semi-formal market in take-off and landing slots has developed (Lesh 2019: 15). Similarly, an airline can also lease a slot from another.

Of all the measures aimed at improving the slot allocation process, the toleration of secondary trading is by far the most sensible and effective one. Suppose a slot is worth £12 million to Airline A, and £10 million to Airline B, but is allocated to Airline B (either because of grandfathering, or because of the slots coordinator's judgement). This is clearly a misallocation: if Airline A values that slot more highly, this is because Airline A can put it to better use than Airline B. The secondary market for slots now makes it possible for Airline A to approach Airline B, and offer to effectively purchase the slot from it. In this way, the initial misallocation can be corrected later.

Secondary trading has its limits, both legally and in practice. For a start, bilateral swaps at a given airport cannot, in and of themselves, change the slot shares of different airlines. If Airline A has 100 slots at a particular airport, and Airline B has 10, then no amount of slot-for-slot swapping can ever change that ratio. And if Airline C has no slots at all at that airport, it cannot even engage in secondary slot trading, simply because it has nothing to trade. In order to benefit from the opportunities of the secondary market, one needs to be in the 'primary market' already.

Airlines cannot swap slots at a ratio other than 1:1. They cannot, for example, trade one highly valuable slot for two or three less valuable ones. Nor can they simply purchase a slot directly, and pay for it solely in cash rather than in kind (i.e. with another slot).

In practice, there are ways around these restrictions. Indirectly, slot trading can lead to changes in the slot shares of different airlines, and it can take forms that look very much like direct purchases. Secondary slot trading has been a success. In the recent past, there have been successful examples of newcomers not just breaking into the UK market, but also upscaling, and becoming major players at some of the UK's largest airports. Where this has happened, the slot trading system has played a role (Pickett & Hirst 2020: 34). For example, EasyJet only started operating at Gatwick Airport in 2002, and initially only had a minor presence there. A decade later, it had become the dominant player at Gatwick, relegating the previous incumbent, British Airways, to second place (Pickett & Hirst 2020: 24–25). EasyJet has also become the dominant player at Luton and Bristol, and a major player at Manchester and Stansted, among others. Ryanair, which also used to be a newcomer, is now the dominant player at Stansted, and a major player at Manchester, Luton, Bristol and Birmingham, among others.

Even so – as a corrective to the problems with the primary allocation mechanism, secondary trading has only gone so far. There seems to be an inverse relationship between the level of congestion and the volume of slot trading, which means that slot trading works least well where it is most needed. At busy airports and/or during busy hours, the secondary slot market is an illiquid market characterised by small numbers of transactions and, even then, most of these take place within alliances (CMA 2019: 15 & 27). This also means that prices fluctuate wildly: there is no such thing as a 'going rate' for a slot. It is more like the market for rare artworks, where the price can be determined by a single buyer and a single seller.

It is beyond the scope of this paper to explore what exactly is holding back secondary trading, or whether it could be improved further. But whatever the exact reason, it is safe to say that the secondary market, while unambiguously a huge benefit, is only a partial solution.

A market-based alternative

The airport slot allocation rules for coordinated airports in the UK are legacy rules inherited from the EU. Brexit has now made it possible to deviate from those rules.

‘Deviating’ can mean very different things, in practice. It is possible to replace the current allocation rules with a more liberal, market-based alternative – but it is also possible to replace them with a more dirigiste, interventionist alternative. In this area, as in so many others, Brexit is very much a double-edged sword.

From a liberal perspective, the current system is very far from ideal. But it is easily possible to imagine even worse alternatives. The current system is, at least, broadly rules-based rather than discretionary, and independent of politics. A politicised system would almost certainly produce worse outcomes.

We might well be moving towards such a system. In 2018, the then Secretary of State for Transport, Chris Grayling MP, said:

[W]e would expect to reserve up to about 15% of slots on the new runway [at Heathrow] for domestic connections. They are a really essential part of the case for this. I have been very clear that there has to be capacity that is available only for domestic connections [...] We will make provision [...] to ensure that there is specific reserved capacity for regional connections within the United Kingdom (House of Commons Transport Committee 2018: Q478).

Elsewhere, the government also signalled its intention to ‘protect slots to support at least 14 domestic routes’ (cited in WPI Economics 2019a: 16).

WPI Economics (2019a: 21–23) also explores (without specifically endorsing them) a number of options that would align the allocation criteria of the slots coordinator, ACL, more closely with various government objectives. What these proposals all have in common is that they would, to varying degrees, politicise the slot allocation process, turning it into a tool for regional and/or industrial policy.

Elsewhere (and specifically in the context of Heathrow expansion), WPI Economics (2019b) proposes to tweak the slot allocation system in such a way that it can be used to systematically build up a second hub carrier at Heathrow, as a rival to IAG. This, too, could be considered a form of industrial policy.

What all proposals of that kind have in common is that they start with a pre-conceived notion of what the ‘ideal’ slot allocation pattern would look like, and then try to use the slot allocation system as a tool for bringing that outcome about. A market-based alternative, in contrast, would start from the presumption that we cannot know the optimal outcome in advance, and that the main purpose of the allocation system should be to find that out in the first place. Proponents of dirigiste solutions see the slot allocation system as a vehicle to get us to an already known destination. Proponents of liberal solutions believe that that destination is, as yet, unknown, and that the purpose of the allocation system is to find out where we should go.

In this spirit, economists have variously proposed systems of periodic slot auctions (see e.g. Bichler et al. 2021; WPI Economics 2019a: 24–27) or, alternatively, a system of congestion charges for busy runways (see Donohue & Hoffman 2007; Pickett & Hirst 2020: 39–40).

Under the first proposal, airlines would have to bid for the right to use a slot during a particular period. That right would be sold to the highest bidder, usually on a time-limited basis (depending on the specifics of the proposal) and, at the end of each period, the cards would be reshuffled anew. If the current slot holder wants to continue using its slots, it would have to bid for them anew, under the same conditions as any other bidder. Current usage rights would not confer future usage rights.

Under a system of dynamic runway pricing, take-off and landing charges would vary with demand for runway capacity, being highest during the busiest hours. This represents an indirect form of slot pricing. The slot itself may still be notionally free, but since it cannot be used without paying

the runway charges, those charges become a de-facto slot price. If the current holder of a popular slot is not prepared to pay the higher charge, it would have to transfer it to a competitor who is.

The two approaches could lead to similar (or, in theory, even identical) outcomes. For that to happen, congestion pricing would have to be sufficiently fine-grained (as opposed to, say, just distinguishing between peak and off-peak), and the charges would have to approximate the slot prices that would have resulted from an auction. What the two methods have in common is that they are open-ended processes.

We can, however, take that logic several steps further. Rather than proposing a discovery process for the optimal allocation of runway slots, in this paper, we are proposing a 'meta-discovery process', that is, a discovery process to find out what the best discovery process is.

If we accept that airport operators are the owners of the runways and the associated physical and logistical infrastructure, then it should be up to each one to make its own decisions on these matters. It should be up to the airport operator to decide whom it wants to permit to use its runways, at what time, at what price, for how long and under what conditions. If so, then it should also be up to each airport operator to decide *how* it wants to decide. Airport operators *might* very well choose slot auctions as a mechanism to allocate their slots. They *might* very well choose a system of dynamic runway pricing. They *might* choose to simply stick to the current system. Or they might choose something completely different.

This meta-discovery process could lead to a great diversity of allocation methods. Airport A could use a slot auction, Airport B could use a slot auction with a completely different type of auction design, Airport C could use a system of slot rental akin to car rental with fast turnover, Airport D could use a system of long-term slot leases with less turnover, Airport E could introduce congestion charges without altering the slot allocation process as such, Airport F could simply voluntarily stick to the old EU/IATA rules, Airport G could use an altogether different method and Airport H could use a combination of several methods. But a meta-discovery process could also lead to the emergence of a standard method, with very little diversity, or none at all.

Airport operators may choose to run their own allocation system in-house, but they could also choose to subcontract it to a specialised intermediary

that works with several airports. This could, for example, be a way of dealing with the interdependency of airports. ACL could survive in such a system, if it reinvents itself as one such intermediary.

It would be a process of learning and adaptation, and this would be true even if all airport operators in the country quickly converged on a single method. For example, even if periodic slot auctions immediately became the standard method, there would still be a need for a discovery process to work out the details. For a start, what does 'periodic' mean? An annual auction? An auction every seven years? Every twelve years? Staggered auctions? Or would the length be endogenous, with slots of varying durations going on offer? Flight schedules are determined long in advance, and airlines need stability and predictability in order to be able to plan ahead. They would be less inclined to invest in a place that does not offer them that. There is a trade-off between dynamism and predictability, and whoever runs the slot allocation system would have to take that into account.

Further, an auction is not an off-the-peg solution. Auctions come in all kinds of shapes and sizes. There are dozens of different types of auction design, some of which can be further divided into sub-types.

But whichever methods were to be adopted, what they all have in common is that they would lead to a more realistic pricing of a scarce resource (where 'realistic' means 'reflecting that scarcity'). Irrespective of how exactly we arrive at that pricing system, a number of beneficial effects seem likely.

1. Incumbent airlines would lose their privileged position. It would put an end to the current practice of subsidising some airlines by granting them free use of a scarce resource as a gift from the regulator. They would now have to pay scarcity prices for their slots just like any other company.

To say that incumbents would lose their *privileged* position is not the same as saying that they would lose their *dominant* position. The dominant players of today may very well remain the dominant players of tomorrow – but not simply *by virtue* of being the dominant players of today. They could retain that position, but they would have to earn it anew. The industry would become less path-dependent, as today's slot allocation would no longer determine tomorrow's.

So the argument is not that a market-based slot allocation system would be certain to lead to a radical shake-up of the industry, dethroning incumbents and boosting their competitors. That *could* be a result: newcomers *could* use the new system to buy their way into the market; and minor players *could* use it to buy extra slot space, upscale and become major players. But it need not be. The new slot allocation pattern could look very similar to the current one or, in theory, even identical. It could well be that the current outcome already resembles the one that a market system would also produce. The point of introducing a market-based allocation system is not necessarily to shake up the industry, but to test whether it needs a shake-up at all.

2. Irrespective of the impact on industry structure and competition, under a realistic pricing system, slots are much more likely to end up with whoever is best placed to make use of them, as revealed by their willingness to pay for it. The new system would be able to tap into that knowledge. This should lead to a more efficient slot usage, which has the same effect as easing capacity constraints.
3. To the extent that the new system makes it easier for airlines to relocate from one airport to another, it should fuel competition between airports.
4. Under the current system, some airlines benefit from airport capacity constraints, since this increases the value of their slots. In a market-based system, that would no longer be the case. If airlines had to pay realistic prices for their slots, they would actively *want* them to be abundant, and cheap. Their incentives would become more aligned with those of their customers. This should make them more supportive of airport expansion.
5. All sorts of anomalies would disappear. As mentioned, in the current system, mergers and acquisitions can be an indirect way of acquiring airport slots. If airlines no longer had de facto ownership of airport slots, that motive for mergers and acquisitions would disappear. Airlines and slots would no longer come as a package deal. If Airline X wants access to slots that Airline Y is currently using, a takeover of Y by X would be neither necessary nor sufficient. It would not be necessary, because it would now be possible for Airline X to purchase access to those slots directly. And it would not be sufficient, because there is no guarantee that Airline Y will still have access to those slots in the next period.

6. Clumsy tools such as use-it-or-lose-it rules would no longer be needed, and absurd practices such as ghost flights would disappear. If an airline has no reason to believe that it will be able to put a slot to good use, it would not bid for it in the first place. Once it has paid for it, it has to make good use of it in order to cover the cost.

If an airline made a mistake and purchased access to a slot that, it now turns out, it cannot put to good use, it would seek to correct that mistake as soon as it can by getting rid of that slot again, rather than unnecessarily holding on to it. And even if it wanted to hold on to it (say, because it expects conditions to improve), flying an empty or near-empty plane would not be the way to do that.

If an airline had already paid for using a slot, and then all passengers on that flight cancel, the rational thing to do would be to cancel the flight altogether, and treat the cost of the slot as a sunk cost.

7. No matter how efficient the primary slot market, market participants make mistakes, and market conditions frequently change in unforeseeable ways. A secondary slots market would probably still be needed to correct the mistakes of the initial allocation.

But it seems plausible that a functioning primary slots market would also improve the functioning of the secondary market. Once a primary market in airport slots has been established, every take-off and landing slot would have a well-established, well-publicised market price. Those prices would be the result of frequent interactions between multiple market participants, rather than of occasional backroom deals between two parties. It would therefore embody a lot more information, which could inform secondary deals.

The introduction of a market-based system does not have to be disruptive. It would not mean tearing up the EU/IATA system, and starting from scratch again. The EU/IATA system would not suddenly disappear, just as it did not suddenly disappear after Brexit. It would still be there, and it would remain the default option. But airport operators would be allowed to deviate from it.

It is conceivable that airport operators would choose not to make much use of their new freedoms. They may find that the benefits of coming up with a completely different allocation process are not big enough to justify the hassle. They may find that whatever their frustrations with the current

allocation rules, the benefits of being part of a system that is standardised across Europe (and beyond) are so large that they outweigh the gains from introducing a better system.

But even in that case, no harm would have been done. The system would have been put to the market test, and it would have passed it for now. The option of deviating from it would remain, and it could be used at some point in the future.

Second-best solutions

The meta-system described above represents the sort of blue-sky thinking one would expect from a think tank publication, which aims to explore an issue from first principles, rather than come up with a politically and practically feasible solution. It therefore gives short shrift to political constraints and legal constraints, as well as other practical difficulties.

The reality would be messier. A market-based approach of the above variety would have to settle a fundamental question first: who owns an airport slot? Do airlines have implicit property rights over their slots? Or is their right to use those slots simply an administrative convention, which could be withdrawn any time? What may sound like a simple question is, in reality, theoretically ambiguous, not to mention a legal minefield (Pickett & Hirst 2020: 6–7 & 12; Boyfield 2003: 29–34).

It does not end here. If it turns out that airlines do not have property rights over slots, what does this mean for airlines that have recently bought an expensive slot on the secondary market? Should that slot, which they have acquired in good faith, suddenly be rendered worthless? Should they be compensated, and if so, by whom?

If decisions about slot allocation were devolved to individual airport operators, would this lead to a huge increase in their market power? Could conventional competition policy deal with this situation adequately? Should airport operators be entitled to a huge windfall profit, if they suddenly had ownership rights over their slots?

To avoid these thorny issues for now, one could revert to the more conventional proposal of a general slot auction, perhaps organised by ACL, and one could initially limit it to slots that have to be newly allocated

anyway. This would mean using the auction for slots that have been either newly created, or returned to the slot pool.

There is an indirect precedent for slot auctions of this type, because they would be comparable to a spectrum auction, for which there are successful examples from around the world (and, indeed, from the UK itself). As a side effect, they could create sizeable windfall gains for the Treasury, either directly, if it is a state entity that organises the auction, or indirectly, if the proceeds are taxed. The prospect of windfall gains should make them politically more likely to happen in the first place, and increase political incentives to take a more permissive stance on airport expansion.

But while this approach would avoid political and legal difficulties, it would create a problem of a different kind: it would create a two-tier system, in which some airlines receive their slots for free while others have to pay for them. While this is not a completely new issue – we could say that due to the existence of a secondary slot market, this is already the case – it would nonetheless be a major distortion.

The reform should therefore not stop there for long. The ambition should be to extend slot pricing, and phase out grandfathering entirely. This could start with a detailed legal inquiry into the property rights situation of airport slots, in order to clarify in advance what compensation (if any) would be owed to whom, and the drafting of a compensation scheme, before any action is taken on grandfathered slots. There could, of course, be discretionary compensation payments over and above the legally required minimum, as a gesture of goodwill and/or to smoothen the transition – this would be a political choice.

This solution would not realise the full range of benefits of a market-based approach. But it would still be a considerable net improvement. It could still lead to a more efficient allocation of slots, because, as the CMA (2019: 22) puts it:

All of the information that an administrator would have to assess would instead be captured in a price that would truly reflect the value that an airline places on the slot. The auction extracts and uses information otherwise unavailable to an administrator.

It could still make it easier for newcomers to enter the market, and for minor players to upscale. It could still increase airlines' 'switching rates' between airports, thus boosting competition.

An alternative second-best option would be the above-described system of congestion charging, which would deliver similar benefits as an auction. Under this system, it would not be necessary to answer the thorny question of who owns a slot, because incumbents would not technically 'lose' their slots. But they could be priced out of using them: what they lose is not the slot, but the right to use it for free. It would be more like an economic eviction, where tenants do have the option of extending their rental agreement, but where the rent rises to a level that exceeds their willingness to pay. Pricing out an incumbent could benefit a competitor.

Congestion charging would avoid the above-described problem of a two-tier system. Once up and running, it could be applied to all slots, not just newly created ones, so it would not matter whether an airline had only just entered the market, or whether it had held a particular slot for decades. Such charging would also raise a substantial amount of revenue, which could be shared between airport operators, and the Treasury.

Of course, the devil will be in the detail. The purpose of this paper was not to present a fully worked-out solution that could become legislation tomorrow. The point of this paper was to make it clear that the liberalisation of air travel in the UK is best thought of as an unfinished success story. Allowing the market into the remaining islands of central planning would be that success story's logical continuation.

References

Bichler, M.; Cramton, P.; Gritzmam, P. and Ockenfels, A. (2021) It is time to auction slots at congested airports. VoxEU & CEPR (<https://voxeu.org/article/it-time-auction-slots-congested-airports>)

Boyfield, K. (ed.) (2003) *A Market in Airport Slots*. London: Institute of Economic Affairs.

CAA (2019) 2018 passenger survey report. Civil Aviation Authority (<https://www.caa.co.uk/Data-and-analysis/UK-aviation-market/Consumer-research/Departing-passenger-survey/2018-Passenger-survey-report/>)

CAA (2021) Table 09 terminal and transit passengers. Annual airport data 2019. Airport related statistics for 2019. UK Civil Aviation Authority statistics.

CAA & ComRes (2019) UK aviation consumer survey June 2019. Civil Aviation Authority (https://publicapps.caa.co.uk/docs/33/CAP1831%20ComRes_CAA_UKACR_Wave%207_full%20report_v2.pdf)

CMA (2018) Advice for the Department for Transport on competition impacts of airport slot allocation. Competition and Markets Authority.

CMA (2019) Aviation 2050. Response from the Competition and Markets Authority. Competition and Markets Authority.

DfT (2003) The future of air transport. White Paper. London: Department for Transport.

DfT (2019) Traffic at major airports, worldwide. Aviation statistics: data tables. London: Department for Transport ()

DfT (2021) Air traffic, United Kingdom airports: 1950 to 2020. Table TSGB0201 (AVI0101). Department for Transport statistics.

Donohue, G. and Hoffman, K. (2007) Evidence that airport pricing works. Policy Brief 67. Los Angeles: Reason Foundation.

Eurocontrol (2022) Charting the European Aviation recovery: 2021 COVID-19 impacts and 2022 outlook. Think Paper 15. Aviation Intelligence Unit.

Fraport AG (2022) Frankfurt airport air traffic statistics (https://www.fraport.com/content/dam/fraport-company/documents/investoren/eng/aviation-statistics/Air%20Traffic%20Statistics%202021.pdf/jcr_content/renditions/original./Air%20Traffic%20Statistics%202021.pdf)

Gillen, D. and Starkie, D. (2016) EU slot policy at congested hubs, and incentives to add capacity. *Journal of Transport Economics and Policy* 50(2): 151–63.

House of Commons Transport Committee (2018) Oral evidence: Airports National Policy Statement, HC 548.

ICAO (2016) Air transport liberalization and the economic development of the countries. Working Paper. Agenda Item 39: Economic Regulation of International Air Transport –Policy. Presented by the Air Transport Research Society. International Civil Aviation Organization.

ITF (2019) Liberalisation of air transport. ITF Research Reports. Paris: OECD Publishing ()

Lesh, M. (2019) Ready for take-off: Building competition in the aviation industry. Briefing Paper. London: Adam Smith Institute.

Niemietz, K. (2013) Depoliticising airport expansion: Market-oriented responses to the global and local externalities of aviation. IEA Discussion Paper 51. London: Institute of Economic Affairs.

Pheasant, J. and Giles, M. (2007) Slot trading in the EU. *Global Competition Review*, November: 30–33.

Pickett, L. and Hirst, D. (2020) Airport slots. Briefing Paper 9062. House of Commons Library.

WPI Economics (2019a) Ticket to fly: Using the allocation of new capacity to maximise the benefits of Heathrow expansion. (http://wpieconomics.com/site/wp-content/uploads/2019/05/19009_Virgin-Atlantic-Ticket-to-Fly-A4_02h-web-1.pdf)

WPI Economics (2019b) Letting competition fly: The case for two national flag carriers. (<http://wpieconomics.com/site/wp-content/uploads/2019/09/110919-FINAL-letting-competition-fly.pdf>)

The Institute of Economic Affairs
2 Lord North Street
London SW1P 3LB
Tel 020 7799 8900
email iea@iea.org.uk



The logo for the Institute of Economic Affairs (iea) features the lowercase letters 'iea' in a white serif font. A small red dot is positioned above the letter 'i'. Below the logo, the full name 'Institute of Economic Affairs' is written in a smaller, white, sans-serif font.

Institute of
Economic Affairs