Slot allocation at an expanded Heathrow

Behavioural experiment

Prepared for
Department for Transport

7 October 2019

www.oxera.com

oxera
compelling economics

Contents

Figures and tables

Executive summary

Oxera has been commissioned by the Department for Transport (DfT) to conduct a behavioural experiment exploring alternative mechanisms for slot allocation. This is in the context of the release of a significant amount of new capacity with the building of a third runway at Heathrow Airport ('Heathrow').

In June 2018, the government approved the building of a third runway at Heathrow.[1] Provisional figures from Heathrow Airport Limited's (HAL) Master Plan indicate that this new runway—the North West Runway—will lead to an additional 276,000 air transport movements (ATMs) each year, an increase of almost 58% relative to the existing cap of 480,000 ATMs.[2]

The purpose of the behavioural experiment ('the experiment') is to explore empirically the impact of different slot allocation mechanisms on the government's objectives, which are to: maximise efficiency through effective competition and connectivity that best meets the needs of consumers; and use the new capacity to safeguard domestic connectivity.

Context

The slot coordinator for Heathrow is Airport Coordination Limited (ACL), which operates under the relevant UK and EU legislation. ACL uses an administrative mechanism to allocate slots to airlines, in which it decides between airline slot requests using a number of criteria, including the following.

- **Grandfather rights**. Where an airline already holds a slot, it may continue to hold that slot in future seasons if it has used the slot sufficiently.
- **The new entrant rule**. Where there is a 'pool' of new slots to be allocated (i.e. slots for which no airline is exercising grandfather rights), 50% of the new slots should be allocated to new entrants.
- **Additional criteria**, including:
  - **requirements of the travelling public and other users**. ACL reviews whether there is increasing consumer demand for flights to a country or city already served, or for flights to a new destination that is not currently served by the airport;
  - **competition**. ACL focuses on competition at a route level (i.e. airport to airport). For example, if there is a particular route flown by only one airline, and another airline proposes to also fly that route, this would be viewed positively by ACL when considering the airline's request.

The experiment empirically tested the impact of changes to the current slot allocation mechanism.

The experiment

We conducted a behavioural experiment in order to determine the impact of different slot allocation mechanisms. The use of behavioural experiments is common in policy design. The controlled experiment context allows us to robustly identify the impact of different policies, controlling for other factors. The experiment was designed in accordance with academic best practice, in

[1] Department for Transport (2018), 'Government sets out next steps for Heathrow expansion', 26 June. A DCO (Development Consent Order) application has yet to be approved.
[2] Heathrow Airport Limited (2019), 'Heathrow Expansion: Preliminary Environmental Information Report', volume 1, chapter 6: 'DCO Project description', Table 6.14.

collaboration with the Centre for Decision Research and Experimental Economics (CeDEx) at the University of Nottingham.

The design of a behavioural experiment requires a certain level of simplification from reality (as participants must be able to make their decisions in a limited timeframe). However, complexity was incorporated wherever necessary in order to capture the key dynamics of airlines' long-term strategies for slot allocation and use. The experiment therefore provides robust insight into the outcomes associated with the release of new slots at Heathrow, while deliberately not capturing every possible level of granularity in its design. In other words, the experiment design was selected to give insight into the key policy design decisions, rather than to provide a complete forecast of all airline behaviour over the next few decades.

There are three main ways in which behavioural experiments can be undertaken—in a laboratory, online, and in the field.[3] The choice among these options depends on the nature of the research question, with different types of experiment suited to different scenarios. We decided to conduct this experiment in a laboratory environment.

There is greater control over factors that may influence decision-making in laboratory experiments than in online experiments. In a laboratory experiment, the experimenter can ensure that there is no cheating or communication among participants, and that the participants are paying full attention. The experimenter can also ensure that the conditions in which the experiment takes place are identical for each participant and for each session. Furthermore, the experimenter can check the identity of participants to ensure that those who take part are those who have been invited.[4]

We note that laboratory experiments can be less 'realistic' than field experiments, as the participants know that they are being observed and are making 'as if' decisions (rather than decisions for real). However, field trials are not always possible, and in any case it is often desirable to conduct testing in a controlled environment before proceeding to a field trial. For example, the results of a laboratory experiment can identify risks associated with different policy options, and can help to refine the policies to be tested in a field trial.

The experiment was conducted with students in the laboratory at CeDEx at the University of Nottingham. The use of students as participants is common in behavioural experiments.[5] One alternative to using students would have been to use industry professionals or experts. Apart from the practical issues that could have arisen in achieving a large enough sample size to obtain statistically significant results, the use of industry professionals/experts would have resulted in methodological concerns. For example, such participants may have had prior beliefs about an airline's preferences for slots, which are likely to have differed between participants in the different treatments, and could therefore have undermined the results of the experiment.

The design of the experiment was informed by research and analysis to understand how airlines might behave when requesting slots for Heathrow's

[3] In laboratory and online experiments, participants are asked to behave 'as if' the decisions were taken 'for real', whereas field experiments involve observing real behaviour (without participants necessarily being aware of the experimentation).
[4] In a laboratory experiment the experimenter can answer questions from participants if they arise, and can hand out hard copies of the instructions.
[5] For a review, see Brandts, J. and Potters, J. (2018), 'Experimental industrial organisation', in L.C. Corchón and M.A. Marini (eds), *Handbook of Game Theory and Industrial Organisation, Volume II*, Edward Elgar Publishing.

third runway. We also conducted extensive stakeholder engagement, including with public sector bodies, private sector firms, industry experts and a number of airlines.

This analysis informed the experiment design in terms of:

- representative airlines;
- airline preferences and payoffs from the actions taken in the experiment;
- airline constraints in the usage and acquisition of slots;
- the number of slots and time periods;
- representative routes.

Treatments

Behavioural experiments estimate the impact of different factors by changing one factor at a time relative to a baseline scenario, which ensures that different effects cannot be conflated. The baseline scenario and each variation are referred to as a 'treatment'. Each treatment varies (from another treatment) in terms of only one factor. Any given factor in the experiment that is held constant across the treatments should not affect the results.

The experiment involved four treatments to test the impact of three different changes to the status quo slot allocation mechanism:

- **Treatment 1: baseline (or 'control')**. The baseline treatment was designed to be as close as possible to the status quo for slot allocation in the UK;
- **Treatment 2: limited duration rights**. Treatment 2 tested the impact of granting slots for a limited period of time only, before the slots needed to be handed back to the slot coordinator—i.e. removing grandfathering rights;
- **Treatment 3: no new entrant rule**. Treatment 3 tested the impact of removing the new entrant rule that is in place in the baseline treatment;
- **Treatment 4: auction**. Treatment 4 tested the impact of using an auction instead of the administrative mechanism. There are multiple ways to implement an auction, and in this treatment we considered a 'combinatorial auction' where participants were able to create a package of slots.

The experiment focused on the allocation and use of the new slots at Heathrow with the third runway, rather than the allocation and use of the existing slots.[6] Thus the results should be interpreted primarily as insight into the allocation of new slots, and therefore the marginal impact of the new slots on the status quo.

Summary of results

In summary, we find that the no new entrant rule and auction treatments increase route-level competition. The auction is the most effective way of increasing allocative efficiency, but reduces the share of slots held by new entrants. While trading increases efficiency, it does not fully mitigate the effects of an inefficient starting allocation.

[6] The experiment data cannot be linked to actual data at the level of granularity of airlines and destinations, because airlines in the experiment cannot be linked to specific airlines in reality, and destinations in the experiment do not capture the full range of destinations currently offered from Heathrow.

In more detail, we find the following.

- Both treatment 3 (no new entrant rule) and treatment 4 (auction) increase route-level competition (the primary measure of competition considered). However, treatment 4 reduces the share of slots held by new entrant airlines. In treatment 2 (limited duration), the level of competition is similar to in treatment 1 (baseline), but the market share of new entrants is the highest.
- Given the mix of routes operated by airlines, the effect of competition from the release of new slots is greatest on long-haul routes.
- We also analyse allocative efficiency (i.e. whether capacity is allocated in a way that maximises social welfare) as the sum of producer surplus and consumer surplus. Trading improves allocative efficiency, although it does not fully mitigate inefficiencies that can arise in the current allocation mechanism. While auction and the baseline provide the same opportunities for trading, the auction results in higher allocative efficiency. This is because the auction allocates slots to airlines that are best placed to meet passenger demand, and therefore have the largest incentive to bid for slots. Treatment 3 (no new entrant rule) also leads to higher allocative efficiency than treatment 1 (baseline), as it allows slots that would otherwise be allocated to new entrants to be allocated more efficiently—although it is still less efficient than the auction treatment. Treatment 2 (limited duration) leads to a relatively inefficient outcome as there are fewer opportunities for trading.
- We also look at productive efficiency (i.e. whether the airport infrastructure is fully utilised) by considering the share of long-haul flights and the average seat-km per flight. Productive efficiency is lowest in treatment 3 (no new entrant rule) due to a change in the mix of airlines, with fewer Middle Eastern and Asian airlines acquiring slots. On the other hand, treatment 3 (no new entrant rule) increases the number of domestic flights.
- Airlines alter their route choice after winning slots where there is an administrative allocation mechanism (treatments 1, 2 and 3), but such behaviour may be less common in reality as airlines are likely to want to avoid harming their relationships with the slot coordinator.

The key results of each treatment as compared to the baseline are summarised in the table below.

Key results: comparing treatments 2, 3 and 4 to the baseline

As compared to the baseline (treatment 1)		Treatment 2: limited duration rights	Treatment 3: no new entrant rule	Treatment 4: auction
Competition	Route-level HHI	—	↑↑	↑†
	Airport-level HHI	—	↑	—
	Share of slots held by new entrants	—	—	↓↓
Efficiency	Allocative efficiency (producer surplus)	↓	↑	↑↑
	Allocative efficiency (consumer surplus)*	—	↑	↑
	Productive efficiency (average-seat km per flight)	—	↓	↓
Connectivity	Share of long haul flights	—	↓	—
	Share of domestic flights	↑	↑	↑
	Number of flights on new routes	—	↓	↓↓

Note: '↑↑' symbolises a large and statistically significant difference relative to the baseline. '↑' symbolises a moderate and statistically significant difference relative to the baseline. '—' symbolises no difference to the baseline (i.e. not statistically significant at the 5% level using the Mann–Whitney test). '↓' symbolises a moderate and statistically significant negative difference relative to the baseline. '↓↓' symbolises a large and statistically significant negative difference relative to the baseline. * Consumer surplus calculated (conservatively) as a robustness check, confirming the direction of results, but the consumer surplus results alone are not statistically significant. † Significant at the 5.8% level.

Source: Oxera.

1 Introduction

Oxera has been commissioned by the Department for Transport (DfT) to conduct a behavioural experiment exploring alternative mechanisms for slot allocation. This is in the context of the release of a significant amount of new capacity with the building of a third runway at Heathrow Airport.

In June 2018, the government approved the building of a third runway at Heathrow Airport ('Heathrow').[7] Provisional figures from Heathrow Airport Limited's ('HAL') Master Plan indicate that this new runway—the North West Runway—will lead to an additional 276,000 air transport movements (ATMs) each year, an increase of almost 58% relative to the existing cap of 480,000 ATMs.[8]

The purpose of this behavioural experiment ('the experiment') was to consider how different slot allocation methods affect the distribution of slots among airlines and therefore how they meet the government's objectives, which are to maximise efficiency through effective competition and connectivity that best meets the needs of consumers, and to use the new capacity to safeguard domestic connectivity.

The experiment was conducted at the Centre for Decision Research and Experimental Economics at the University of Nottingham.

The experiment, and this report, are focused on empirically exploring the impact of different slot allocation mechanisms on the government's objectives, in order to contribute to the DfT's White Paper. The report does not undertake an evaluation of the merits of the current slot allocation system and does not provide recommendations on the most appropriate form of slot allocation mechanism going forward.

This report is structured as follows:

- section 2 describes the current slot allocation mechanism at Heathrow;
- section 3 analyses the preferences of airlines, and the constraints they face, when requesting slots;
- section 4 sets out the design of the laboratory environment;
- section 5 describes the slot allocation mechanisms in the experiment;
- section 6 outlines the experiment treatments;
- section 7 sets out the results.

The appendices to this report are structured as follows:

- Appendix A1 lists the stakeholders that we engaged with. The design of the experiment was informed by an extensive stakeholder engagement process.
- Appendix A2 provides screenshots of the experiment;
- Appendix A3 reviews the relevant literature on slot allocation mechanisms and auction design;

[7] Department for Transport (2018), 'Government sets out next steps for Heathrow expansion', 26 June.
[8] Heathrow Airport Limited (2019), 'Heathrow Expansion: Preliminary Environmental Information Report', volume 1, chapter 6: 'DCO Project description', Table 6.14.

- Appendix A4 details descriptive statistics about the participants;
- Appendix A5 sets out further detail on the consumer surplus analysis.

Some information in this report has been redacted to ensure that confidential information is not disclosed. This is indicated by [✂].

2 The current slot allocation mechanism at Heathrow

It is important to understand the current slot allocation mechanism in order to determine the appropriate experiment design. This section considers the current regulation that applies to slot allocation in the EU, and the process that is undertaken by the UK slot coordinator to evaluate competing requests. It also briefly considers slot allocation in the context of a large release of new capacity at Heathrow.[9]

2.1 Regulatory context

Slot allocation in the EU is subject to Regulation 95/93, which applies the Worldwide Slot Guidelines (WSG) as defined by the IATA (International Air Transport Association).[10] While the WSG are reviewed on a regular basis (the most recent version came into effect in August 2019), Regulation 95/93 has remained unchanged since 1993.

The slot coordination process differs depending on the capacity constraints at an airport (i.e. whether demand for slots exceeds supply). An airport where demand consistently exceeds supply is categorised as a 'Level 3' airport. These airports require the slot coordinator to make decisions regarding which airlines should be granted the right to use which slots.[11]

Slots are allocated for two seasons (Winter and Summer), approximately six months in advance of the season. A slot is defined in the WSG as:[12]

> a permission given by a coordinator for a planned operation to use the full range of airport infrastructure necessary to arrive or depart at a Level 3 airport on a specific date and time.

Airlines require pairs of slots (for arrival and departure at an airport).[13]

Regulation 95/93 sets out two principles that the slot coordinator must consider when allocating slots.[14]

- **Grandfather rights (first priority)**. Where an airline already holds a slot, it may continue to hold that slot in future seasons if it has used the slot sufficiently—i.e. if it meets the 'use it or lose it' (UIOLI) criteria (set at 80%). For example, if an airline used a slot 80% of the time in the Winter season of 2018–19, it would retain the right to use that slot in the following Winter season of 2019–20. Thus grandfather rights refer to both the granting of rights in perpetuity and the UIOLI rule.
- **New entrant rule (second priority)**. Where there is a 'pool' of new slots to be allocated (i.e. slots for which no airline is exercising grandfather rights),

[9] The UK is set to leave the EU before the construction of the new runway. At this stage, it is not clear which regulatory frameworks will apply when slot allocation takes place.
[10] Council Regulation (EEC) No 95/93 of 18 January 1993 on common rules for the allocation of slots and Community airports. IATA (2019), 'Worldwide Slot Guidelines', effective 1 August 2019, 10th Edition.
[11] Although the slot coordinator also makes decisions at Level 2 airports, congestion is greater at Level 3 airports. ACL (2019), 'FAQs', webpage. The Level 3 airports in the UK are: Birmingham, Gatwick, Heathrow, London City, Luton, Manchester and Stanstead. A Level 2 airport is a Slot Facilitated Airport where demand is close to capacity and coordination is needed to ensure the demand does not exceed the capacity. A Level 3 airport is a Fully Coordinated Airport; where demand exceeds the capacity and due to the high volume of demand, there is no opportunity to re-schedule flights on a voluntary basis.
[12] IATA (2019), 'Worldwide Slot Guidelines', effective 1 August 2019, 10th Edition, p. 42.
[13] Slots are allocated in five-minute intervals by ACL at all airports in the UK, but the flexibility around this interval varies according to the level of congestion at the airport in question.
[14] Council Regulation (EEC) No 95/93 of 18 January 1993 on common rules for the allocation of slots and Community airports, Articles 2, 8, 10.

50% of the new slots should be allocated to new entrants. New entrants are defined as airlines that:

- hold less than 3% of all the slots available on that day at that airport, or less than 2% of all the slots available on that day at that airport system (i.e. group of airports serving a city), and:
  - hold (or have been allocated) fewer than four slots at that airport on that day; or
  - where the airline will operate direct flights to another EU airport (or airport system) where no more than two other airlines operate direct flights on that route, and hold (or have been allocated) fewer than four slots at that airport on that day and on that particular route.

Slots allocated to a new entrant cannot be traded or leased to another airline for two years. New entrants are defined at the airline level, meaning that new entrant slots cannot be transferred between airlines that are part of a broader group.

Further guidance may be required for the slot coordinator to allocate specific slots. The WSG provide the following high-level guidance.[15]

- **Additional criteria (third priority)**. On a case-by-case basis, the slot coordinator may consider the following when deciding which airline to award a slot to:
  - **effective period of operation**—airlines that require the slots for more weeks within the season are given preference;
  - **curfews**—priority is given to an airline whose schedule is constrained by a curfew at the airport at the other end of the route;
  - **time spent on waitlist**—requests that are pending on the waitlist should take priority;
  - **type of service and market**—different types of services (scheduled, charter and cargo) and markets (domestic, short-haul and long-haul) should be balanced, and the airport route network should be developed;
  - **competition**—competitive factors should be considered;
  - **requirements of the travelling public and other users**—the needs of consumers and shippers should be met;
  - **frequency of operation**—having more flights per week should not in itself imply a higher priority for slot allocation;
  - **local guidelines**—the slot coordinator should take account of any local guidelines.

The exact interpretation of the additional criteria is not specified in the WSG (e.g. it is unclear how competitive factors are evaluated). In the next section we describe how ACL follows these criteria in the allocation of slots at UK airports.

[15] IATA (2019), '[Worldwide Slot Guidelines]', effective 1 August 2019, 10th Edition, p. 67.

2.2 Slot allocation at Level 3 airports in the UK

The slot coordinator for all UK airports is Airport Coordination Limited (ACL), an independent company (limited by guarantee) operating under the relevant UK and EU legislation.

ACL follows Regulation 95/93 and the WSG at all UK airports, with the most complex decisions over slot allocation occurring at Level 3 airports. Our understanding of ACL's process, based on discussions with ACL staff, is shown in Figure 2.1.

Figure 2.1 Slot allocation process at Level 3 airports

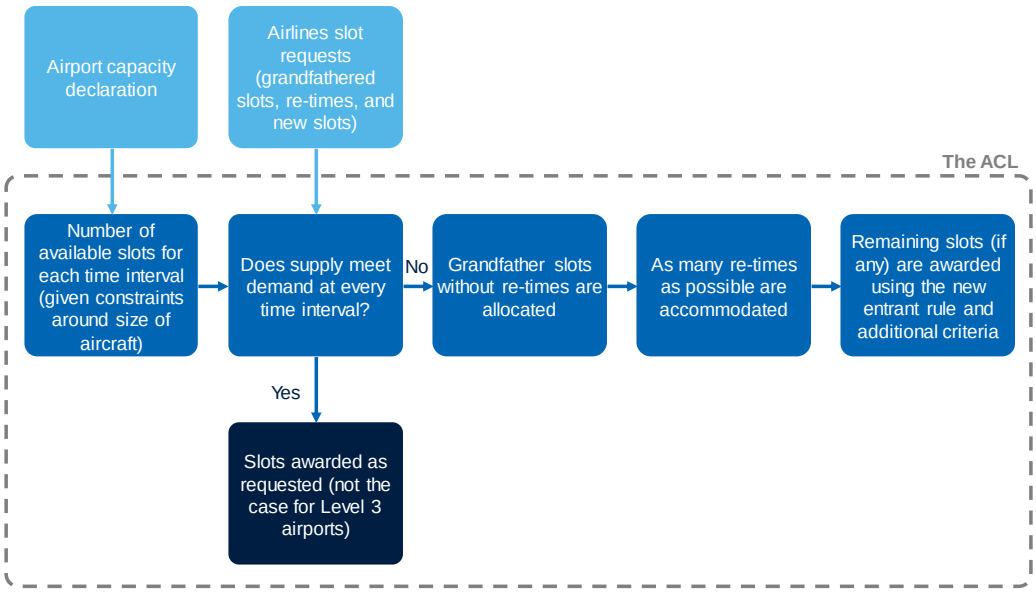

Note: 'Re-times' refer to requests by airlines to change the timing of slots that they already hold.

Source: Oxera.

We note that the WSG has recently agreed to make changes to the re-time priority, whereby requests to re-time are treated in the same way as requests for new slots. This would limit an airline's ability to automatically re-time slots when new capacity becomes available. These changes may be included in the next version of Regulation 95/93 and may be implemented by ACL.

The steps in the slot allocation process at Level 3 airports are described further below.

2.2.1 Capacity declaration and calculation

At the start of each scheduling process, undertaken six months in advance of each season, the airport makes a capacity declaration regarding the total capacity of each terminal at the airport.[16] This takes account of a number of constraints:

- the runway, in terms of the maximum number of arrivals and departures (and total ATMs) across the day;
- any banned aircraft, as an airport may disallow certain types of aircraft;

[16] For example, see Heathrow's capacity declaration for Summer 2018: https://www.acl-uk.org/wp-content/uploads/2017/09/S18-Declaration-Letter-Appendices-for-ACL-2.pdf.

- parking by terminal, in terms of physical stands for aircraft;
- passenger flow by terminal, with the airport providing ACL with gate capacity and average load factors to apply in the capacity calculation;
- baggage handling capacity by terminal;
- any specific constraints relating to particular types of aircraft, as some very large aircraft require different terminal equipment, of which there may be a limited supply at the airport.

ACL uses these constraints to build a picture of the overall supply of slots at the airport, by terminal.

2.2.2 Matching supply and demand

ACL then receives slot requests from airlines. Each slot request includes the time of the slot, the route, and the aircraft type. Airlines state their ideal time for a slot, and a window around that time that they would be willing to accept.

There are three types of slot request: grandfather slots without re-time requests (at the exact time as in the previous season); grandfather rights with re-time requests (a requested change to the timing or aircraft); and new slot requests.

Having received the slot requests, ACL calculates whether supply meets demand. Given the airlines' stated aircraft, ACL also calculates whether any of the constraints are binding. Supply does not meet demand at all times of the day at Level 3 airports, meaning that ACL proceeds to the next step.

Grandfather rights and re-times

Starting with grandfather slots (without re-time requests), ACL builds a picture of the demand for slots at each terminal and for each time interval. ACL then calculates how many re-times can be accommodated. If a re-time cannot be accommodated (because there is no capacity at the time requested, or because there is no capacity to increase the size of the aircraft), the airline receives the grandfathered slot without amendments. It is our understanding that a significant proportion of grandfathered slots are 're-timed' each season.

ACL completes this process manually and engages with airlines at early stages of the process to ensure that it has a good understanding of airline requests. We understand that, where possible, ACL considers the turnaround time between arrival and departure of the aircraft at the airport.[17]

If, after the re-times have been accommodated, there is still capacity at the airport, ACL proceeds to the next step.

New slots (new entrant rule and additional criteria)

ACL awards 50% of the new slots to new entrants (as defined above). Where there is still excess demand for slots, ACL applies the additional criteria. It is our understanding that ACL applies the additional criteria on a case-by-case basis, reviewing the merits of the different slot requests. The relative weight of the different factors is not defined and may vary.

[17] Airlines request slot pairs and therefore reveal the preferred turnaround time. ACL considers the preferred turnaround times, such that an airline is not allocated a slot pair that it would struggle to service.

It is also our understanding that ACL interprets certain additional criteria set out by the WSG as follows.

- **Requirements of the travelling public and other users**. ACL reviews whether there is increasing consumer demand for flights to a country or city already served, or for flights to a new destination that is not currently served by the airport. For example, if an airline proposes to fly to a new city then this would be looked upon favourably by ACL. Alternatively, if an airline proposes to fly a route where there is increasing consumer demand, this would also viewed favourably.
- **Competition**. ACL focuses on competition at a route level (i.e. airport to airport). For example, if there is a particular route flown by only one airline, and another airline proposes to also fly the route, this would be viewed positively by ACL in considering the airline's request. Alternatively, if two airlines fly a particular route, and the airline with the lower number of weekly flights on that route proposes increasing its flights, this would be viewed as a positive factor in ACL's consideration.[18]

When there are two requests for a slot, of which one is for a new route and the other would add competition to an existing route (all else being equal), ACL decides on a case-by-case basis.

Airlines are not bound to fly to the destination that they specify in their requests to ACL.[19] The only exception is where an airline qualifies for new entrant status and obtains a new entrant slot on the basis of the route that it would offer. In this case ACL has the power to restrict the airline from altering the route.

It is also our understanding that, if an airline were leasing slots at a particular time to other airlines, and that airline requested more slots at that time, ACL would be less likely to grant this request. However, ACL may not know whether a transfer is permanent or a lease, and therefore it may be difficult to apply this strategy.

2.2.3 Slot trading

Once slots are allocated by ACL, airlines are allowed to trade or lease their slots to other airlines.[20] As shown in Figure 2.2, a number of slot transfers occur at Heathrow each year.

[18] Similarly, the UK Competition and Markets Authority (CMA) has stated: 'Our understanding is that ACL's current assessment of competition is to give priority to allocating slots to airlines that want to enter routes that are poorly served, ie routes with the fewest number of carriers or unserved routes.' Competition and Markets Authority (2019), 'Aviation 2050 Response from the Competition and Markets Authority', June, para. 4.20. We also understand that ACL may consider whether there is evidence of sufficient consumer demand to justify the extra operations on that route.
[19] While any change of route (or aircraft) requires ACL approval, it can currently refuse a request only on the basis of feasibility at the airport.
[20] Slot trades are technically 'swaps' of slots between airlines, although we understand that an airline that does not want to (or cannot) give up an existing slot can gain (and then swap) slots from ACL that hold no value. These slots are also known as 'inoperable slots' as there is no ATM attached.

Figure 2.2 Number of slot transfers at Heathrow (Summer season)

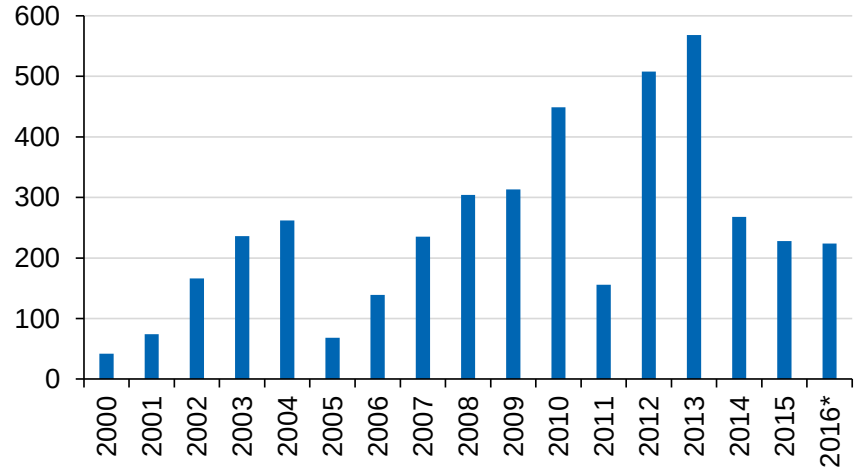

Note: Each slot transfer refers to one arrival or departure per week. It is our understanding that the swap of one slot for another is counted in this data as one slot transfer. Data refers to Summer season slots only.

Source: ACL (2016), 'ACL Slot Coordination', presentation, September.

2.3 Application of the current slot allocation mechanism to a large release of new slots

We note that there may be operational amendments to the ACL process and timings of slot allocation in the case of a large release of new slots. However, there is at least one example of where a slot coordinator in Europe has followed its existing slot allocation mechanism for a large release of new slots (Frankfurt Airport in 2011).[21]

If the current rules are applied to a large release of new slots (such as with Heathrow's new runway), the pool of new slots will be allocated according to the three priorities identified above: re-times of grandfathered slots; the new entrant rule; and the additional criteria.

In terms of timing, the third runway ATMs will not be released all at once. Instead, Heathrow is consulting on phasing in slots on a yearly basis. It aims to release an initial 5,000 ATMs in 2022 and to achieve the full release by 2050, as shown in Table 2.1.

[21] Steer Davies Gleave (2011), 'Impact assessment of revisions to Regulation 95/93', March.

Table 2.1 Forecast phasing of extra capacity at Heathrow

Year	Total ATMs	Additional ATMs
2022	485,000	5,000
2023	495,500	10,500
2024	500,000	4,500
2025	505,000	5,000
2026	505,000	0
2027	567,000	62,000
2028	607,500	40,500
2029	631,500	24,000
2030	665,000	33,500
2031	685,000	20,000
2032	700,000	15,000
2033	715,000	15,000
2034	730,000	15,000
2035	740,000	10,000
2036	742,000	2,000
2037	744,000	2,000
2038	746,000	2,000
2039	748,000	2,000
2040	750,000	2,000
2050	756,000	6,000
Total		**276,000**

Source: Heathrow Airport Limited (2019), 'Heathrow Expansion: Preliminary Environmental Information Report', volume 1, chapter 6: 'DCO Project description', Table 6.14.

2.4 Conclusion

In this section we considered the current regulation that applies to slot allocation in the EU, and the process that is undertaken by the UK slot coordinator to evaluate competing requests. Grandfathering rights, the new entrant rule and the additional criteria which are set out in Regulation 95/93 are key elements that feed into our experiment.

In the next section, we present the results of our research on airline slot preferences, which also guided our experiment design.

3 Airline interests in slots at an expanded Heathrow

This section sets out the research and analysis undertaken to understand how airlines may behave when applying for slots for Heathrow's third runway, in order to inform the experiment design.

First, in section 3.1, we consider the range of potential airlines that may be interested in obtaining slots at Heathrow, in order to incorporate a representative set of airlines in the experiment.

Second, in sections 3.2 and 3.3, we consider airlines' preferences for slots, in terms of timing and routes, as well as the (financial and non-financial) constraints they face when requesting slots. This information determines the objective function of the airlines, which was communicated to participants in the experiment to ensure that they behaved in a similar way to how airlines would behave in a real-world slot allocation scenario at Heathrow.

Finally, in section 3.4, we analyse airlines' aircraft usage in order to be able to make inferences regarding the number of passengers per slot in the experiment.

The data sources and information that we have relied on are detailed in Appendix A1. We have also conducted extensive stakeholder engagement, including with public sector bodies, private sector firms, industry experts and a number of airlines. The results of this stakeholder engagement inform our findings in this section of the report. The list of stakeholders we engaged with can be found in Appendix A1.

3.1 Types of airline

In order to simulate requests for slots by airlines in the experiment, it is necessary to understand the overall demand for slots at Heathrow and which airlines are likely to be interested in operating at Heathrow following the expansion.

Currently, 81 airlines operate at Heathrow. The largest airline is British Airways, which holds over half the slots at the airport. Other airlines that hold a significant number of slots are Aer Lingus, American Airlines, Lufthansa, and Virgin Atlantic. Figure 3.1 below shows the percentages of slots held by the largest five airlines at the airport.

Figure 3.1 Slot holdings at Heathrow (%), Summer 2018

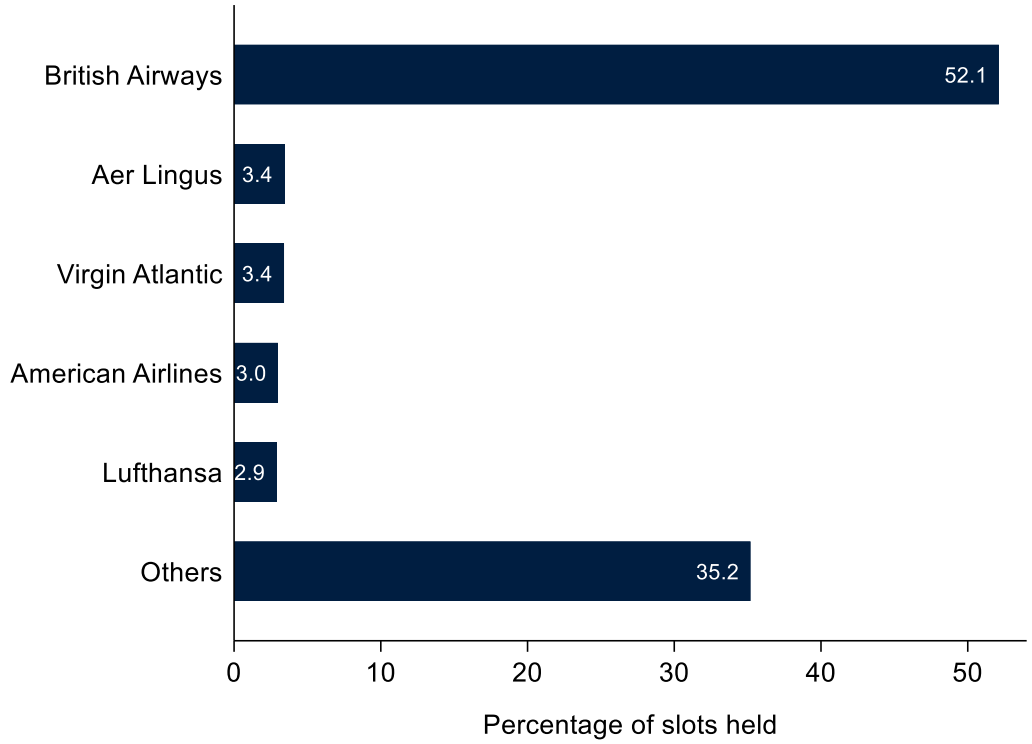

Note: Aer Lingus and British Airways are part of the IAG group. Some of these slots may be leased rather than owned, and/or there may be additional slots that airlines hold which are leased to other parties.

Source: Oxera analysis of OAG data.

Given the capacity constraints at Heathrow, a number of airlines currently operating at the airport may be interested in acquiring additional slots once the new runway is built. In addition, there may be airlines that do not currently operate from Heathrow but which would be interested in doing so after the release of additional slots. For example, only a handful of low-cost carriers (LCCs) currently operate at Heathrow. However, given the growing presence of LCCs in Europe and at other London airports in the last decade, there may be a number of LCCs that would be interested in operating at Heathrow. The LCCs with the largest slot holdings at other London airports, but which do not hold slots at Heathrow, are:[22]

- Ryanair and Jet2 at Stansted Airport;
- easyJet and Norwegian at Gatwick Airport;
- easyJet and Wizz Air at Luton Airport.

Some of these airlines also operate from other large European hub airports (e.g. easyJet operates at Paris CDG and Amsterdam Schiphol airports).

Figure 3.2 below sets out the DFT's passenger demand forecast for Heathrow with the North West Runway. The figure provides an indication of the likely passenger demand for routes to different regions. Most current demand at Heathrow is for travel within Europe or to other OECD countries. However, there is forecast to be increasing demand for travel to many regions of the world, including the Far East and the Middle East.

[22] As indicated by the ACL Summer 2019 Start of Season report for each airport.

Figure 3.2 Constrained passenger demand forecast from Heathrow to different regions (with third runway)

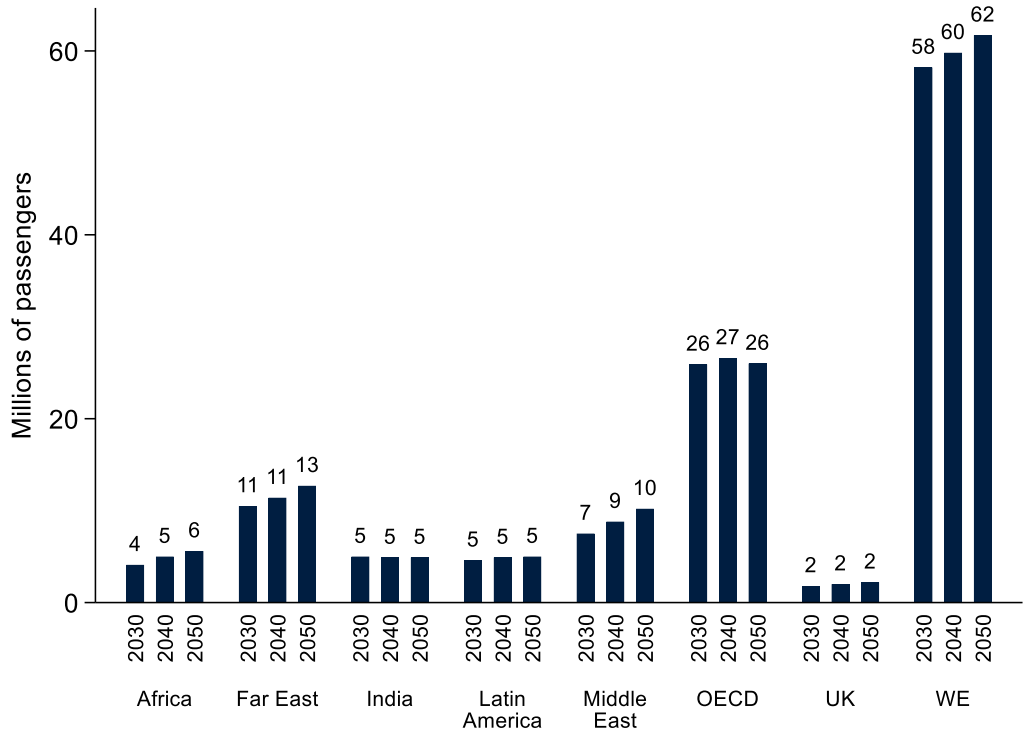

Note: Following the DfT's NAPDM forecasting methodology, WE (Western Europe) refers to all EU and non-EU countries in Europe (including Russia). OECD refers to OECD countries outside Europe (the USA, Canada, Mexico, Japan, New Zealand and Australia).

Source: Department for Transport (2017), 'UK Aviation Forecasts', October.

The forecasts indicate that demand for flights to India and Latin America is not expected to grow significantly and will continue to constitute a small proportion of demand for flights at Heathrow. While the forecasts indicate that the demand for flights to Africa will grow significantly in percentage terms, the demand for flights to Africa will remain a small proportion of total flights. Given the limited demand for flights, and the fact that a number of routes are already offered to these destinations from Heathrow, it is unlikely that African, Indian and Latin American carriers will require a significant number of new slots at Heathrow.

Other types of airline such as freight and charter airlines have indicated some interest in obtaining slots at Heathrow, although on a much smaller scale than traditional carriers. [✂] In addition, the business model operated by freight carriers results in a preference for slot times during the night, which is not possible at Heathrow due to night flying restrictions.

In addition to the above, we have considered a number of other features of airlines that may affect their desire and ability to obtain slots at an expanded Heathrow:

- **business model**—whether the airline is an LCC or a full-service carrier;
- **flight duration**—short-haul long-haul, or both;
- **flight destination**—e.g. Europe, North America, the Middle East;
- **preferences for slot times**—patterns of slot holdings;
- **constraints**—airlines' ability to obtain and use slots.

We did not consider that it would be feasible to represent every airline currently operating at Heathrow, or with a potential interest in operating at Heathrow, in the experiment. We therefore categorised airlines into broad groups that cover the range of airlines. The categorisation was carried out on the basis of the airline features mentioned above, and is explained in more detail in sections 3.2 and 3.3.

Table 3.1 sets out seven high-level groups that collectively describe most of the airlines currently operating at Heathrow or that may do so in the future. While there will naturally be variation within each high-level group, the groupings allow for clear observations to be made about airline preferences and constraints.

Table 3.1 Airline groups

Group	Example airlines
Large Heathrow hub	British Airways
Mid-sized Heathrow hub	Virgin Atlantic
Short-haul low-cost	easyJet, Ryanair
Long-haul low-cost	Norwegian
North Atlantic	American Airlines, Delta Air Lines
European	Air France-KLM, Lufthansa
Middle Eastern and Asian	Cathay Pacific Airways, Emirates

Source: Oxera.

These groups are based largely on airlines' current business models. While we did take account of stakeholders' views about their future plans for use of any additional slots at Heathrow, we avoided speculation about very different future business models. We consider that this is reasonable given that the initial slots with the new runway are expected to be made available in the next few years.

3.2 Airline preferences

In order to provide participants with information regarding airlines' slot preferences, including timings and routes, we gathered information on a number of factors that determine these slot preferences. In particular, we analysed airlines' existing patterns of slot holdings at Heathrow, as well as at other London airports and other European hub airports, such as Schiphol Airport.[23] Our findings on airline preferences were also informed by stakeholder engagement.

Our research indicated that the following factors are likely to be key determinants of slot preferences across different types of airline, several of which are closely correlated:

- passenger demand for certain destinations will influence airlines' choice of route;
- the number of airlines flying a route will affect the profitability of that route, which will influence airlines' choice of routes;
- passenger demand for flights to a certain destination may determine optimal flight times (e.g. flights to a key business centre may indicate that optimal flight times are in the morning and at the end of the day);

[23] We looked at other airports because the capacity constraints at Heathrow may result in airlines not holding the slots that they would most prefer.

- airline business models, in terms of customer base, will determine passengers' willingness to pay for certain flight times. For example, business travellers may have higher willingness to pay for a specific morning departure time;
- competition for slots will affect the cost of acquiring the desired slots (on the secondary market or, in cases of allocation through auction).

There may also be non-financial factors (e.g. to meet the objectives of the government of the airline's home country). Below we consider the preferences for the different groups of airlines in more detail.

3.2.1 Large hub carrier

Heathrow serves as the hub airport in British Airways' (BA) hub-and-spoke network. BA is currently the only carrier that uses Heathrow as its main hub and operates at a large scale in terms of number of flights and destinations, with over 50% of all slots.[24] BA is part of the International Airlines Group (IAG), which also includes airlines such as Aer Lingus, Iberia and Vueling. We focus on BA rather than IAG in this section.

As of Summer 2017, BA flew to 139 airports in 61 countries. As shown in Figure 3.3, over 75% of its flights serve a European airport, with the remainder flying to a mixture of destinations in the Americas, Asia and, to a lesser extent, Africa.

Figure 3.3 BA flight destinations from Heathrow by region (Summer 2018)

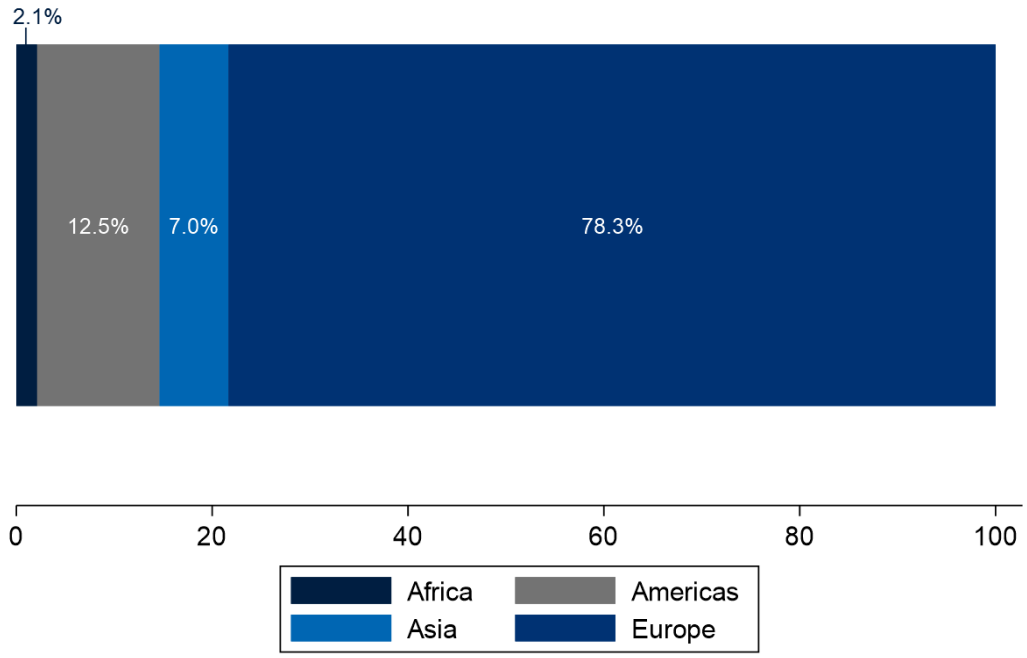

Source: Oxera analysis of OAG data.

In our engagement with BA, the airline indicated that [✂]% of its passengers connect through Heathrow. A large portion of these passengers are domestic/short-haul passengers that transfer onto a long-haul flight at Heathrow.

[24] The ACL Start of Season reports indicate that BA held 4,943 out of 9,664 slots in Winter 2018 and 4.714 out of 9,236 slots in Summer 2019.

In the hub-and-spoke model, airlines often operate in 'waves' (or 'banks') at their hub airport, such that groups of arrivals precede groups of departures in order to facilitate efficient transfers. One example of this is KLM, which operates a hub-and-spoke model at Schiphol. Figure 3.4 shows KLM's short-haul arrivals and long-haul departures, with five waves of each. The short-haul arrival peaks occur one to two hours before the long-haul departure peaks, facilitating the transfer from short- and medium-haul arrivals[25] onto long-haul flights to passengers' ultimate destinations.

Figure 3.4 KLM slots at Schiphol: short-haul arrivals and long-haul departures (Summer 2017)

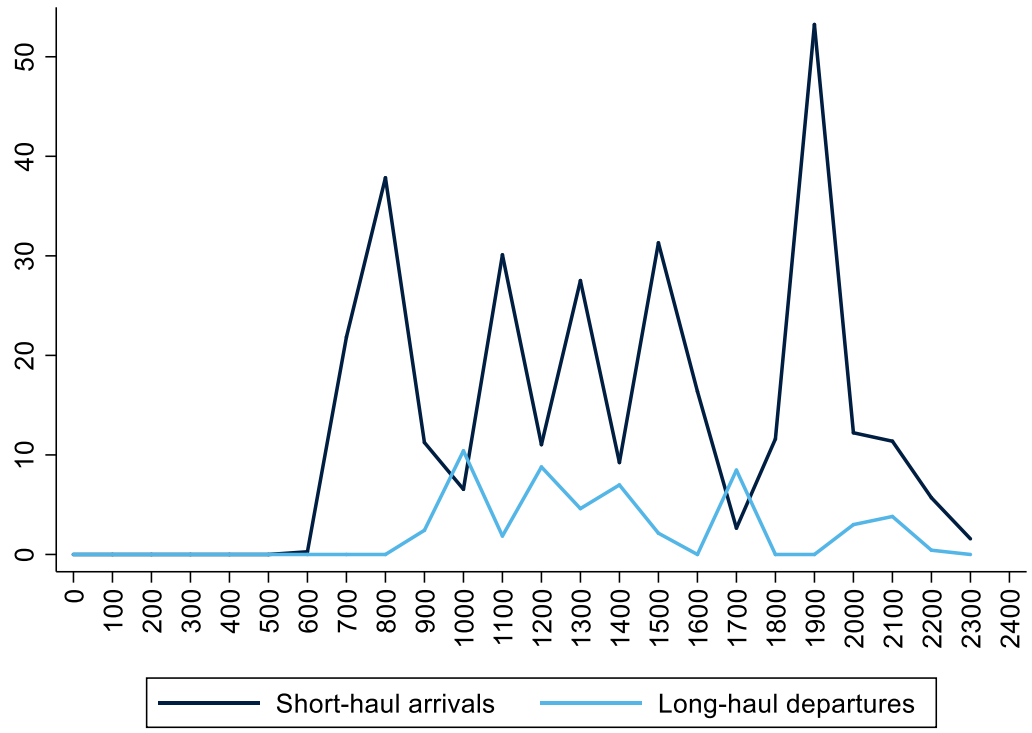

Note: Short-haul is defined as flights of up to 360 minutes in duration, with long-haul flights being over 360 minutes. Each point represents departures leaving in the hour following the time stated on the x axis (e.g. 600 refers to flights departing between 06.00 and 06.59).

Source: Oxera analysis of OAG data.

In contrast, BA does not appear to currently operate a large-scale wave pattern at Heathrow (although there may be smaller waves of flights that are not visible in the aggregate data). This is shown in Figure 3.5.

[25] Throughout the remainder of the report, when we refer to short-haul this includes both short- and medium-haul flights and is defined as flights of up to 360 minutes in duration.

Figure 3.5 BA scheduled flights from Heathrow (Summer 2018)

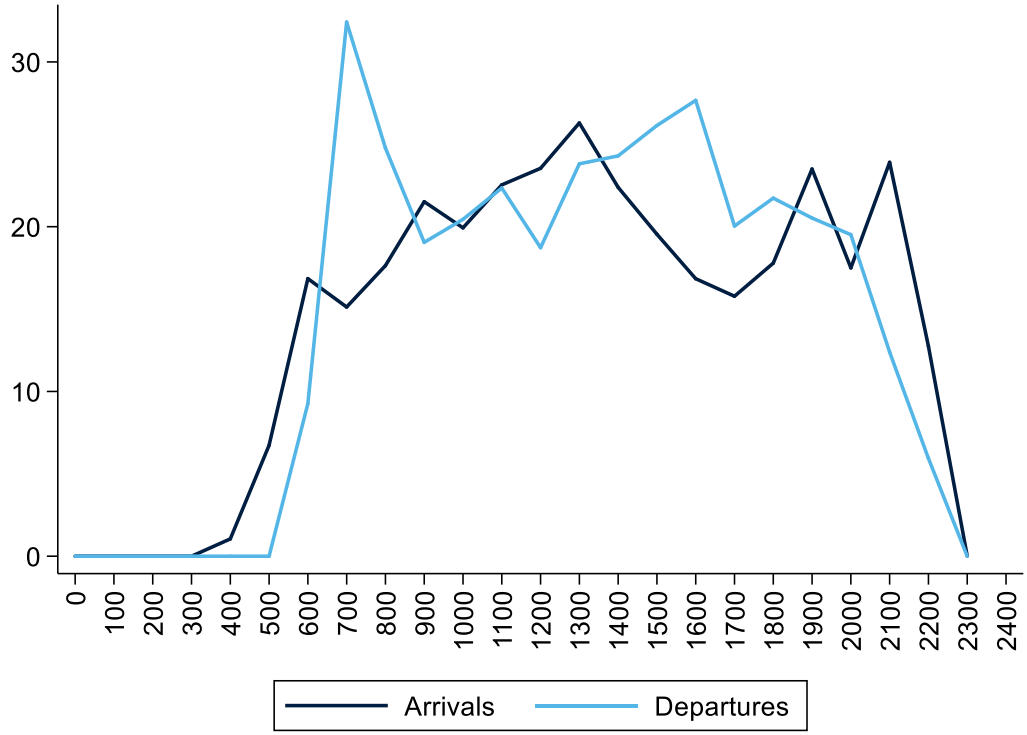

Note: Each bar represents departures (arrivals) leaving (landing) in the hour following the time stated on the x axis (e.g. 600 refers to flights departing between 06.00 and 06.59.).

Source: Oxera analysis of OAG data.

We consider that BA's flights at Heathrow are unlikely to be organised using a wave structure with the new runway, given its existing portfolio of slot holdings and current terminal constraints. [✂]

We also see that BA has a ratio of short-haul to long haul-flights of 2.5:1, indicating that short-haul flights 'feed' long-haul flights.[26] In our engagement with BA, it indicated that it would like to [✂].

3.2.2 Mid-sized Heathrow hub carrier

The second group contains only one airline: Virgin Atlantic. This is because Virgin is unlike most long-haul full service carriers, which generally operate flights exclusively to their base country from Heathrow. For example, Air Canada flies from Heathrow to Canada (only), and Delta Air Lines flies from Heathrow to the USA (only).

As shown in Figure 3.6, 79% of Virgin's flights from Heathrow are to the Americas, while 21% are to destinations outside the Americas (i.e. Africa and Asia). All of Virgin's current flights at Heathrow are long-haul—it does not fly to any European airports.

[26] Based on Summer 2018 OAG data on flights at Heathrow.

Figure 3.6 Virgin's destinations from Heathrow by region (Summer 2018)

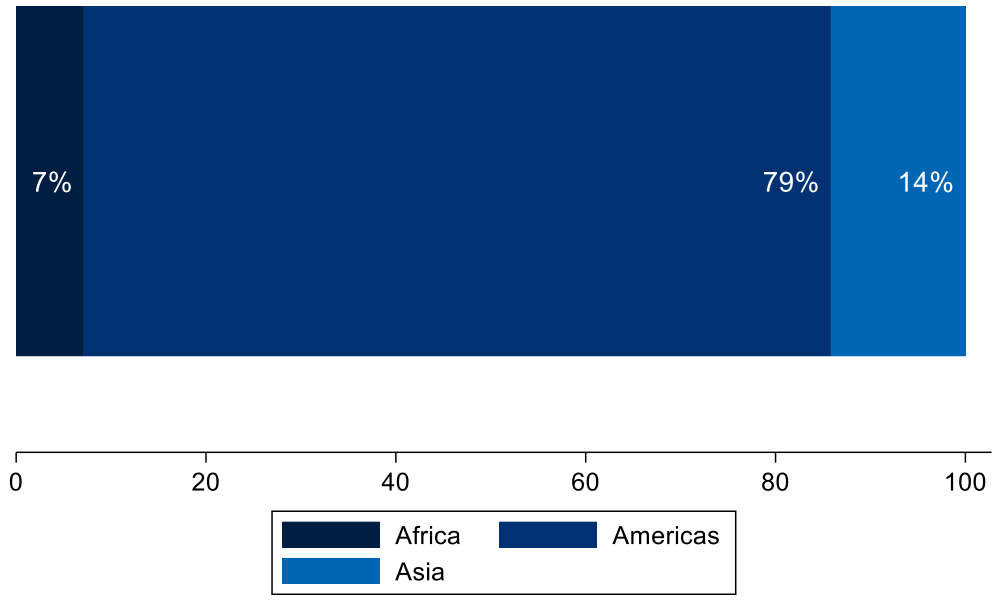

Source: Oxera analysis of OAG data.

We understand that Virgin intends to intend to operate a hub-and-spoke model from Heathrow.[27] Under this model, [✂]. Virgin is in a partnership with Delta, which also operates at Heathrow. Delta noted that [✂]. Due to this business strategy, Virgin has different slot preferences to the other airlines.

As such, for the purposes of the experiment, Virgin is currently the only airline that is classified as a mid-sized Heathrow hub carrier. Virgin holds around 3% of the available slots at Heathrow and uses Heathrow as its main hub.[28] Virgin has publicly indicated that it wishes to obtain 150 of the new Heathrow daily slots (equivalent to 109,200 ATMs, or 40% of all the new slots).[29]

As Virgin uses Heathrow as a hub for flights to several continents, it may also have different preferences for slots than, for example, North American airlines. The variety of destinations offered means that Virgin may need to hold slots at different times of the day. [✂].

3.2.3 Short-haul low-cost carriers

The third group of airlines are short-haul LCCs. Short-haul LCCs typically operate a business model based on providing low-cost air travel to European destinations. This business model relies on high aircraft utilisation (e.g. limited turnaround time), resulting in a preference for multiple slots (or rotations) each day. In order to meet the required number of rotations, an aircraft typically operates its first route early in the morning, and its last route in the evening.

Flybe, Eurowings (previously Germanwings) and Vueling are LCCs that currently operate from Heathrow. An examination of their slot holdings at Heathrow shows that departures and arrivals are clustered around 8am, 2pm and 8pm (see Figure 3.7).

[27] Virgin Atlantic (2019),' Virgin Atlantic plans to increase flights from London Heathrow'.
[28] ACL, Heathrow Airport (LHR) Summer 2019 Start of Season Report, 8 March 2019, p. 3.
[29] Collingridge, J. (2019), 'Sir Richard Branson's Heathrow plea', *The Sunday Times*, 23 June.

Figure 3.7 LCC flights from Heathrow in 2018

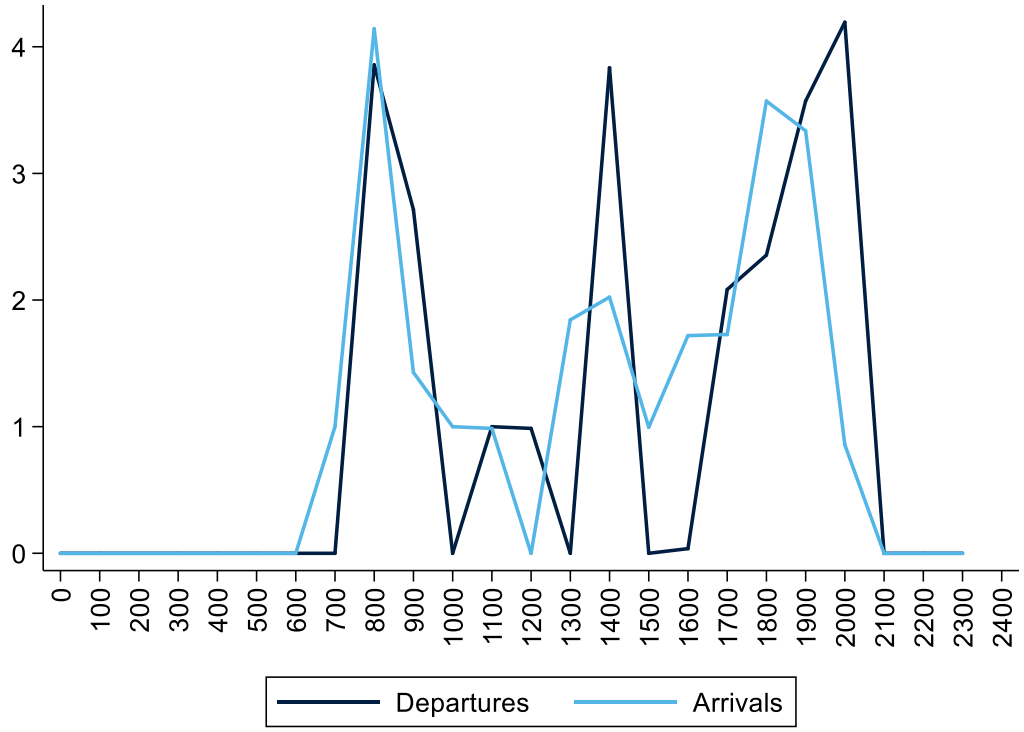

Source: Oxera analysis of OAG data.

The number of rotations required depends on the destination and therefore the flight time. For example, Flybe operates four daily flights from Heathrow to Newquay with around 3.5 hours in between each Heathrow departure. Eurowings operates three daily flights to several German destinations, with morning, mid-afternoon and evening departures.

Wizz Air confirmed to us that [✂].

Other short-haul LCCs that are not currently operating at Heathrow but that might have an interest in obtaining slots are easyJet and Ryanair. easyJet held 81,466 Summer slots at Gatwick in 2018, 43% of all slots available. It held 46% of Winter slots at Luton in 2018. Ryanair currently holds 77,990 slots at Stansted, 66% of the available slots.[30]

Figure 3.8 shows easyJet's flights from London airports. The data shows that easyJet generally operates at least three rotations, with most flights departing at 0600, and two more clusters of departures at midday and at 1800. In terms of arrivals, there are clusters around midday, 1800 and 2200. Given that easyJet's operations at other London airports are less constrained than those for LCCs operating at Heathrow, it is likely that these times are optimal flight times for its current route pattern.

[30] Calculations based on the ACL Summer 2019 Start of Season report for each airport.

Figure 3.8 easyJet flights from London airports (Summer 2018)

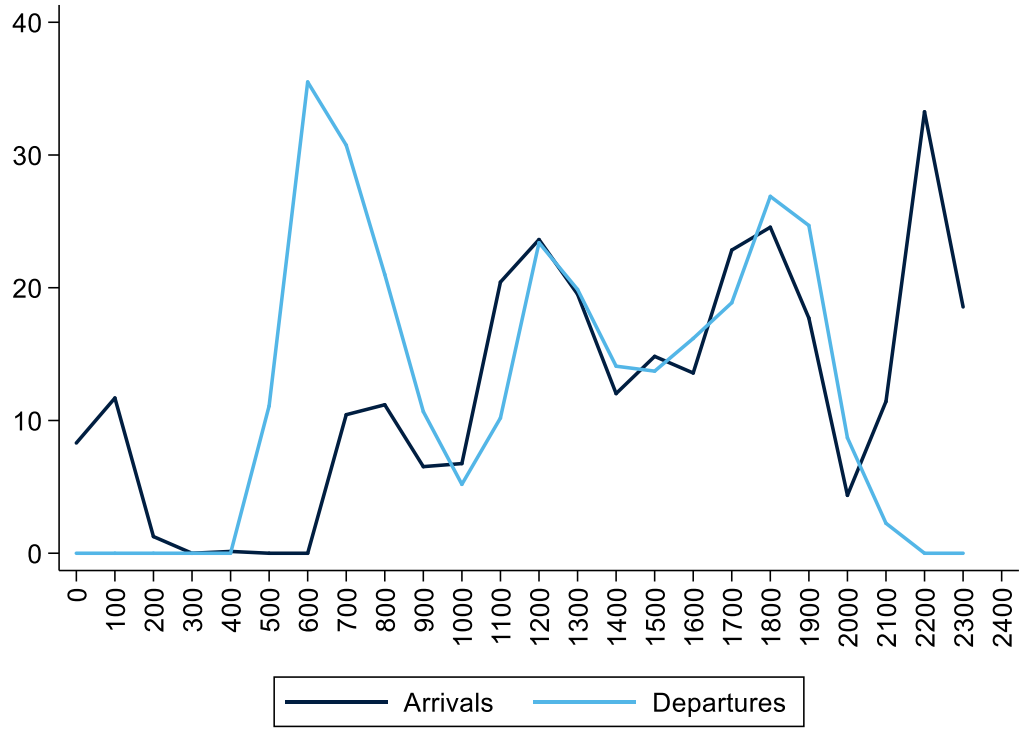

Source: Oxera analysis of OAG data.

easyJet has publicly stated its ambition to launch 19 new routes from Heathrow. It intends to use 30 aircraft to operate 30,000 to 55,000 routes per year.[31] In the response to the Airport Commission's consultation, easyJet confirmed that the airport infrastructure is suitable for it to viably operate from Heathrow, just as it does at Paris-Charles de Gaulle and Schiphol.[32]

3.2.4 Long-haul low-cost carriers

Long-haul LCCs such as Norwegian offer low-cost flights to long-haul destinations such as North America. Currently, Gatwick is the only UK airport from which Norwegian offers low-cost long-haul flights. [✂]

Norwegian's flights to New York from Gatwick operate mostly in the early morning, departing between 0600 and 1000. In addition, it offers a departure at the end of the day. [✂].

3.2.5 North Atlantic carriers

The fifth group of airlines are North Atlantic carriers. The two largest North Atlantic carriers that operate at Heathrow are United Airlines and American Airlines, which each hold around 2.5% of the slots.[33] North Atlantic carriers fly exclusively to US and Canadian airports, and each serves a number of routes.

These carriers have a high number of arrivals around 7am, followed by a large number of departures around 10am.

[31] Tobin, L. (2016), 'EasyJet aims to cash in on Heathrow's third runway', *Evening Standard,* 26 October, https://www.standard.co.uk/business/easyjet-aims-to-cash-in-on-heathrows-third-runway-a3379156.html.
[32] easyJet considers that there is sufficient stand capacity at Heathrow to support an easyJet base of 30 aircraft at Terminal 4.
[33] ACL, Heathrow Airport (LHR) Summer 2019 Start of Season Report, 8 March 2019, p. 3.

Figure 3.9 Departure times for North Atlantic carriers at Heathrow, other London airports and Amsterdam Schiphol (2017)

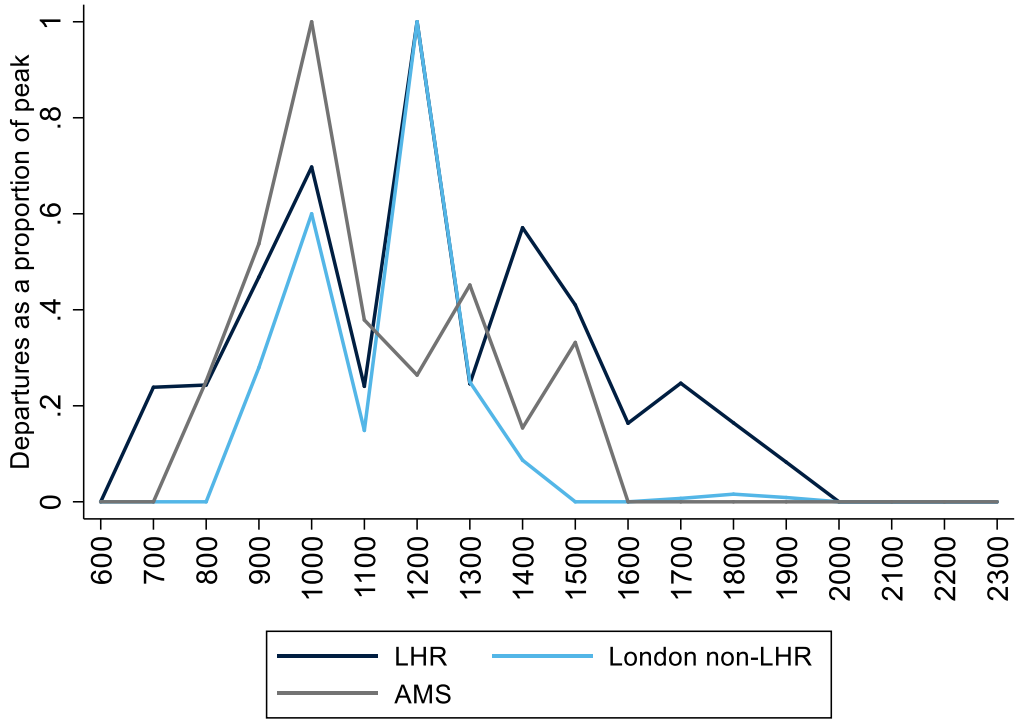

Note: LHR, London Heathrow. AMS: Amsterdam Schiphol. The North Atlantic airlines considered are Air Canada, Air Transat, American Airlines, Delta Air Lines, US Airways and United Airlines. For ease of understanding, only departures are plotted.

Source: Oxera analysis of OAG data.

In our engagement with United Airlines, it indicated that [✂]. Delta Airlines indicated that [✂]. In addition, North Atlantic carriers may wish to hold some extra departures around 6pm, accommodating business passengers that wish to fly back to North America at the end of the working day.

North Atlantic carriers currently fly only to destinations in North America. In our engagement with United Airlines, [✂]. American Airlines has indicated that, [✂].

In our engagement with Delta, [✂]. American Airlines indicated that [✂]. United Airlines also indicated that it [✂].

3.2.6 European carriers

The sixth group of airlines are European full-service carriers.[34] These are airlines based in Europe that use Heathrow to fly short-haul routes to a small number of European destinations (typically in their base country or countries).

The largest European carriers operating at Heathrow are Lufthansa, Aer Lingus and SAS with approximately 5%, 3% and 2% of slots, respectively.[35] Both Lufthansa and Aer Lingus serve only one destination country, while SAS serves a number of Scandinavian countries. However, the large majority of

[34] These European airlines are Aegean Airlines, Aer Lingus, Aeroflot Russian Airlines, Air Austral, Air Baltic, Air Europa, Air France, Alitalia, Austrian Airlines, Belavia, Brussels Airlines, Bulgaria Air, Croatia Airlines, Finnair, Iberia, Icelandair, KLM, LOT, Lufthansa, SAS Scandinavian, SATA, Swiss/Crossair, TAP Air Portugal, TAP Portugal, TAROM and Thomas Cook Airlines.
[35] Percentages calculated based on the information provided in the ACL Summer 2019 Start of Season report for Heathrow.

European carriers currently operate routes to only one destination from Heathrow.

The short-haul nature of these flights means that they typically use slots throughout the day, as illustrated in Figure 3.10.

Figure 3.10 European carrier flights from Heathrow (Summer 2018)

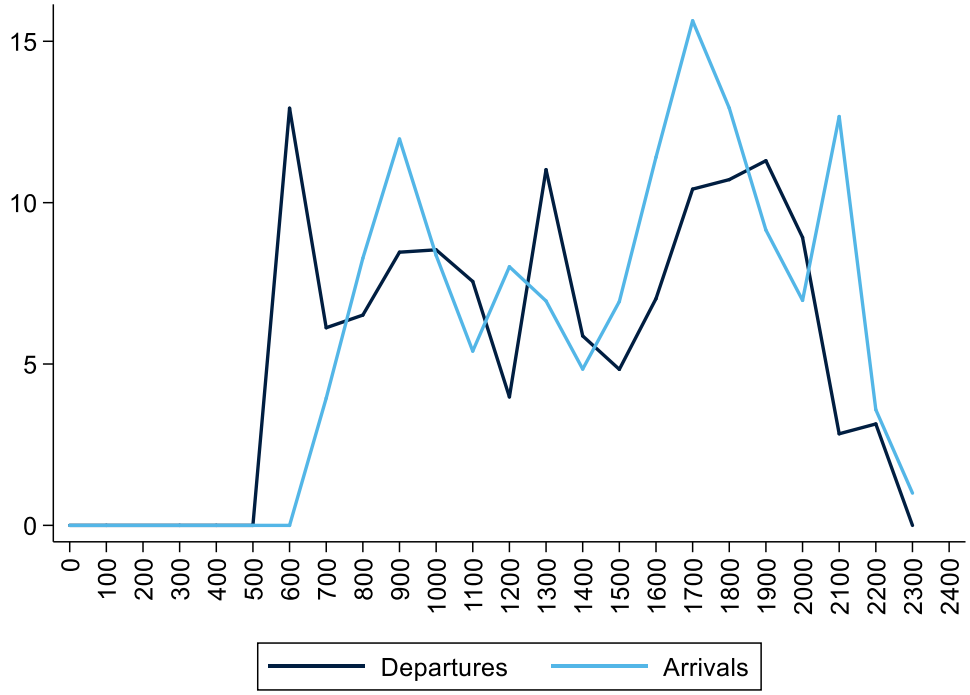

Source: Oxera analysis of OAG data.

The exact number of flights offered throughout the day by an individual airline is highly dependent on the destination. Flights to nearby large hubs such as Amsterdam, Dublin and Frankfurt are likely to be offered very frequently. For example, Lufthansa operates flights to Frankfurt from Heathrow every hour from 6:30am until 7:30pm, except for midday.

Destinations that are located further away tend to be operated fewer times a day. For example, SAS operates five to six flights a day to several Scandinavian airports from early morning until the evening. Furthermore, destinations that are located further away and which tend to have a higher proportion of leisure, rather than business, passengers tend to have fewer daily flights.

We spoke to Lufthansa, which indicated that [✂]. KLM noted that [✂].

3.2.7 Middle Eastern and Asian carriers

The seventh and final group of airlines are Middle Eastern and Asian carriers that fly long-haul routes to the Middle East and Asia from Heathrow. Given that airline preferences for slot times depend in part on the local arrival time at the destination, Middle East and Asian airlines prefer different slot times to, for example, North Atlantic carriers.

We recognise that there is variation within this group (given the size of the Middle East and Asia). However, Figure 3.11 shows that both Asian and Middle Eastern airlines hold slots throughout the day with large peaks of

departures concentrated towards the late evening, and smaller peaks of arrivals in the morning and in the afternoon.

Figure 3.11 Middle Eastern and Asian carrier flights at Heathrow (Summer 2017)

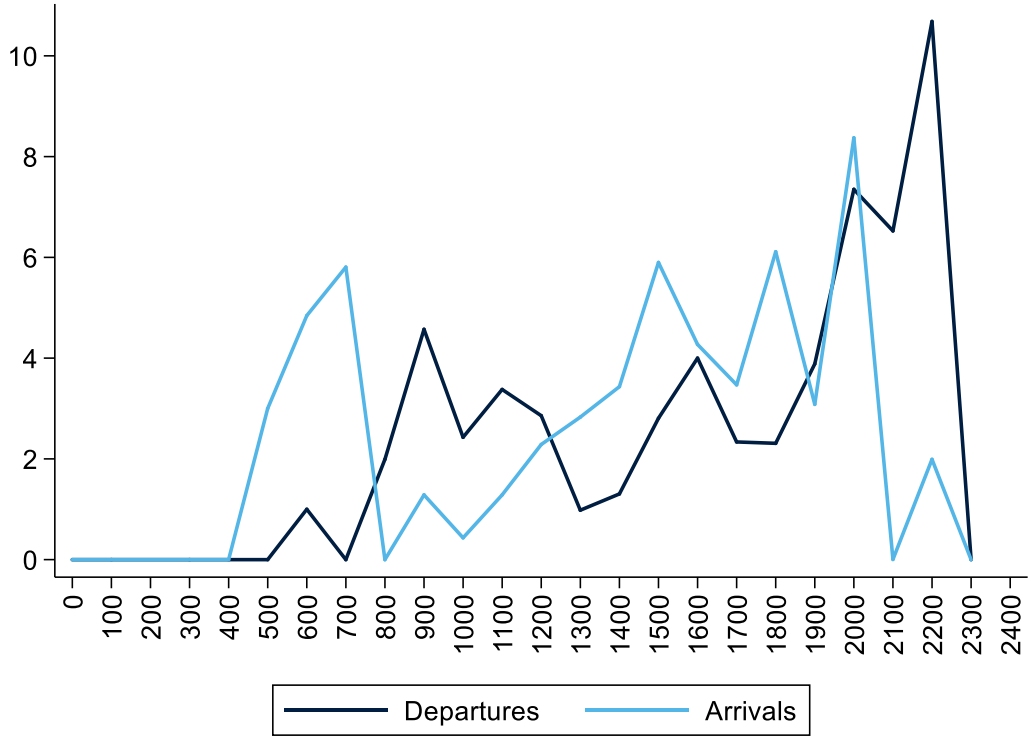

Note: The Asian airlines are Cathay Pacific, China Airlines, China Eastern, China Southern, EVA AIR, EVA Airways, Garuda Indonesia, Korean Air, Malaysia Airlines, Singapore Airlines and Vietnam Airlines. The Middle Eastern airlines are Arkia – Israeli airlines, El Al Israel Airlines, Emirates, Etihad, Qatar, Oman Air, Royal Jordanian and Syrian Arab Airlines.

Source: Oxera analysis of OAG data.

Industry experts informed us that, despite the variation in destinations offered, Middle Eastern and Asian airlines tend to have similar preferences for slot times. This is partly because passengers often transfer at Middle Eastern airports on their way to Asian destinations, meaning that the time zone of the end-destination for passengers is similar. For example, a flight departing from Heathrow at 10pm will arrive at an airport in the Middle East in the morning and in South East Asia towards the end of the day.

Furthermore, industry experts informed us that the wave structure of Middle Eastern airports is often designed around scarce European slots, meaning that they are optimised for the evening slots that the carriers already hold. Nevertheless, Middle Eastern and Asian carriers may also wish to hold early morning slots for arrivals at Heathrow.

The only carrier in this group that we spoke to is ANA. ANA told us that [✂].

3.2.8 Summary of preferences

In summary, all groups of airlines value the first slots of the day at Heathrow highly. However, there are airline group-specific preferences, such as the demand by short-haul LCCs for multiple rotations and the demand by Middle Eastern and Asian airlines for evening slots. The attributes of the airlines in

each of the groups and their preferences for slots are set out in Table 3.2 below.

Table 3.2 Airline groups

Group	Business model	Typical flight duration	Typical flight destination	Preferred times	Example airlines
Large Heathrow hub	Full-service	Short-haul and long-haul	Worldwide	Throughout the day in waves	British Airways
Mid-sized Heathrow hub	Full-service	Long-haul	North America, Asia, Africa	Peaks in morning	Virgin Atlantic
Short-haul low-cost	Low-cost	Short-haul	Europe	Multiple rotations	easyJet, Ryanair
Long-haul low-cost	Low-cost	Long-haul	North America	Peaks in morning	Norwegian
North Atlantic	Full-service	Long-haul	North America	Peaks in morning	American Airlines, Delta Air Lines
European	Full-service	Short-haul	Europe	Throughout the day	Air France-KLM, Lufthansa
Middle Eastern and Asian	Full-service	Long-haul	Middle East and Asia	Throughout the day with peaks in evening	Cathay Pacific Airways, Emirates

Note: the preferred times are representative of the airline groups and do not necessarily represent the preferences of example airlines or airlines that we spoke to as part of the stakeholder engagement.

Source: Oxera.

3.3 Airline constraints

Airlines and other stakeholders have indicated that there are a number of constraints that they need to take into account when obtaining slots at Heathrow, as follows.

- Air Service Agreements (ASAs) between countries, which are a prerequisite for airlines to fly a route. In the experiment, we did not allow airlines to operate routes that would not be allowed under current ASAs.
- Airport capacity constraints (e.g. terminal capacity). Several stakeholders indicated that (after runway constraints) one of the primary constraints is terminal capacity. Some stakeholders considered that there is not currently sufficient capacity to be able to deal with the increased passenger flow resulting from the increased number of flights at Heathrow with the new runway. However, in the experiment we assumed that the slot allocation mechanism would ensure that the first phases of slots will be awarded in such a way that the combined aircraft capacity does not exceed the terminal capacity. In other words, we assumed that terminal capacity would not constrain demand in the experiment.
- Available aircraft given airline fleet, which affects the feasibility of slot holdings. A number of airlines expressed the potential constraints that they might face in ensuring that they acquire the appropriate fleet. This forms a constraint mainly in terms of timing, because there tends to be a lead time of around three to four years for aircraft purchasing. Aircraft purchase also represents one of the biggest costs for an airline, which means that there

may be financial constraints. This constraint may be somewhat alleviated by the potential for airlines to lease aircraft, or if capacity is allocated well in advance of the runway becoming operational. In the experiment we incorporated airline-specific constraints around the ability of an airline to operate a large number of new slots from Heathrow.

- Available crew given the number of staff employed. Airlines indicated that ensuring that they have sufficient crew members, particularly pilots, could also be a constraint in the short run. However, unlike the current slot allocation process, the newly available slots may be allocated a number of years in advance. This means that obtaining sufficient crew is likely to form a key constraint only when an airline gains a very large number of slots relative to its current operations. Therefore, in the experiment we incorporated airline-specific constraints around the ability of an airline to operate a large number of new slots from Heathrow.
- Choice of route. This may determine feasible flight times, given slot availability and curfews at the destination airport.

In addition, airlines may face financial constraints in obtaining slots in the secondary trading market or if slots are allocated through an auction mechanism. A number of factors play a role in determining an airline's financial constraints and its ability to raise capital:

- total revenue;
- net income;
- the airline's weighted average cost of capital;
- whether the airline has access to government funding.

Table 3.3 provides an overview of selected financial metrics for the seven airline categories.

Table 3.3 Overview of financial constraints of different airline groups

Group	Total revenue (billions)[1]	Net income (billions)[2]	Post-tax WACC[3]	Majority government share[4]	Reference airline(s)
Large Heathrow hub	£12.2	£1.4	6.1	No	British Airways
Mid-sized Heathrow hub	£2.8	£0	2.1	No	Virgin Atlantic
Short-haul LCC	£5.9–£6.3	£0.5–£1.3	6.2–10.3	No	easyJet, Ryanair
Long-haul LCC	£3.6	£-0.1	5.3	No	Norwegian
North Atlantic	£32.2–£34.6	£1.1–£3.1	5.1–6.9	No	American Airlines, Delta Air Lines
European	£6.3–£31.5	£1.8–£1.9	2.8–4.2	No	Air France-KLM, Lufthansa
Middle Eastern and Asian	£10.1–£19.3	£0.3–£0.6	4.9–8.4	Yes	Cathay Pacific Airways, Emirates

Source: [1] Airlines' annual reports. [2] Airlines' annual reports. [3] Bloomberg data, individual airlines' post-tax WACC. [4] ICAO's (International Civil Aviation Authority) list of government-owned and privatised airlines. Company currency converted to GBP using 12-month average exchange rates reported on 12 August—see https://www.ofx.com/en-au/forex-news/historical-exchange-rates/yearly-average-rates/.

We considered that these measures should be interpreted with caution for the purposes of the experiment, given that the experiment covered long-term slot allocation taking place several years in the future. We therefore also incorporated the views of industry experts in calibrating the budget constraints in the experiment.

Furthermore, we took account of airline statements about their budget constraints. For example, Norwegian stated that it has a very limited budget and would not be able to participate in an auction.

3.4 Aircraft usage

Aircraft size is another dimension affecting the use of slots at Heathrow—for example, larger aircraft require more terminal capacity.

There is a strong relationship between flight duration and the number of seats on the aircraft. The DfT has forecast the use of aircraft in terms of the number of seats by destination over the next 20 years, as shown in Figure 3.12.

Figure 3.12 Forecast aircraft size by destination

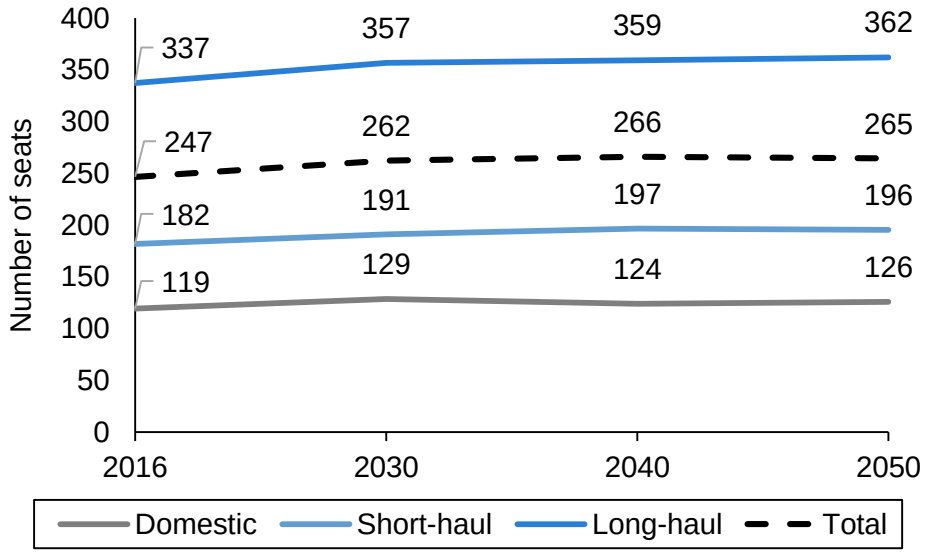

Note: Short-haul is defined by the DfT as flights within Europe (including non-EU countries, such as Eastern Russia).

Source: Department for Transport (2017), 'UK Aviation Forecasts', October.

The overall trend shows that, on average, aircraft used on flights to short-haul destinations have just over half the seat capacity of aircraft used on flights to long-haul destinations. This trend is expected to be stable over time, even though both short-haul and long-haul aircraft will continue to increase in size.

Given the correlation between aircraft size and destination, it is expected that the fleet of an airline will be influenced by its routes. Where slots are allocated in advance, this will allow for long-term planning and for the fleet to be adapted to accommodate the airline's strategy. We therefore asked participants for their choice of route in the experiment, and used this to estimate the number of seats on the aircraft.

3.5 Conclusion

The above research and analysis has informed our understanding of how airlines may behave when requesting slots for Heathrow's third runway. The

number of airline groups, their preferences and their constraints were inputs to the experiment design, which is described in sections 4–6.

4 Design of the laboratory environment

In this section we describe the experimental environment, including how the laboratory experiment was selected and set up, the practical considerations for the experiment, and the choice of participant pool.

4.1 The laboratory environment

Unlike surveys that elicit 'stated preferences' based on how people or organisations say they will behave, behavioural experiments elicit 'revealed preferences' based on observed behaviour. There are three main ways in which behavioural experiments can be undertaken—in a laboratory, online, and in the field. In laboratory and online experiments participants are asked to behave 'as if' the decisions were taken 'for real', whereas field experiments involve observing real behaviour (without participants necessarily being aware of the experimentation).

The choice between these options depends on the nature of the research question, with different types of experiment suited to different scenarios. We decided to conduct this experiment in a laboratory environment for a number of reasons.

First, and most importantly, there is greater control over factors that may influence decision-making in laboratory experiments compared with online experiments. In a laboratory experiment the experimenter can ensure that there is no cheating or communication between participants, and that the participants are paying full attention. The experimenter can also ensure that the conditions in which the experiment takes place are identical for each participant (e.g. they use the same type of computer) and for each session (e.g. they use the same computers). Furthermore, the experimenter can check the identity of participants to ensure that those who take part are those who have been invited.

Second, in a laboratory experiment:

- the experimenter can answer questions from participants if they arise;
- the experimenter can hand out hard copies of the instructions.

However, lab experiments are typically more time-consuming than online experiments (e.g. because it takes time to welcome participants into the lab and pay them at the end). We took this into account when designing the experiment and considering its duration.

We also note that laboratory experiments can be less 'realistic' than field experiments, as the participants know that they are being observed and are making 'as if' decisions (rather than decisions for real). However, field trials are not always possible, and in any case it is often desirable to conduct testing in a controlled environment before proceeding to a field trial. For example, the results of a laboratory experiment can identify risks associated with different policy options, and help refine the policies to be tested in a field trial.

4.2 Practical design considerations

4.2.1 Location and software

The experiment was conducted at the University of Nottingham in the Centre for Decision Research and Experimental Economics (CeDEx). CeDEx operates an experimental laboratory with capacity for 32 participants, which

means that up to three groups of ten participants were able to use the lab simultaneously. The lab is designed to ensure that participants cannot communicate with each other during the experiment, as each participant has their own computer terminal with screens around the desk. Each session of the experiment was overseen by two experimenters who were able to answer questions from participants.

LIONESS software was used for the experiment programming, allowing for best-practice data collection at both participant and group level.[36]

4.2.2 Ethical approval

The experiment was conducted according to the highest ethical standards, in line with the University of Nottingham's Code of Research Conduct and Research Ethics. The experiment obtained ethical approval from the Nottingham School of Economics Research Ethics Committee.

4.2.3 Pilot

Before running the experiment, we ran a 'pilot' of each treatment.[37] The pilots enabled us to check:

- that the experiment functioned;
- the time taken to run the experiment;
- that participants understood the instructions;
- the payoffs earned.

We made a number of minor changes after the pilots before running the experiments. Data from the pilots is therefore not used in the data analysis or in forming conclusions.

The pilots took place in June and July 2019, with the experiment taking place in July 2019.

4.3 Duration

The experiment was designed to take under two hours. This provided enough time (and rounds) to ensure that participants could learn and that enough data would be generated, but not so long that participants would lose focus.

Each screen in the experiment had a countdown timer to ensure that the experiment ran to time. The timers were set using the timing data from the pilots to ensure that participants had an appropriate amount of time to understand the information and to make their decisions.[38]

4.4 Participants

4.4.1 Using students as participants

CeDEx has a volunteer participant database of undergraduate and postgraduate students, with over 3,000 registered participants and new

[36] https://lioness-lab.org/
[37] A pilot is a pre-test of the experiment where certain features of the experiment are calibrated.
[38] We did not place a timer on most screens in the pilots, so that we could observe the time taken by participants (without constraints). The timers for the experiment were in general set to be greater than the average time taken (to allow for most participants to navigate their way through without being restricted), but less than the maximum time taken (to ensure that the experiment ran on time). More restrictive timers were placed on screens where the information had been shortened/simplified following feedback from the pilot.

students registering every semester. Participants were recruited and randomly assigned to treatments via the ORSEE (Online Recruitment System for Economic Experiments) programme as is best practice in experimental economics.[39] The descriptive statistics of the sample are shown in 7.1.

Each participant could take part in the experiment only once (e.g. they could not come back and take part in a different treatment) and participants in the pilot were not used again in the experiment.[40] This ensured that all participants had the same opportunity to learn, and that the design of one treatment (or outcome of one session) did not impact the behaviour of a participant in another treatment (or session).[41]

The use of students as participants is common in behavioural experiments, including experiments in behavioural industrial organisation.[42] One alternative to using students would have been to use industry professionals or experts. Apart from the practical issues that could have arisen in achieving a large enough sample size to obtain statistically significant results, the use of industry professionals/experts would have resulted in methodological concerns for the following reasons:

- such participants may have had prior beliefs about an airline's preferences for slots, meaning that their behaviour may not have been in line with the instructions. These beliefs are likely to have differed between participants in the different sessions, which could have undermined the results of the experiment;
- such participants may have seen the experiment as part of a wider 'game', in which they could influence government policy through the results of the experiment. In other words, they may have brought their own views about slot reform to bear, undermining the results of the experiment.

With regard to professionals who work for airlines, these methodological issues would have arisen regardless of whether the participants played the role of the airline they work for, or a different airline.

An alternative approach to a behavioural experiment would have been to use agent-based modelling (ABM), whereby a computer simulation of the slot allocation mechanism is used. While there have been ABM studies on slot allocation (see Appendix A3), there are a number of advantages of behavioural experiments, including:

- ABM would require a more complete specification of airline preferences and strategies, which would limit the richness of the results. For example, in the experiment we allow airlines to mislead the slot coordinator (in terms of not flying the requested route), but we did not specify whether participants should or should not mislead the slot coordinator;

[39] http://www.orsee.org/web/
[40] This ensured that participants were all given the same opportunity to learn, and avoided the sample being 'contaminated' by participants who were aware of the precise mechanism being tested by the experimenter.
[41] While there are experiments which test different treatments with the same participants, the aim of these experiments is to assess the participants' reaction to the *change of treatment*, rather than the difference between the treatments per se. In this case, we are more interested in the difference between treatments, rather than the speed of participants' change of strategy. Further, the use of repeated identical rounds in the experiment facilitated participants' learning for any given treatment.
[42] For a review, see Brandts, J. and Potters, J. (2018), 'Experimental industrial organisation', in L.C. Corchón and M.A. Marini (eds), *Handbook of Game Theory and Industrial Organisation, Volume II*, Edward Elgar Publishing.

- behavioural experiments allow for non-standard preferences to impact the results. These non-standard preferences, such as loss aversion, may be reflected in reality by firms.

4.4.2 Anonymity

In accordance with best practice, participant decisions in the experiment were anonymous (i.e. names of participants were not recorded in the data). In each session the participants knew the airline that they were representing but did not know which other airlines were being played by which participants. Furthermore, participants' payoffs were anonymous, which meant that they were able to make their decisions without regard to their future interactions with other participants.

4.4.3 Incentivisation

In accordance with best practice, we incentivised the participants with a cash payment at the end of the experiment.[43] This ensured that the participants were motivated to try their best to act in accordance with the guidance and instructions provided to them, and to maximise their airline's payoffs.

The payoff consisted of three elements (which was explained during the experiment):

- a participation fee (£3);
- a payoff from the slot allocation task (participants were told that this would be up to £14);
- a payoff from the risk aversion question at the end of the experiment (which varied according to participant decisions).

The payoffs received are outlined in section 7.

Following academic best practice, the payoffs to the participants were described in terms of Experimental Currency Units (ECUs). The use of ECUs also allowed for different conversion rates to be applied for different participants so that every participant could expect to receive a reasonable return for their effort, despite the fact that some airlines could expect to generate substantially higher payoffs.[44] Each participant was shown their ECU to £ conversation rate at the start of the experiment. The conversion rates are shown in section 5.5.1.

The payoff from the slot allocation task was determined based on one of the three identical rounds that participants completed. The round was randomly selected, and this was explained to participants at the beginning of the experiment.[45]

[43] For example, see Gneezy, U. and Rustichini, A. (2000), 'Pay enough or don't pay at all', *The Quarterly Journal of Economics*, **115**:3, pp. 791–810.
[44] The use of different conversion rates between participants has been used in experiments where there are asymmetric contests. For example, see Chowdhury, S.M., Kovenock, D. and Sheremeta, R.M. (2011), 'An experimental investigation of Colonel Blotto games', *Economic Theory*, April, **52**:3, pp. 833–861.
[45] For example, see Baltussen, G., Post, G.T., Van Den Assem, M.J. and Wakker, P.P. (2012), 'Random incentive systems in a dynamic choice experiment', *Experimental Economics*, **15**:3, pp. 418–443.

4.5 Experiment structure

The structure of the experiment followed academic best practice, with clear instructions followed by a practice round and then three repeated rounds. The experiment structure is shown in Figure 4.1.

Figure 4.1 Experiment structure

Source: Oxera.

When participants entered the lab, instructions were presented on their screens. Participants were also provided with a hard copy key information sheet, which contained all the information that they would need when making their decisions during the experiment. The screens informed participants where the information could also be found on the information sheets.[46] The information sheet included their role, payoffs, constraints, and a description of how the slots would be awarded (i.e. how the slot coordinator made decisions, or how the auction worked).

In designing the on-screen and paper instructions (see appendix A2), we made the instructions very clear. We did this in a number of ways, including the following.

- Using simple language. For example, we avoided lengthy or unusual words.
- Avoiding industry jargon where possible. For example, instead of referring to 'the ACL algorithm' or 'the slot coordinator', we referred to 'the computer'.
- Presenting only necessary information. Only information that was required for participants to make decisions was included. For example, wider context on the nature of slot reform was excluded.
- Using images where possible. For example, screenshots of important screens were presented in the instructions, so that participants understood how to make decisions before they reached the screens.
- Limiting the opportunity for participants to make clear 'mistakes' over the choice of route. As discussed in section 3, airlines are restricted in their choice of route by internal strategy and external constraints (e.g. air service agreements). Therefore we restricted participants from selecting routes where their airline would or could not fly.

The instruction screens were as follows.

- **Welcome screen.** Information about the experimental environment, in terms of consent, privacy, and the absence of deception.
- **Thanks for your participation.** Information about the overall context (use of the new runway at Heathrow), and information about payoffs.
- **Slots.** Information about runway slots and how many slots at different time periods would be allocated. Information was also given about the ten airlines in the experiment and which airline the participant was representing.

[46] Participants were allowed to write on their key information sheets (as each session had new copies).

- **Rounds.** Information about the structure of each round, and the nature of the practice round.
- **Screenshots.** Screenshots of the bidding, trading and route selection stages, with explanations of how participants would make decisions.
- **Your airline.** Information about the participant's role, including the airline's strategy, available routes, budget constraint, and new entrant status.
- **Your airline's payoffs.** Information about the airline's payoffs, including both the base element and the special element.
- **The base element of your payoffs.** A chart showing the base element of the airline's payoffs.
- **The computer.** Information about how the computer allocates slots, and the nature of new entrant slots. Information was also provided about each airline's pre-existing slot holdings.
- **Test questions.** Three test questions checking participant understanding. Participants could not proceed without correctly answering all three questions.

Immediately following the test questions, participants took part in a 'practice round' to familiarise themselves with the experiment. This practice round did not affect participant payoffs, and the participants were therefore free to try different strategies without the pressure of needing to optimise their decisions.[47]

Following the practice round, there were three repeated rounds. We included three rounds so that participants could learn from the impact of their decisions in each round and potentially adopt different strategies as a result. The three rounds were identical and independent (the decisions taken in one round and the outcomes of that round did not affect the other rounds), and each participant remained as the same airline throughout the rounds. This ensured that participants had the opportunity to learn how to optimise decisions as a given airline, without having to learn different roles.

In addition to the key information sheets, each decision-making screen (i.e. where participants made their decisions in the rounds) had a payoff simulator. This automatic calculator allowed participants to enter any combination of time and route to find the payoff that it would generate.[48] Furthermore, each decision-making screen had the following:

- clear instructions on how to make decisions, and what different pieces of information meant;
- a sentence at the bottom of the screen stating what screen the participant would see next;
- a reminder of the participant's airline role;
- a reminder of the round number.

Towards the end of the experiment participants were presented with a series of questions about their characteristics and preferences. This information is used

[47] Data from the practice round is not used in the data analysis.
[48] The payoff simulator ran on an ex ante slot-by-slot basis, meaning that it did not take into account the 'special elements' of the payoff (see section 6.2). This was clearly explained to participants.

in the data analysis to test whether these factors drive the results (see section 7). Participants answered questions about the following topics:

- **Preferences regarding risk aversion and loss aversion.** We asked risk aversion[49] and loss aversion[50] questions as these attributes could be drivers of behaviour in the context of bidding and trading.
- **Level of focus.** Following academic best practice, participants were asked a question to test whether they were paying close attention at this stage of the experiment.[51]
- **Personal characteristics.** Participants were are asked to state their gender, age, the course that they were studying, and whether they were at postgraduate or undergraduate level.
- **Prior knowledge of slot allocation and the airline industry.** Participants were asked if they had ever worked in the aviation industry and if they had any previous familiarity with slot allocation. Participants were also asked how many flights they took each year.

[49] Following academic best practice, we incentivised participants to answer the Holt and Laury (2002) question eliciting the level of risk aversion. Holt, C.A. and Laury, S.K. (2002), 'Risk Aversion and Incentive Effects', *The American Economic Review*, **92**:5, December, pp. 1644–1655.
[50] 'Imagine there was a gamble with a 50% chance of losing £100 and a 50% chance of winning a positive reward. How big would the positive reward need to be for you to accept the gamble?'
[51] 'You buy a bat and ball for £1.10. The bat costs £1 more than the ball. How much does the ball cost?' See Frederick, S. (2005), 'Cognitive Reflection and Decision Making', *Journal of Economic Perspectives*, **19**:4, pp. 25–44.

5 Design of the slot allocation environment

Based on the research and analysis described in sections 1–3, in this section we explain the design of the slot allocation environment in the experiment. In particular, we set out the number of airlines, the number of slots and time periods, and the number of representative routes that were used in the experiment. We also explain the airline payoffs and constraints that we established for the ten airlines, and the design of the slot coordinator mechanism.

5.1 Number of airlines

To ensure that the range of airlines that are likely to be interested in acquiring slots at Heathrow was captured in our experiment, participants were asked to represent a specific airline type and were required to bid for slots according to that airline's objectives and constraints.

We included ten participants in the experiment, with one or two participants for each of the airline groups set out in section 3.[52] The ten airlines are listed in Table 5.1.[53] There was not a one-to-one relationship between the representative airlines in the experiment and real airlines, as we aggregated information from multiple airlines when designing the representative airlines.

This means that it would not have been possible to assign each of the airlines in the experiment the name of a real airline. We therefore gave each of the ten airlines a fictional name. The use of fictional airline names also removed the risk that participants had preconceptions or biases about airlines, and so behaved differently as a result. However, we did use the name of Heathrow Airport and real destinations (e.g. Frankfurt), to help ensure that the context felt realistic to the participants. The airlines' names were selected such that participants could identify their role in the experiment (e.g. East Asia is an airline flying to China).

Table 5.1 Representative airlines

Airline group	Number of airlines	Airline names
Large Heathrow hub	1	National Airlines
Mid-sized Heathrow hub	1	Jet Flights
Short-haul low-cost	1	Speedy Flights
Long-haul low-cost	1	Quick Wings
North Atlantic	2	Atlantic Airlines, North American Airways
European	2	Euro Airlines, Royal Airways
Middle Eastern and Asian	2	West Asia Airways, East Asia Airways

Source: Oxera.

The new entrant rule, as described in section 2, means that airlines that do not hold many slots (or indeed any slots) are collectively reserved access to 50% of new slots. Each new entrant airline may gain up to two daily slot pairs under

[52] We did not run the experiment unless we had exactly ten participants (one to play the role of each airline). More than one participant per airline would have added unnecessary complexity to the experiment, given that we were not examining the decision-making process within airlines. More than one participant per airline would also have required more time in the experiment for the participants to communicate with each other (which would have meant having less time available for other, more important aspects).
[53] While the airlines within each group were similar in terms of strategy, preferences and constraints, there were some differences between them (such that no two airlines were identical in every respect).

the new entrant rule (after which it no longer qualifies as a new entrant). This dynamic was included in the experiment.

Three of the ten representative airlines qualified for new entrant status—Speedy Flights, Quick Wings and East Asia Airways. Each of these three airlines held no pre-existing Heathrow slots, and therefore could obtain up to two slot pairs of new entrant slots.

5.2 Slots and time periods

Given the constraints of running a behavioural experiment, we simplified some of the aspects of slot allocation while retaining the most important elements.

The experiment explored the allocation of slots for a 'representative day'. The results can then be extrapolated to a weekly basis, although we recognise that in reality not all slots are held by the same airline every day of the week.

Slots were also grouped together in time periods. Within each time period, the slots had the same value (i.e. they were homogeneous). This simplification was necessary in order to make the experiment easily understandable for participants. In any case, we are not interested in the precise timing of slots at the five-minute level (which can often be adjusted through annual re-times).

Following the stakeholder engagement, we decided to have five time periods: 0600–0800; 0800–1000; 1000–1400; 1400–1800; and 1800–2200. The first two time periods were two hours in duration, while the last three time periods were four hours in duration. The difference in duration was to reflect the fact that stakeholders told us that the very first slots in the morning would be the most valuable to airlines, and therefore that the first time slots should be separated out from the rest.

Departures and arrivals were allocated in pairs within time periods in the experiment. This is because, in reality, airlines request slot pairs (and slots are useful only when paired). While we recognise that turnaround times differ between airlines, many pairs will be used within a time period. Alternatively, the experiment could have allocated departures and arrivals separately, but this would have introduced another element of complexity and would not have added greatly to our understanding of the impact of different slot reforms. Another option would have been to allocate only one of departures or arrivals, but this would have been less realistic.

We also recognise that, for airlines that base aircraft at Heathrow, the first slot of the day may be a departure (rather an arrival) and therefore the time between slots may be greater. However, industry experts informed us that these airlines typically have other slots at Heathrow already, and could therefore adjust their existing slot portfolio in order to use these pairs of departures and arrivals.

There was a trade-off in determining the number of slots in the experiment, with a larger number increasing complexity and a smaller number reducing the amount of granularity that we could observe in the data.

We designed the experiment to allocate 48 daily slot pairs, which corresponds to 34,944 slots (ATMs) each year. As discussed in section 2, the release of slots at Heathrow will be phased. Phase 1 (2022–26) is expected to involve the release of 25,000 ATMs, while another 62,000 ATMs will be released in 2027

as the start of Phase 2.[54] The allocation of 48 daily slot pairs is therefore analogous to Phase 1 and some of the first slots to be released as part of Phase 2.[55]

The allocation of slots across the different time periods is set out in Table 5.2 below.

Table 5.2 Slots and time periods

	0600–0800	0800–1000	1000–1400	1400–1800	1800–2200
Duration (hours)	2	2	4	4	4
Number of slots	6	6	12	12	12

Source: Oxera.

As discussed in section 5.6 below, the slot coordinator reserved up to 50% of slots for new entrants, with each new entrant able to two slot pairs before losing new entrant status. Therefore, in total, given that there were three new entrants, the new entrants could obtain up to six slot pairs out of the 24 slot pairs reserved for new entrants by the slot coordinator. The demand for new entrant slots was therefore less than the supply of new entrant slots, which is what a number of stakeholders told us they expect to happen (given the design of the new entrant rule). The excess supply of new entrant slots was added to the pool of general slots.

In the experiment, participants only made decisions with regards to new slots that are being made available with the new runway. Including current slot holdings would have added a 'slot management' dimension to the experiment, the effect of which could be conflated with the slot allocation mechanisms that were tested. It would also have made the experiment more complex for participants to understand.

5.3 Representative routes

5.3.1 Number of routes

We included ten participants (airlines) in the experiment. Based on this, and the number of slots awarded (48 daily pairs), we determined the number of routes to include in the experiment. We defined routes at the city level.

As shown in Table 5.3, there was an overall ratio of 1.55 airlines per route and 3.41 departures per route at Heathrow in Summer 2018. 1.55 airlines per route would imply approximately 15 routes for ten airlines, and 3.41 departures per route would imply 14 routes given 48 slot pairs.

Table 5.3 Number of airlines and departures per route (Summer 2018)

Region	Airlines per route (mean)	Daily departures per route (mean)
Domestic	1.38	6.90
Europe	1.38	4.33
Middle East and Asia	1.69	2.33
North America	1.94	3.53
Other	1.33	1.18
All	**1.55**	**3.41**

[54] Heathrow Airport Ltd (2019), 'Preliminary Environmental Information Report: DCO Project Description', volume 1, chapter 6, p. 6.81.
[55] The experiment focused on the new capacity, abstracting away from re-times of existing slots.

Source: Oxera analysis of OAG data.

It is expected that the number of routes flown will increase following the release of new slots.[56] We therefore decided to have 20 routes in the experiment.

5.3.2 Number of routes by region

Given that there were 20 routes in the experiment, we needed to determine how many routes would be offered to different regions.

As shown in Table 5.4, routes in Europe account for 41% of all routes offered at Heathrow, and ATMs to these destinations collectively account for 52% of all ATMs. However, as the new capacity at Heathrow is expected to lead to a greater number of long-haul (and therefore not European) routes, we included only seven European routes in the experiment.[57] Middle East and Asian regions collectively account for a similar proportion as North America (17–25% each). We therefore included five routes for North America and five routes for Middle East and Asia.

While routes in the UK account for only 4% of all routes offered at Heathrow and flights to these destinations account for only 8% of flights, we included a larger proportion (three routes) in the experiment. This was because domestic connectivity is a policy aim for government, and so the option to fly to several domestic routes should be available to airlines (even if the profitability of doing so is low—see below).[58] This also allowed us to assess the effect of different slot allocation mechanisms on the share of domestic flights.

Table 5.4 Number of routes and departures per day by region (Summer 2018)

Region	Number of routes	Total departures per day
Domestic	4%	8%
Europe	41%	52%
Middle East and Asia	25%	17%
North America	17%	18%
Other	12%	4%
All	**100%**	**100%**

Source: Oxera analysis of OAG data.

5.3.3 Number of new routes

Given the number of routes in each region, we decided on the number of new routes available at Heathrow in the experiment (i.e. routes that are not currently flown by any airline).

As shown in Table 5.5, a number of new routes have been introduced each Summer season over the last few years, with the largest number of new routes being introduced in Summer 2018 (12 new routes), equivalent to 6% of all routes flown from Heathrow. However, the proportion of new routes is expected to be greater with the new runway at Heathrow as slots become less

[56] For example, the Airports Commission stated that the opening of new routes would be a key strategic benefit from the new runway: 'new routes and services will enable trade to increase with a wider range of countries and regions across the world.' Airports Commission (2015), 'Final Report', July, para. 6.28.
[57] Department for Transport (2017), 'UK Aviation Forecasts', October.
[58] Department for Transport (2018), 'Aviation 2050 The future of UK aviation', December.

scarce.[59] We therefore decided to have five new routes (25% of the routes) in the experiment.

Table 5.5 Number of new routes from Heathrow

Summer season	Number of new routes opened (compared with the previous Summer season)
2013	7
2014	8
2015	6
2016	11
2017	11
2018	12
Total	**55**

Source: Oxera analysis of OAG data.

5.3.4 Summary of representative routes

We selected the 20 routes shown in Table 5.6. The 20 cities were chosen to be representative of different destinations in a region, and were selected to be consistent with the ten representative airlines in the experiment. For example, the strategy of Euro Airlines was to fly to German cities from Heathrow. Two of the representative routes were therefore to German cities.

Table 5.6 Representative routes in the experiment

Region	Number of routes	Existing routes	New routes
UK	3	Belfast, Edinburgh, Inverness	n.a.
Europe	7	Amsterdam, Frankfurt, Munich, Palma, Santorini	Rotterdam, Bern
North America	5	Seattle, New York, Chicago, Los Angeles	Orlando
Middle East and Asia	5	Beijing, Dubai, Shanghai	Al Ain, Hangzhou
Total	**20**	**15**	**5**

Note: routes included in the experiment were selected for their representativeness and ease of understanding by participants. They do not necessarily represent the intentions of the example airlines or airlines that we spoke to as part of the stakeholder engagement.

Source: Oxera.

5.4 Airline payoffs

Airline payoffs were described to participants in the experiment as the sum of a base element and a special element. The base element included factors that did not depend on the slots that other airlines obtained. The special element included these interactions and complementarities. Figure 5.1 sets out the factors that were included as part of the base and special elements, while further information is provided below.

[59] For example, the Airports Commission has stated: 'The airport accommodates more flights on its two runways than any equivalent airport in the world, but has for many years been operating at the limit of the capacity that can be provided. This has reduced its ability to accommodate new routes […]'. It also stated that: 'due to capacity constraints at Heathrow, the airlines operating at the airport, in particular BA and its partners whose hub operation is based there, find it difficult to expand their current networks. Slots at Heathrow are very rarely available and therefore, whenever a new route is launched, it has to be at the expense of a different service which needs to be cancelled or replicated elsewhere.' Airports Commission (2015), 'Final Report', July, p. 19, 81.

Airline payoffs are reflected though the payoffs on the new slots, rather than changing their payoffs for existing slots. While we recognise that the choice of where to fly will impact the payoff of existing slots, the exact nature is difficult to determine. For example, more slots operated to one destination may help an airline reach critical mass on a route, or it may cannibalise its existing routes to some extent. In general, we expect that the impact on profitability of existing slot(s) is likely to be outweighed by the direct profitability of the new slot(s). Therefore we did not include this interaction in the experiment.

Figure 5.1 Airline payoffs

Base element	Special element
• slot time • choice of route	• route-level competition • route-level volume effects • slot complementarities • maximum number of slots

Source: Oxera.

5.4.1 Base element

Time period

As discussed in section 3, slot time is an important determinant of the value of the slot. This is because demand is greater for slots at particular times—e.g. because of typical business hours.

There is relatively little publicly available evidence on the value of slots by time period, as information on slot trades tends to be scarce and airlines are not required to inform ACL about the value of the slot trade. We also understand that slots can be traded for both money and non-monetary assets (e.g. slots at another airport), making valuation difficult without full disclosure of the transaction.

There are, however, several sources which show the slot values over the course of the day:

- a 2015 presentation by Heathrow which states that midday slots are worth 30% less than early morning slots at Heathrow, and that evening slots are worth 50% less than midday slots.[60] However, it appears that these results are based on a relatively small sample (five slot trades);[61]
- a 2019 study conducted for HAL, which found the same result with a sample of eight slot trades (excluding two of the eight as outliers).[62] This analysis relies on slot prices as reported in the media.

We also asked stakeholders about the value of slots by time period. Our discussions with stakeholders demonstrated that (at a high level) airlines have similar time preferences in that morning slots are particularly valuable. The value of the first slots in the morning appear to be much greater than slots later in the morning. The strength of apparent preference for early morning slots implies a greater difference in slot value by time of day than identified in the analysis listed above. This could be because the analysis above is based on a

[60] Sarah Whitlam, Head of Network Development at Heathrow Airport Limited (2015), 'An introduction to slot trading', PowerPoint presentation, July.
[61] It is unclear which slot trades are used for the analysis.
[62] Frontier Economics (2019), 'Estimating the congestion premium at Heathrow', May.

small sample, or because it does not take into account non-monetary aspects of trade.

Stakeholders also stated that evening slots are more valuable than slots in the middle of the day (although they are less valuable than morning slots). The exception to this pattern is Middle Eastern and Asian airlines, which have a higher value for slots late in the day than other airlines (such that evening slots are of equal value to the most valuable morning slots). This is because of both the existing 'wave' design of Middle Eastern hub airports, and desirable arrival times in the Far East.

Therefore we designed the payoff multiplier by time period to reflect the feedback from stakeholders. The value of the evening slots was designed to be greater for Middle Eastern and Asian airlines, reflecting the strong preference for these slots.

The slot value by time period in the form of a multiplier for the different airline types is shown in Table 5.7.

Table 5.7 Payoff multiplier by time period

	6am–8am	8am–10am	10am–2pm	2pm–6pm	6pm–10pm
Payoff multiplier (all airlines except Middle East and Asian airlines)	4.0	2.0	2.0	1.0	1.5
Payoff multiplier (Middle East and Asian airlines)	4.0	2.0	2.0	1.0	4.0

Source: Oxera.

Route choice

The choice of route is an important determinant of airline payoffs, and route-level profitability varies by airline. In setting route-level profitability we used RDC data[63] on routes from Heathrow, cross-checking this with data on routes from other London airports.

In applying this data in the experiment we made a number of adjustments. First, RDC data did not include all the types of airline represented in the experiment. For the airlines where we did not have data, we set the route-level profitability as an adjustment from the airlines for which we have data using insight from the stakeholder engagement. Second, new routes were set to be as profitable as the least-profitable current route in that region for that airline.[64] This is because it can be expected that airlines have already chosen to fly their most profitable routes. For example, if we observe that an airline flies to Munich but not to Bern, we can infer that flying to Bern is no more profitable for the airline than flying to Munich (or else the airline would fly to Bern).

Third, RDC data gives route profitability on a stand-alone basis (i.e. it does not account for the way in which short-haul flights feed long-haul flights). For example, a short-haul route may appear unprofitable on a stand-alone basis, but an airline may continue operating the route due to the transfer traffic it provides for profitable long-haul routes. This is likely to be the case for those airlines that use Heathrow as their hub airport, or that use links from Heathrow to feed their hub airport. We have therefore adjusted the data to reflect this dynamic. The route-level payoffs are shown in Table 5.8.

[63] RDC is an aviation consultancy that collects data on airline profitability.
[64] New routes are routes not currently offered from Heathrow.

Table 5.8 Route profitability

	National Airlines	Jet Flights	Atlantic Airlines	North American Airways	Speedy Flights	Quick Wings	Euro Airlines	Royal Flights	East Asia Airways	West Asia Airways
Inverness	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]
Edinburgh	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]
Belfast	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]
Amsterdam	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]
Rotterdam	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]
Frankfurt	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]
Munich	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]
Palma	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]
Santorini	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]
Bern	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]
Seattle	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]
New York	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]
Chicago	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]
Los Angeles	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]
Orlando	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]
Beijing	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]
Shanghai	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]
Hangzhou	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]
Dubai	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]
Al Ain	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]	[✂]

Source: Oxera analysis of RDC data.

Airlines' payoffs based on the combination of slot time and route choice were presented visually to participants, as shown in Figure 5.2. The on-screen payoff simulator also calculated the payoff for all combinations of route and time period.

Figure 5.2 The base element, as shown visually to participants (National Airlines)

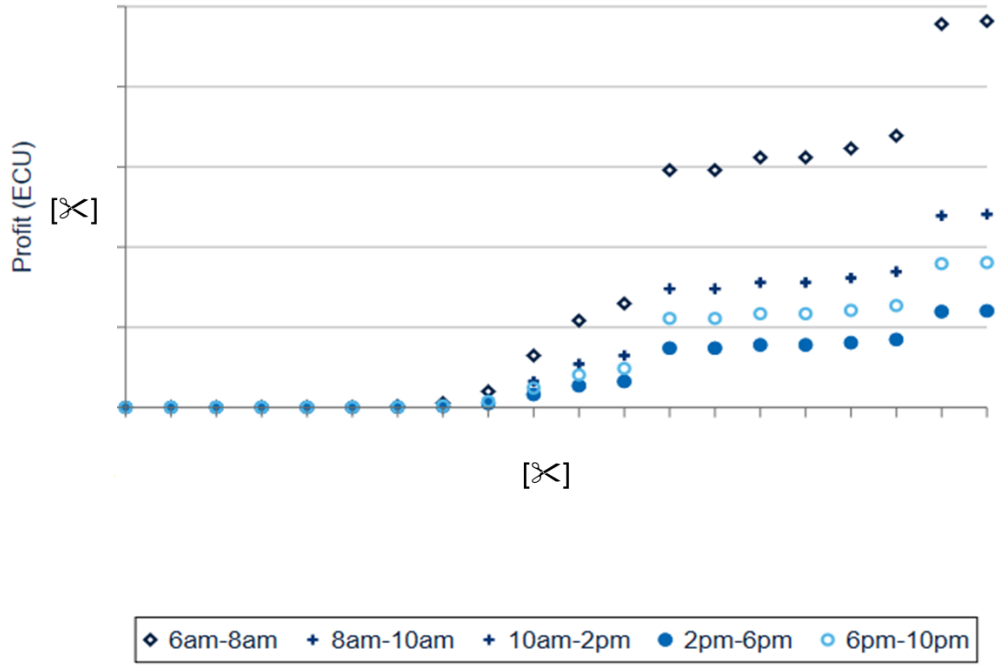

Note: This chart was included in both the screens and the key information sheets. Where the chart markers overlap, the payoffs were similar. The key information sheets also included a table with this data, so that participants knew the precise figures.

Source: Oxera.

5.4.2 Special element

Route-level competition

Route-level competition is a key determinant of slot profitability,[65] with route-level competition likely to be strongest within time periods. However, there is significant variation in the magnitude of this effect depending on the particular route, the number of airlines, and the identity of those airlines.

There is a well-established literature estimating the impact of route-level competition on fares (see Bruekner et al., 2013, for a review).[66] For example, in the American market, entry by LCCs onto a route has been found to reduce fares by up to 34%, with entry by 'legacy' airlines reducing fares by up to 5%.[67] Another study found that the entry of a low cost carrier onto a route at the

[65] For example, see Borenstein, S. (1989), 'Hubs and High Fares: Dominance and Market Power in the U.S. Airline Industry', *The RAND Journal of Economics*, **20**:3, pp. 344–365.
[66] Brueckner, J.K., Lee, D. and Singer, E.S. (2013), 'Airline competition and domestic US airfares: A comprehensive reappraisal', *Economics of Transportation*, **2**:1, March, pp. 1–17.
[67] Brueckner, J.K., Lee, D. and Singer, E.S. (2013), 'Airline competition and domestic US airfares: A comprehensive reappraisal', *Economics of Transportation*, **2**:1, March, pp. 1–17.

same airport lowers fares by 46%, with entry onto a route at an airport that is 'near' the airport in question lowers fares by 15%.[68]

Another study, reported in the 2007 European Commission Decision on the proposed Ryanair/Aer Lingus merger, found that Ryanair's presence on airport-to-airport routes out of Dublin reduced Aer Lingus' fares and load factors by 5–10%.[69] Further, Ryanair's presence on city-to-city routes out of Dublin reduced Aer Lingus' fares and load factors by 0-5%. The presence of other airlines on Aer Lingus' fares was found to be less than that of Ryanair.

Stakeholders told us that the impact of competition is likely to be greater when airlines fly routes at similar times of day, and that the impact of using multiple slots on a route is likely to increase the level of competitive pressure on other airlines. In other words, we might expect the impact of competition to be heightened within a particular time period (e.g. 6–8am), and when a competing airline operates more slots on that route.

We therefore made a simplifying assumption based on the range of estimates in the empirical literature regarding the strength of the effect of competition on airline payoffs: each additional slot used by another airline for the same route and in the same time period reduced the slot payoff by 15%. This effect applied in a log scale (i.e. three slots used by other airlines on that route at that time period reduced the slot payoff by 39%).[70]

Route-level volume effects

Similarly, route-level volume effects are a key determinant of the profitability of a slot. For example, if an airline were to use many slots during the same time period on the same route then profitability may decrease unless there is significant unmet (or increasing) demand.

This effect is likely to be less than the impact of competition between airlines, as within-airline effects will not involve price competition. Stakeholders confirmed this to be the case, although it is difficult to observe this empirically as airlines can be expected to operate more flights on more profitable routes (all else being the equal). We therefore made a simplifying assumption in the experiment regarding the strength of the route-level volume effect: each additional slot used by that airline for the same route and in the same time period reduced the payoff of the slots on that route by 5%.[71]

[68] Morrison, S.A. (2001), 'Actual, Adjacent, and Potential Competition Estimating the Full Effect of Southwest Airlines', *Journal of Transport Economics and Policy*, **35**:2, May, pp. 239–256.
[69] Results as reported in European Commission (2007), 'COMMISSION DECISION of 27/06/2007 declaring a concentration to be incompatible with the common market and the EEA Agreement (Case No COMP/M.4439 – Ryanair / Aer Lingus)', C(2007) 3104, Annex III 'Price Correlation Analysis for Geographic Market Delineation'.
[70] The effect of competition is based on incremental slots obtained from the third runway and not current slot holdings. For example, suppose an airline faces two competitor flights on the same route and in the same time period. After the third runway, the competitor adds another flight. The effect of competition on payoffs is 15% rather than 39%. This is because the RDC data on route-level profitability already takes account of the impact of existing competition.
[71] This effect applied in a log scale (i.e. three slots used by other airlines on that route at that time period reduced the slot payoff by 14%). The impact of route-level volume effects was applied after the impact of competition. The effect of volume considered new slots allocated by the slot coordinator in the experiment, excluding slots that airlines already held. This was because the data on route-level profitability already takes account of the impact of existing volume.

Slot complementarities

As described in section 3, stakeholders told us that the value of slots was complementary in certain situations. Where this complementarity was clear, we included it in the experiment, as described below.

- **Ratio of short-haul to long-haul flights (for National Airlines and Jet Flights)**. Due to the nature of using Heathrow as a hub airport (and the way that short-haul flights 'feed' long-haul flights), if the airline achieved a final ratio of short-haul to long-haul flights between [✂]:1 and [✂]:1 (i.e. a ratio of [✂]:1 was included), then the payoff increased by 25% on all slots. Short-haul flights were those to the rest of the UK and Europe. Long-haul flights were those to Asia, the Middle East and the USA. It is difficult to empirically observe the magnitude of this factor for hub airlines. However, based on the stakeholder engagement, we expect that the ratio will have a material effect on profitability. [✂]. Therefore the factor was chosen to be more important than the impact of one slot of competition (i.e. greater than 15%). However, stakeholders also expected that this ratio would not outweigh the incentive to fly profitable long-haul routes. Therefore the factor was chosen to be small enough such that route-level profitability would remain the primary driver of route choice for these airlines.
- **Rotations per day (for Speedy Flights)**. Due to the nature of short-haul low-cost operations, if the airline achieved rotations of three slots in a day then the payoff on those three slots increased. Specifically, the payoff increased by 25% for every group of three slots in these three time periods: 6am–8am; 10am–2pm; 6pm–10pm.[72] The impact of rotations at Heathrow on profitability is difficult to empirically observe, but it is likely to be important. Therefore the factor was set to be more important than the effect of one slot of competition (i.e. greater than 15%). However, we expect that low-cost short-haul carriers would operate slots at Heathrow even if the slots do not perfectly fit their preferred rotations pattern. Therefore the factor was chosen to be small enough such that route-level profitability would remain the primary driver of route choice.

Maximum number of slots

As described in section 3, stakeholders informed us that there would be constraints in how many slots they would be able to use in any given phase of the slot release at Heathrow. These constraints would be due to a variety of factors (including fleet and crew availability).

Furthermore, a number of airlines suggested that they would not want more than a certain number of slots at Heathrow, even if the slots were costless to obtain. This is because these airlines have hubs at other airports, and a large number of Heathrow slots would be of limited use (as the slots would be used primarily to feed their hub airport). They also face limits in their ability to finance the acquisition of new slots.

The special element of the payoffs therefore included a maximum number of slots that an airline could fly and earn a payoff from.[73] However, this did not

[72] Similar flight patterns by low-cost short-haul carriers have been observed at other London Airports. For example, see Airports Commission (2015), 'Final Report', July, para. 9.45.
[73] The maximum number of slots element did not take account of pre-existing slots held at Heathrow. The use of a single cut-off point (the maximum number of slots) for each airline was a simplifying assumption, as in reality there may be more gradual diminishing marginal returns to holding slots.

restrict airlines from requesting or obtaining more slots that could then be traded. The maximum number of slots were determined as follows:

- BA indicated that [✂] and industry experts stated that the large hub carrier would be expected to be able to operate a large number of new slots at Heathrow;
- Virgin Atlantic indicated it would like [✂]% of the new slots;
- American Airlines indicated that it would like up to [✂]% more slots that they currently hold;
- EasyJet have publicly indicated that it would like 30,000–55,000 ATMs (equating to 11–20% of the new runway capacity). Further, industry experts indicated that low-cost short-haul airlines may have the ability to switch a large proportion of their flights at other London airports to Heathrow at short notice;
- Norwegian indicated that it would like [✂] daily pairs, with possibly more in the longer term;
- KLM, Lufthansa and Finnish Airlines indicated that [✂];
- Industry experts expected that the Middle East and Asian airlines would not want more than two daily pairs at the first release of slots.

The maximum number of slots that could be flown by each airline is shown in Table 5.9. The total demand for slots was 96 slots, which exceeds the supply of slots of 48.

Table 5.9 Maximum number of slots, by airline

Airline	Maximum number of slot pairs that can be flown
National Airlines	[✂]
Jet Flights	[✂]
Speedy Flights	[✂]
Quick Wings	[✂]
Atlantic Airlines	[✂]
North American Airways	[✂]
Euro Airlines	[✂]
Royal Airways	[✂]
West Asia Airways	[✂]
East Asia Airways	[✂]
Total	**96**

Source: Oxera.

The airlines' payoffs did not include a minimum number of slots required in order to earn a payoff. This is because a number of stakeholders told us that there was still value to holding even one slot pair at Heathrow. Furthermore, given that the 48 daily slot pairs represented the first slots to be released at Heathrow, airlines would expect to increase their slot allocations over time (in the subsequent phases of slot releases).

5.5 Airline constraints

As described in section 3, airlines face budget constraints in the acquisition of slots in an auction and in the trading of slots. Based on our discussion with

various stakeholders, analysis of airlines' financial performance and experimental pilots, the budget constraints were set to be:

- relative to slot value. The budget constraints were set to provide a constraint in the acquisition of slots (i.e. so that an airline could not acquire all the slots), but to be large enough to allow an airline to acquire a number slots (if an airline paid its full valuation for the slots). The budget constraints therefore ranged from ECU 12.50 to ECU 275.00;
- relative to other airlines' budget constraints. Given the financial information presented in section 3.3, different types of airline are likely to have different abilities to acquire slots. Furthermore, airlines and industry experts expressed views over different airlines' abilities to raise large sums of money to finance slot acquisition at Heathrow in the long term.

The budget constraints in the experiment are shown in Table 5.10 and also described below.

- BA has a relatively high total revenue, and a positive net income. Moreover, Heathrow is the hub airport for BA and therefore stakeholders expected that BA will be able to raise significant funds for the purchase of new slots at Heathrow. Therefore BA had a large budget for new slots.
- Virgin Atlantic wishes to make Heathrow an important hub airport and therefore could be expected to be able to raise a large amount of funds for new slots at Heathrow.[74] However, Virgin has low total revenue and net income compared to other airlines. Therefore Virgin had a more limited budget than BA and the North Atlantic airlines.
- Low-cost short-haul airlines have relatively low total revenue and net income compared to other airlines. However they have a wide range for the WACC with the top of the range higher than other airlines. Therefore the low-cost short-haul carrier had a limited budget compared to long-haul carriers, but a higher budget than low-cost long-haul airline and the European airlines.
- Norwegian has a negative net income, and industry experts said that the low profitability would limit the budget for new slots. Norwegian also said that it would have very limited available budget for new slots at Heathrow. Therefore the budget of the low-cost long-haul airline was the lowest of all airlines.
- North Atlantic carriers have large total revenue and net income compared to other airlines. However, as Heathrow is not their hub airport, industry experts said their available budget for new slots would be more limited than BA. Therefore the budget of North Atlantic carriers was larger than most other airlines, but lower than BA and the Middle Eastern/Asian airlines.
- European carriers have a wide range of total revenue and positive net income. However, as Heathrow is not their hub, and given their limited willingness to pay for new slots at Heathrow, the budget of European carriers was lower than most airlines.
- The Middle East and Asian carriers have a wide range of total revenue and net income. However, industry experts stated that financial constraints were

[74] Virgin Atlantic (2019),' Virgin Atlantic plans to increase flights from London Heathrow'.

likely to be less relevant for these airlines, partly due to their state backed ownership and wider strategic motivations.

The release of new Heathrow slots may impact airlines' ability to raise finance and thus acquire new slots (through trading or an auction). This is because airlines currently have the opportunity to use slot holdings at Heathrow as assets against which they can raise finance.[75] However, the release of new Heathrow slots could be expected to decrease the value of existing Heathrow slots, thus reducing airlines' ability to raise finance for the acquisition of new slots. All else being equal, we would expect that this effect would be greatest for airlines with the largest holdings of Heathrow slots.[76]

However, the magnitude of the reduction in value of existing Heathrow slots is unclear. It may be the case that the initial phases of the slot release do not have a large impact on the value of existing slots (as the marginal increase in capacity is relatively low). On the other hand, the forthcoming increase in Heathrow capacity may already be impacting the value of existing slots. Therefore, given the inherent uncertainty, we did not explicitly take account of this dynamic in the experiment design.

Table 5.10 Budget constraints, by airline

Airline	Budget constraint (ECU)
National Airlines	[✂]
Jet Flights	[✂]
Speedy Flights	[✂]
Quick Wings	[✂]
Atlantic Airlines	[✂]
North American Airways	[✂]
Euro Airlines	[✂]
Royal Airways	[✂]
West Asia Airways	[✂]
East Asia Airways	[✂]

Source: Oxera.

Any budget left unspent contributed to participant payoffs (otherwise there would have been no opportunity cost to spending the constraint). Participants could not spend more than their budget constraint in the trade stages, meaning that they could not earn a negative payoff.

In reality, it is likely that airlines know how their budget constraint compares with those of other airlines, even if the exact budget constraints of other airlines are unknown. The participants were therefore shown this information on the screen and in the key information sheets. A screenshot of how this was presented visually is shown Figure 5.3.

[75] This applies at other UK airports as well. For example, The Times (2019), 'Norwegian puts Gatwick slots on offer', 3 September.

[76] We also note that this dynamic could reduce the value of slots at other London airports, if slots at other London airports are to some extent substitutable with slots at Heathrow.

Figure 5.3 Screenshot of budget constraint information (representative airline)

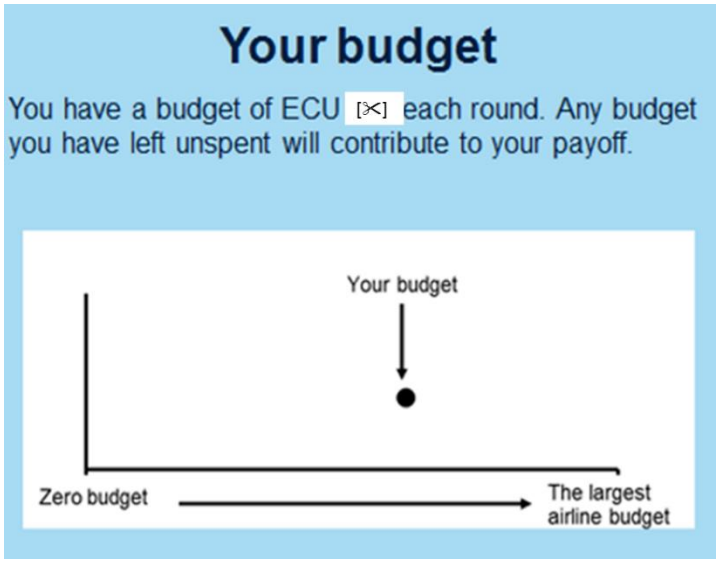

Source: Oxera.

5.5.1 Conversion rates

As described in section 4.4.3, each airline had a different ECU to £ conversion rate. These are listed in Table 5.11 below.

Table 5.11 Conversion rates from ECU to £

Airline	ECU per £1
National Airlines	[✂]
Jet Flights	[✂]
Speedy Flights	[✂]
Quick Wings	[✂]
Atlantic Airlines	[✂]
North American Airways	[✂]
Euro Airlines	[✂]
Royal Flights	[✂]
West Asia Airlines	[✂]
East Asia Airlines	[✂]

Source: Oxera.

5.6 The slot coordinator

The role of the slot coordinator (in treatments 1, 2 and 3 of the experiment—see the next section) was undertaken by an algorithm. We recognise that ACL weighs up the additional criteria on a case-by-case basis. However, for the purposes of the experiment it was necessary to have the slot coordinator act in the same way in every session (otherwise the results could have been due to the slot coordinator behaving differently, rather than differences between the treatments).

The algorithm described below is a close representation of our understanding of how ACL considers the additional criteria (e.g. competition) as set out in section 2. However, it abstracts away from considerations that are outside the experiment environment (e.g. the effective period of operation). We also

recognise that the rules ACL follows may not be codified to the extent required by an algorithm. Therefore the algorithm is necessarily an approximation of how ACL would take decisions.

5.6.1 Description of the algorithm

In assessing requests for slots, the algorithm proceeded as follows (also shown in Figure 5.4 below).

- First, the algorithm assessed aggregate supply and demand for slots in each time period. If demand exceeded supply in any time period, ACL considered new entrant slots (50% of slots). Airlines classified as new entrants were able to win up to two slot pairs of new entrant slots.
- Second, the algorithm assessed supply and demand for new entrant slots in each time period. If demand exceeded supply in any time period, the algorithm applied the additional criteria (as described below).
- Third, any remaining unallocated new entrant slots were added to the pool of 'general' slots to be allocated.[77]
- Fourth, the algorithm assessed supply and demand for general slots in each time period. If demand exceeded supply in any time period, the algorithm applied the additional criteria (described in section 5.6.2).

Figure 5.4 High-level description of the algorithm

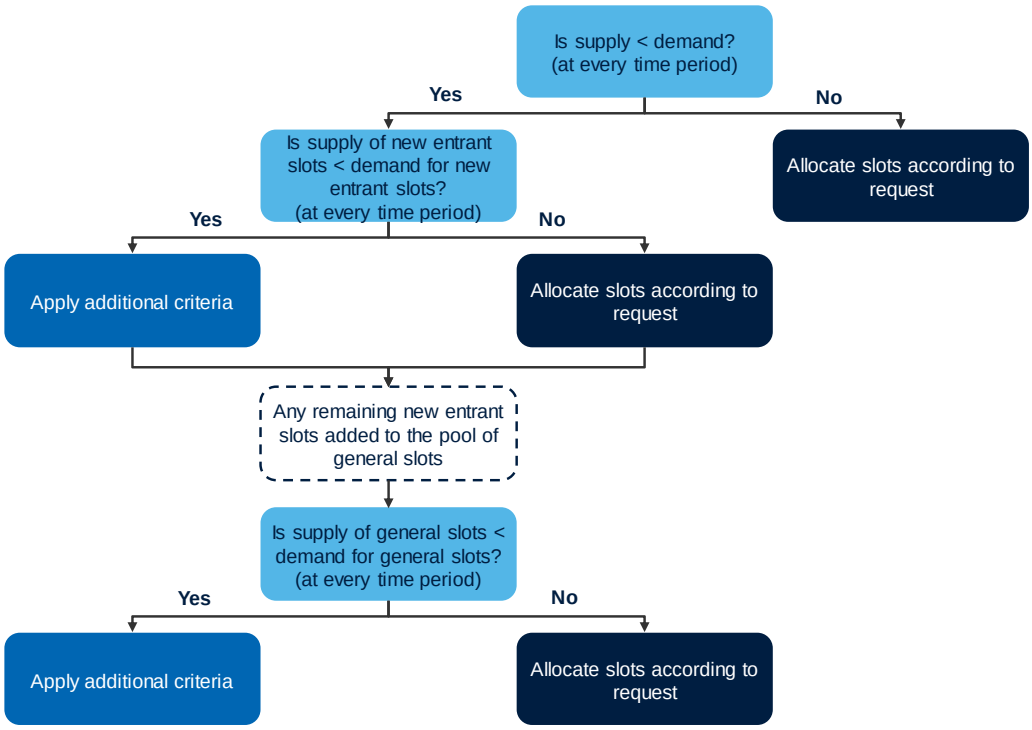

Note: The additional criteria are described in section 5.6.2.

Source: Oxera.

[77] The term 'general' slots is not used by ACL, but was introduced for the experiment so that participants could easily understand which slots were not new entrant slots.

This means that:

- new entrant slots were reserved equally throughout the day (i.e. 50% of every time period). This is because there is no evidence to suggest that ACL would reserve more new entrant slots at a particular time of day. However, the new entrant slots allocated depended on new entrants' requests for slots;
- if there were new entrant slots that were not awarded, these were added to the general pool of slots. It is our understanding that slots would not be left unused.

As shown in Figure 5.4 above, the new entrant slots were allocated before the general slots. This was an iterative process, proceeding one period at a time. The selection of a time period with which to begin the allocation process introduced an unpredictable element to the process. In a sense, this mirrors our understanding of how ACL would consider allocating slots in reality. In other words, ACL does not start by using an algorithm to compute the overall optimal solution given its criteria.

5.6.2 Application of the additional criteria

The algorithm applied the additional criteria based on an approximation of the current ACL rules. The application of the additional criteria, which is shown in Figure 5.6 below, had the following (descending) order of priority.

1. **Adding new routes from the airport**. New routes were prioritised over existing routes from Heathrow. The algorithm gave priority to new routes to the destinations with the largest forecast growth in demand over the coming decades. The DfT unconstrained demand forecasts, shown in Table 5.12, indicate that demand for the Middle East and Asia will grow fastest, followed by demand for Europe and then for North America. The algorithm therefore gave first priority to new routes in the Middle East and Asia, second priority to new routes in Europe, and third priority to new routes in North America. The algorithm considered new routes to be 'new' regardless of how many slots it allocated to each new route. This meant that the algorithm prioritised capacity to new routes beyond the first flight to that destination.[78]

Table 5.12 Unconstrained passenger demand forecast, by region (million passengers departing Heathrow per annum)

Region in the experiment	Unconstrained demand forecast for 2030	Unconstrained demand forecast for 2050	Percentage increase in demand
Middle East and Far East	18.0	22.9	+27.2%
Western Europe	58.2	61.7	+6.0%
North America	25.2	24.9	-1.2%

Note: Western Europe refers to all EU and non-EU countries in Europe (including Russia and Eastern European countries). North America refers to Canada and the USA. Middle East and Far East excludes India, New Zealand and Australia as they were not included as destinations in the experiment.

Source: Department for Transport (2017), 'UK Aviation Forecasts', October.

[78] This is a simplification of how we might expect the slot coordinator to consider new routes in reality. For example, ACL may stop considering a route as 'new' after allocating a certain number of slots for that route. However, it is unclear which order of slot requests ACL would consider in making this decision, and it is also unclear how many slots would be required before the route ceased to be considered as 'new'.

2. **Increasing the number of airlines on existing routes with priority for routes where there are fewer current airlines**. The algorithm started from the position of the pre-existing routes, and dynamically updated it in accordance with the slot allocations. The same information was also given to participants in the instructions (on the screens and in the key information sheets)—see Figure 5.5 below.

   The data on pre-existing routes was drawn from Summer 2018 OAG data on flights from Heathrow to destinations defined at the city level. Some of the data is the result of the combination of several airlines in a given airline group because there is not a one-to-one relationship between airlines in reality and the ten representative airlines in the experiment. There is also not a one-to-one relationship between the routes flown by airlines and the 20 representative routes in the experiment. For example, the North Atlantic airlines (Atlantic Airlines and North American Airways) were each told that they had existing flights to three of the five North American airports in a pattern representative of how a number of North Atlantic airlines fly to North American destinations in reality.

Figure 5.5 Screenshot of the table of pre-existing routes

Route	Country	New route from Heathrow?	National Airways	Jet Flights	Atlantic Airlines	North American Airways	Speedy Flights	Quick Wings	Euro Airlines	Royal Flights	West Asia Airways	East Asia Airways
Inverness	UK		1									
Edinburgh	UK		10									
Belfast	UK		4									
Amsterdam	The Netherlands		8							10		
Rotterdam	The Netherlands	*Yes*										
Frankfurt	Germany		6						12			
Munich	Germany		6						8			
Palma	Spain		2									
Santorini	Greece		1									
Bern	Switzerland	*Yes*										
Seattle	The USA		2	1								
New York	The USA		10	7	4	5						
Chicago	The USA		2		4	3						
Los Angeles	The USA		3	3	2	1						
Orlando	The USA	*Yes*										
Beijing	China		1									
Shanghai	China		1	1								
Hangzhou	China	*Yes*										
Dubai	United Arab Emirates		3	1							6	
Al Ain	United Arab Emirates	*Yes*										

Source: Oxera.

3. **Increasing the number of slots held by the airline with the fewest number of slots on that route**. The algorithm started from the pre-existing routes, and dynamically updated it in accordance with the slot allocations.

To the extent that there were airlines with identical priority based on these criteria, the allocation was randomly decided.

Figure 5.6 Application of additional criteria to each slot pair

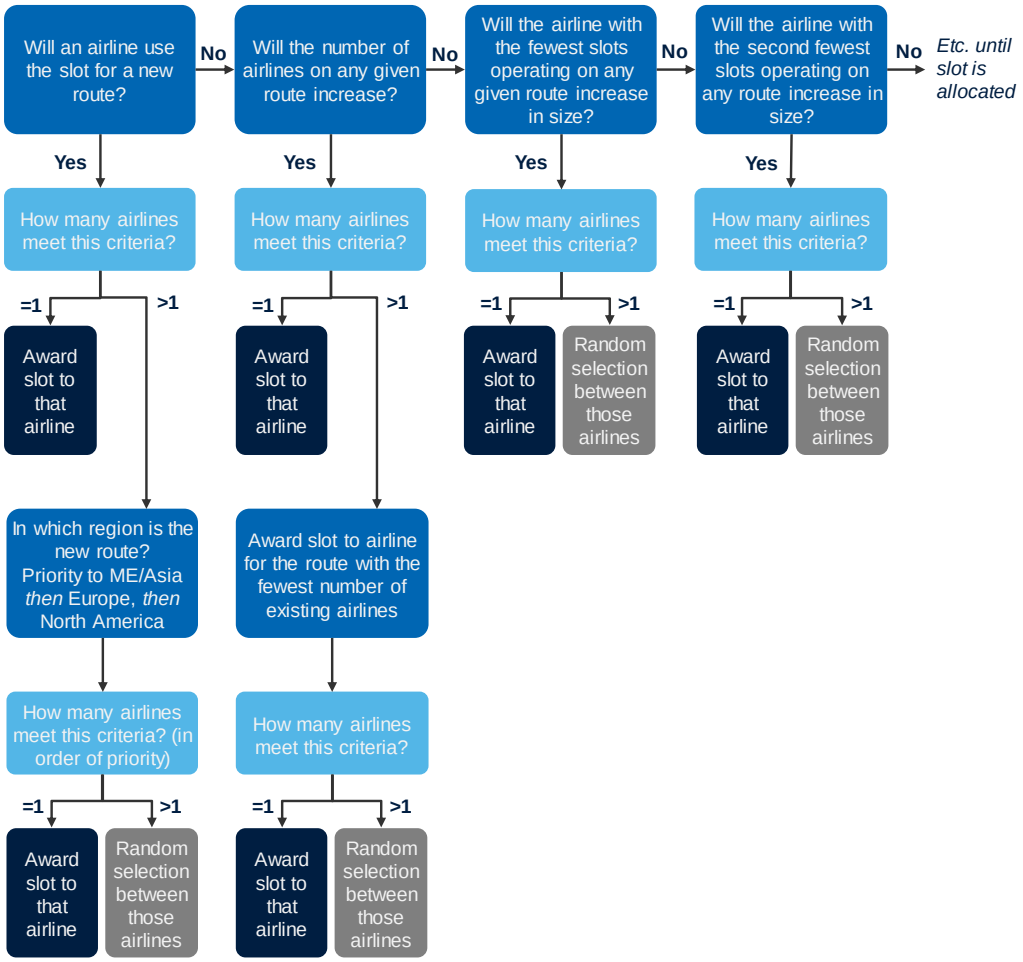

Note: New routes were defined at the city level (for example, as there were no pre-existing flights from Heathrow to Bern, an application to use a slot for Bern would count as a 'new route').

Source: Oxera.

5.6.3 Describing the slot coordinator to participants

The slot coordinator algorithm determined the outcome of the bidding stage(s) in all but one of the treatments (see section 7). It was therefore important that participants understood the role and nature of the slot coordinator.[79]

Participants were provided with high-level information about what the algorithm prioritised (analogous to what factors ACL considers most important), without revealing the exact way in which the algorithm did this. This reflects reality in that the decision-making process is not fully transparent to airlines. Airlines informed us that they were aware of the main considerations that ACL takes account of when allocating slots using the additional criteria (e.g. prioritising new routes), but that the precise way in which the different priorities is weighed is unclear.

The information regarding the slot coordinator algorithm was stated as follows in the instructions.

[79] The slot coordinator algorithm was described as 'the computer' to participants in order to communicate that it was an automated decision-making process but to avoid complex language.

The rules that the computer follows are partially known by the airlines:

- The computer gives priority to new routes from Heathrow (routes where no airline currently flies).
- New routes to Asia are given higher priority than new routes to Europe. New routes to Europe are given higher priority than new routes to the USA.
- Priority is given to routes your airline is not already flying to (but other airlines may fly to).
- Priority is given to the smallest airline on a route (e.g. if Airline A has 1 slot on that route, and Airline B has 10 slots on that route, Airline A would have priority).

In reality, airlines have the opportunity to state that they will use the requested slot for one route, but after obtaining the slot actually use the slot for another route.[80] This dynamic was included in the experiment, as participants were told that they did not have to fly the route that they stated to the slot coordinator.

[80] 'Once a slot is allocated and used for a specific period of time, the airline is able to change use to operate on a different route to the one that was agreed with the administrator when the slot was initially allocated. There is no mechanism to sanction an airline if a slot is not used as an airline said it would be, or other incentives for airlines to be truthful in their submissions.' Competition and Markets Authority (2019), 'Aviation 2050 Response from the Competition and Markets Authority', June, para. 4.18.

6 Design of the experiment treatments

In this section we outline the four treatments in the experiment. The first section (6.1) provides a high-level overview of the four treatments, while the subsequent sections (6.2–6.5) explain the key features of the experiment design for each treatment in more detail.

Screenshots of the experiment and key information sheets can be found in Appendix A2. A review of the existing literature on slot allocation mechanisms can be found in Appendix A3.

6.1 The four treatments

Behavioural experiments estimate the impact of different factors by testing one factor at a time, which ensures that different effects cannot be conflated. Each treatment therefore varies (from another treatment) in terms of only one factor. Any given factor in the experiment design that is held constant across the treatments should not affect the results. Otherwise there would be considerable uncertainty over which factor is driving the results.

In our experiment, we included a baseline treatment that represents the current slot allocation system. Treatments 2, 3 and 4 then tested the impact of changing one factor in each treatment relative to the baseline (or 'control') treatment.

The relationship between the four treatments in this experiment is shown in Figure 6.1.

Figure 6.1 The four treatments

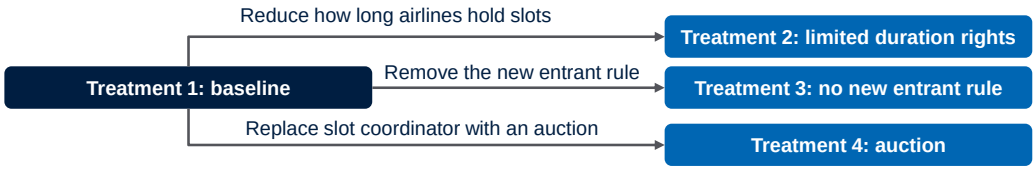

Source: Oxera.

These four treatments were chosen as being of the most interest to the DfT, and each was selected to give insight into a particular change to the slot allocation mechanism. The design of the treatments was informed by the relevant literature (see Appendix A3).

Treatments 2 and 3 represent changes to the administrative mechanism for allocating slots: awarding slots for a limited duration only (treatment 2); and awarding slots without the new entrant rule (treatment 3).[81] Treatment 4 replaces the administrative mechanism with a market-based mechanism (an auction).[82]

- **Treatment 1: baseline (or 'control')**. The baseline treatment was designed to be as close as possible to the status quo for slot allocation in the UK. Our understanding of the status quo was informed by discussions with ACL and

[81] The design of the new entrant rule and the grandfather rule have been critiqued in the literature. For example, see Haylen, A. and Butcher, L. (2017), 'Airport Slots', House of Commons Library Briefing Paper 488, 12 June.
[82] As recommended by the Competition and Markets Authority (2018), 'Advice for the Department for Transport on competition impacts of airport slot allocation', December.

other stakeholders, and by looking at publicly available information (such as the WSG).[83]

- **Treatment 2: limited duration rights**. In the status quo, slots are granted in perpetuity as long as the use it or lose it thresholds are met (and unless the airline chooses to sell the slot to another airline). Treatment 2 tested the impact of granting slots for a limited period of time only before the slots need to be handed back to the slot coordinator. Otherwise, treatment 2 was identical to treatment 1 (the baseline).
- **Treatment 3: no new entrant rule**. New entrants are reserved slots by the slot coordinator in the status quo. Treatment 3 tested the impact of removing the new entrant rule. Otherwise treatment 3 was identical to the baseline.
- **Treatment 4: auction**. In the status quo, the slot coordinator awards slots to airlines using an administrative mechanism (as opposed to a market mechanism). Treatment 4 tested the impact of using an auction instead of the administrative mechanism. Otherwise, treatment 4 was identical to treatment 1 (the baseline).

Each treatment is described in more detail below.

6.2 Treatment 1: baseline (or 'control')

There were a number of stages within each round, which enabled us to explore how an airline chooses to use the slots that it is allocated, in terms of trading and choice of route. The round structure is shown in Figure 6.2.

Figure 6.2 Treatment 1 round structure

Source: Oxera.

6.2.1 Bidding stage

As the allocation of slots for the new runway will start with airlines making 'bids' or 'requests' to ACL, each round in the experiment started with a bidding stage. Participants made bids through the use of drop-down menus. Each bid involved selecting the route that the airline would fly with the slot. Participants were not held to these route requests after obtaining the slot (and this was explained to them). The drop-down menus contained only the routes that each airline could fly (e.g. West Asia Airways could select only Dubai or Al Ain). This is shown in Figure 6.3.

[83] We recognise that ACL has not previously awarded a large release of new slots at a capacity-constrained airport in the UK, and there is therefore some uncertainty over how it would implement its current rules with respect to Heathrow's third runway. However, in the absence of information about how it would award slots, we incorporate the current rules into the baseline treatment.

Figure 6.3 Screenshot of the bidding stage

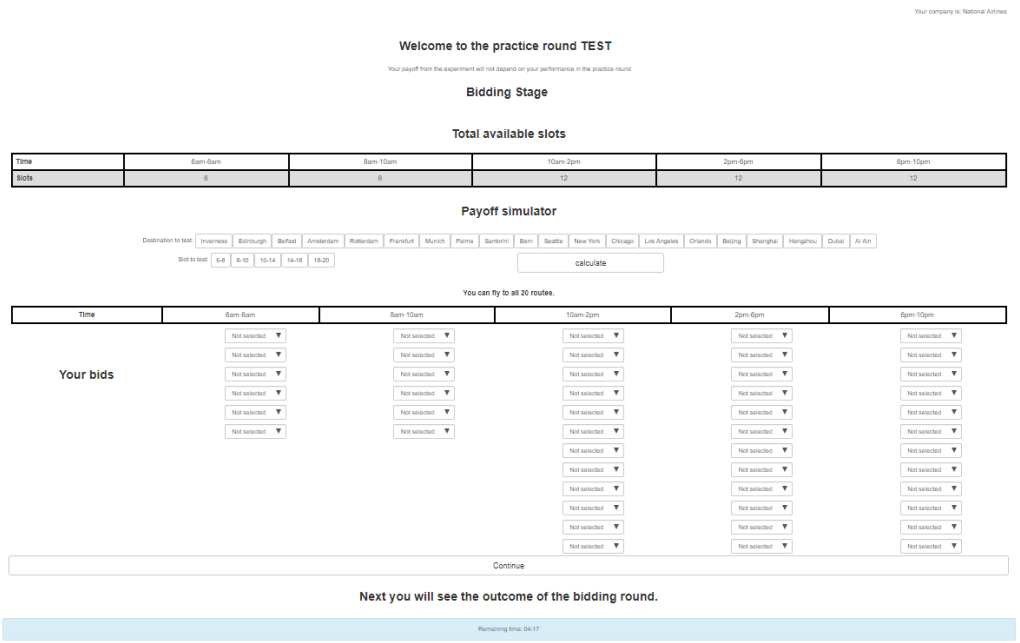

Note: The word ‘TEST’ appears in some of these screenshots because the screenshots were taken when the experiment was in ‘test mode’—it did not appear in the actual experiment.

Source: Oxera.

After the bidding stage, the participants learned which slots they had obtained.

6.2.2 Trade stage 1

As discussed in section 2, airlines operating at Heathrow have the opportunity to trade slots. We therefore provided two opportunities for trading in the experiment (trade stage 1 and trade stage 2). Trade stage 1 occurred before the participants had selected which routes they would actually fly (which might differ from their slot requests to the slot coordinator in the bidding stage), and therefore represented the time immediately after the bidding stage.

Trade stage 1 (and also trade stage 2) were designed as follows.

- **Bilateral agreements**. Various types of trading markets could have been used in the experiment, ranging from efficient clearing markets (where the identity of the counterparty is not known) to forums where bidders agree on a series of bilateral trades. As described in section 2, slot trades are bilateral between airlines, and the counterparty is known. We therefore designed the trading to be a series of buy offers posted on an open forum, where the buying airline’s name was published and participants could choose which airlines to sell slots to. Buy offers could be retracted by participants at any time during the trade stage.
- **Slot sales**. Slot trades are technically slot swaps. However, stakeholders informed us that many of these swaps involve the use of slots that cannot be flown, and are then handed back to ACL. Slot trades therefore become like slot sales in practice, and we designed the trade stage to be slot sales.
- **Time limit**. As described in section 2, while slot trades do occur at Heathrow (and at other UK airports), the number of slots that are traded is relatively limited. We therefore designed the experiment to include a mechanism to limit the amount of trading. One way would have been to

include a transaction fee when trading, whereby a certain cost is paid by the participants. However, it is unclear what level of cost would have been appropriate in the context of slot trades. Therefore we used a time limit in the experiment as a proxy for the 'stickiness' of the secondary market in reality.

- **New entrant slots**. As described in section 2, new entrant slots cannot be traded for two years after they are awarded. In the experiment, new entrant slots could therefore not be traded at trade stage 1. This is shown in Figure 6.4.

Figure 6.4 Screenshot of trade stage 1

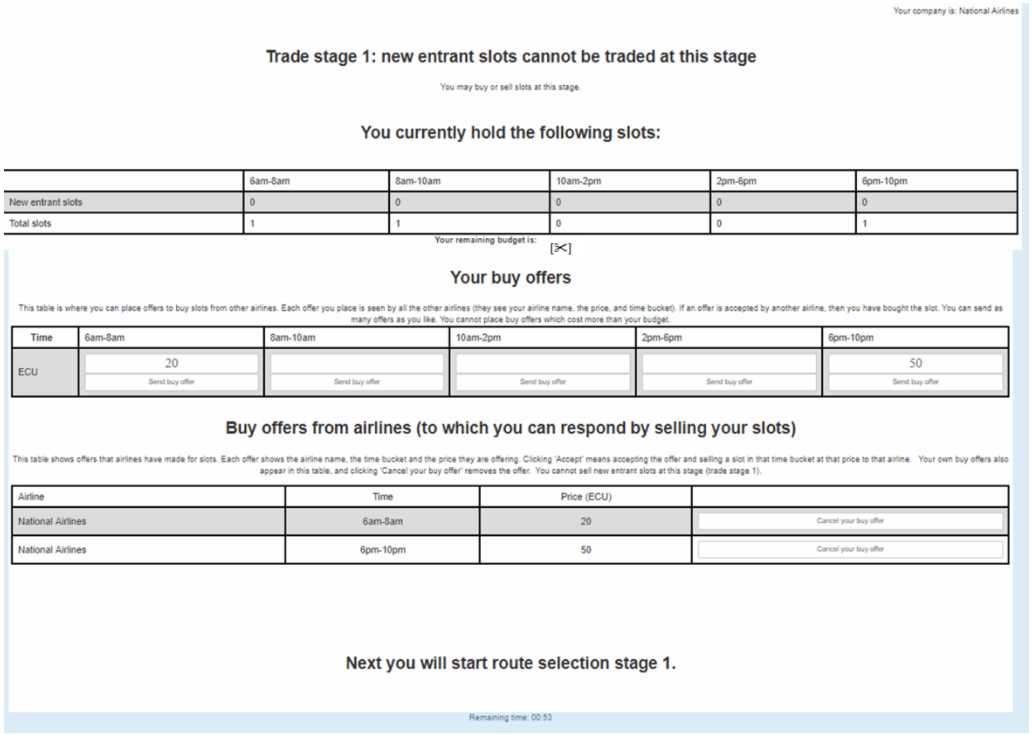

Source: Oxera.

6.2.3 Route selection stage 1

After obtaining slots through either the slot coordinator or trading, airlines need to choose how to use their slots (i.e. where they will fly to), even though they have already stated a route choice to the slot coordinator at the bidding stage. In the experiment, participants chose their routes in route selection stage 1 after trade stage 1. This took the form of a drop-down menu for each slot that they held.

The drop-down menus did not offer pre-selected routes (the default option had no route selected). The alternative would have been to set the default to be the bids that participants made to the slot coordinator in the bidding stage, but this approach has two drawbacks:

- some slots were obtained through trading, so no prior route would have been available;
- this would have signalled to participants that they should keep the same routes as those that they originally told the slot coordinator. Behavioural experiments have shown that individuals tend to be influenced by 'default bias', whereby the default option is strongly influential in their final choice.

Therefore, to avoid the default bias and to indicate to participants that they had a choice over route selection, no default routes were selected.

New entrant and general slots had the same level of flexibility over route selection. As described in section 2, in reality there are certain conditions under which an airline may be eligible for a new entrant slot because of its choice of route.[84] In such circumstances, the slot coordinator can restrict the participant's choice of route after the slot has been awarded. However, it is our understanding that this restriction applies only to a minority of new entrant slots, and therefore we did not include it in the experiment. This stage is illustrated in Figure 6.5.

Figure 6.5 Screenshot of route selection stage 1

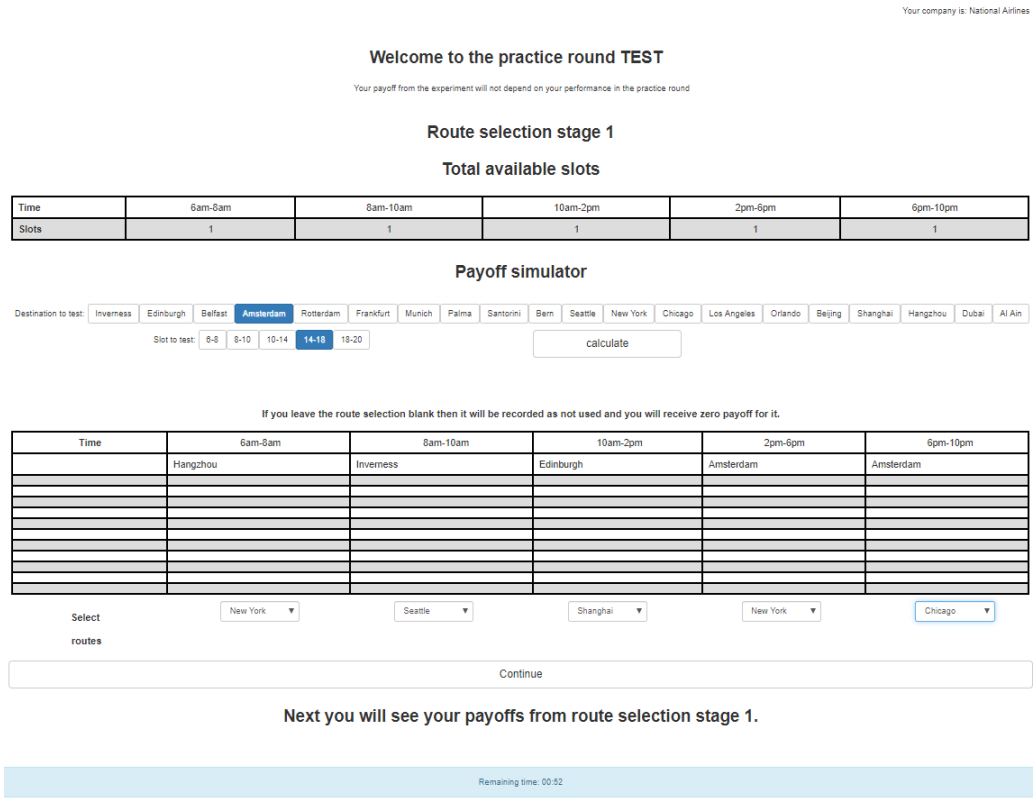

Source: Oxera.

After route selection stage 1, participants saw their payoffs.

6.2.4 Trade stage 2

Given that, in reality, airlines would have a continued opportunity to trade their slots, we included another trade stage in the experiment. Trade stage 2 differed from trade stage 1 only in that new entrant slots could be traded—trade stage 2 represented the time after the two-year restriction on new entrant slots ended.

6.2.5 Route selection stage 2

After trade stage 2, airlines again had the opportunity to decide where to fly with their slots. Route selection stage 2 was identical to route selection stage

[84] Such routes occur within the EU. See Council Regulation (EEC) No 95/93 of 18 January 1993 on common rules for the allocation of slots and Community airports, Article 2, para. (b) (ii).

1. As before, no ‘default’ routes were presented to participants, as airlines were free to change their routes.

Lastly, participants learnt their payoff from the round as a whole (the sum of their payoffs after the two route selection stages).

6.3 Treatment 2: limited duration rights

As described in section 2, in reality slots are awarded in perpetuity under the current slot allocation rules. This was reflected in treatment 1 in the experiment, where (within each round) participants did not have to hand back their slots (although they could sell their slots through trading). Treatment 2 tested the impact of awarding slots with limited duration.

In treatment 2, slots that were awarded in bidding stage 1 were handed back to the slot coordinator after route selection stage 1, after which there was another bidding stage (bidding stage 2) in which airlines could again bid for slots.

Bidding stage 2 replaced trade stage 2, reducing the opportunity for airlines to trade their slots. This is representative of reality, where limited duration slot rights reduce the amount of time in which slots are held and are therefore tradeable. However, we do not associate a specific number of years to the duration that the slots are held in the experiment. After bidding stage 2 there was route selection stage 2, where participants selected their preferred routes with the ‘new’ slots obtained.

Other than the replacement of trade stage 2 with bidding stage 2, treatment 2 was identical to treatment 1. For example, the new entrant rule still applied, meaning that new entrant slots could not be traded at the (only) trade stage. The round structure of treatment 2 in shown in Figure 6.6.

Figure 6.6 Round structure in treatment 2: limited duration rights

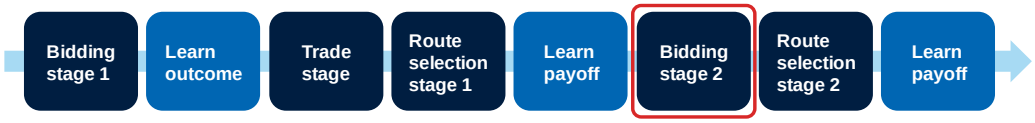

Note: The difference relative to treatment 1 is highlighted in red.

Source: Oxera.

In practice, awarding slots with limited duration may mean that airlines have less time to recover fixed costs of using the slot. In reality this may affect certain types of airline differently, for example new entrants may face higher fixed costs as they start operating from Heathrow. However, in order to make as few changes as possible between each treatment (and to ensure that the treatment effects cannot be conflated with other factors), the route payoffs in treatment 2 were identical to that of treatment 1.

6.4 Treatment 3: no new entrant rule

As described in section 2, the new entrant rule is a key element in the slot allocation mechanism. Treatment 3 removed the new entrant rule in order to test its impact. This involved the following changes (in comparison to treatment 1):

- the experiment screens and key information sheets did not mention new entrant airlines, new entrant slots, or the new entrant rule;

- the slot coordinator algorithm did not start by allocating new entrant slots to new entrants;
- all slots could be traded at trade stage 1.

Otherwise, treatment 3 was identical to treatment 1 (e.g. it had the same round structure).

6.5 Treatment 4: auction

Treatment 4 tested the impact of replacing the administrative mechanism for allocating slots at the bidding stage with an auction. Treatment 4 therefore had the same round structure as treatment 1, with the difference being the nature of the bidding stage.

6.5.1 Auction design

There are many types of auction, with the appropriate design being dependent on the circumstances (e.g. the nature of the resource being auctioned, or the nature of competition). There is no 'one size fits all', and there is a well-established literature on auctions in behavioural experiments.[85]

The auction design in the experiment was determined following a review of the literature (outlined in Appendix A3) and discussions with stakeholders. Other auction designs may also be appropriate for slot allocation, depending on the precise context.

The auction design in this experiment was as follows:

- **auction format**—one-off sealed-bid combinatorial auction. The combinatorial auction (i.e. package bids) was selected because there are both complementarities and substitutability between slots. The auction was sealed-bid because there was little common value uncertainty (i.e. each airline knew their own value from the slots).[86] Sealed-bid auctions are also simpler to administer than multi-round auctions;
- **packaging**—flexible packaging (where airlines determined which slots would be in their bids). This was to accommodate the fact that airlines will have different preferences with respect to how to bundle slots together to meet their needs (e.g. long-haul versus short-haul carriers);
- **bidding language**—XOR (meaning 'exclusive or'), taking account of both complementarity and substitutability, and ensuring that airlines did not exceed their budgets. This means airlines could only win one of their package bids;
- **winner determination**—was solved using an approximate algorithm (the mathematical determination of the winner(s) in this auction design requires an approximation);

[85] Klemperer, P. (2002), 'What Really Matters in Auction Design', *Journal of Economic Perspectives*, **16**:1, pp. 169–189. For example, see Lusk, J.L. and Shogren, J.F. (2007), *Experimental Auctions: Methods and Applications in Economic and Marketing Research*, Cambridge University Press.

[86] We also note that while airlines knew the mechanism by which competition would impact their payoffs (i.e. a reduction of 15% for each competitor slot on that route in that time period), they did not know how many competitors would choose to fly any given route at any given time period. Therefore there was some value uncertainty over slot value ahead of the auction. However, this uncertainty would have been unlikely to have significantly reduced through observing the prices bid by other airlines at auction (as the other airlines could operate the slot on a variety of routes).

- **payments**—first-price payments. The theoretical efficiency benefits of second-price payments in this auction design are limited (given the use of an approximate winner determination algorithm and XOR bidding). First-price payments are also simpler for participants to understand.

6.5.2 Implementation of the auction

The auction was implemented in the following way.

- **Step 1: new entrant slots**. The first slots to be allocated were the new entrant slots. At step 1, analogous to the baseline treatment, only new entrant airlines could bid for slots, and they could bid for up to only two slot pairs. Winning a package at step 1 did not stop new entrants from bidding or winning packages at step 2. The remaining unsold new entrant slots were then added to the pool of general slots. After Step 1, all airlines were told which new entrants had won which slots and at what price.
- **Step 2: general slots**. The general slots included the remaining new entrant slots, and all airlines could bid for them (including new entrant airlines). There was no restriction on how many general slots an airline could win. After step 2, all airlines were told which airlines had won which general slots and at what price. This is shown in Figure 6.7 below.

Figure 6.7 Screenshot of the bidding stage in treatment 4, step 2

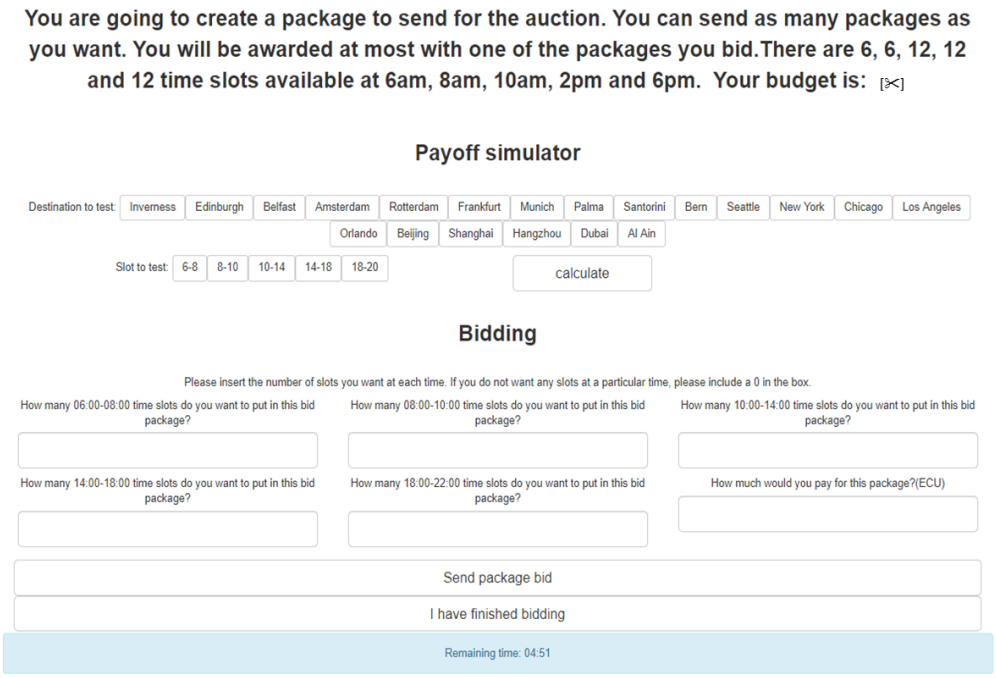

Source: Oxera.

The airline budgets (as described in section 5.5) were kept the same for all treatments, including treatment 4. This was because the constraints faced by airlines (financial and otherwise) could be expected to be the same regardless of the slot allocation mechanism. It is worth noting that while the auction bidding was combinatorial, airlines could subsequently trade individual slot pairs (which was also the case in treatments 1–3).

6.6 Conclusion

The four treatments were designed to robustly test the impact of three different policies affecting the slot allocation mechanism. They can be summarised as follows:

- **Treatment 1: baseline (or 'control')**. The baseline treatment was designed to be as close as possible to the status quo for slot allocation in the UK;
- **Treatment 2: limited duration rights**. Treatment 2 tested the impact of granting slots only for a limited period of time, before the slots needed to be handed back to the slot coordinator—i.e. removing grandfathering rights;
- **Treatment 3: no new entrant rule**. Treatment 3 tested the impact of removing the new entrant rule which is in place in the baseline treatment;
- **Treatment 4: auction**. Treatment 4 tested the impact of using an auction instead of the administrative mechanism. There are multiple ways to implement an auction, and in this treatment we considered a 'combinatorial auction' where participants are able to create package of slots.

Section 7 describes the results of these treatments.

7 Results

In this section we describe the results of the experiment. We start by summarising the descriptive statistics (section 7.1), before turning to analysis of competition (section 7.2), domestic and long-haul connectivity (section 7.3), efficient usage of Heathrow's capacity (section 7.4), trading behaviour (section 7.5) and bidding behaviour (section 7.6). Section 7.7 provides an overall summary of the results.

In general (and unless stated otherwise), the results relate to the allocation and use of the new slots (i.e. the 48 slots in the experiment), rather than all the slots at Heathrow.[87]

7.1 Descriptive statistics

In order for an experiment to be robust, any differences between treatments should be a result only of the treatments themselves, and not due to any intrinsic differences between participants. In our experiment we randomly allocated participants to treatments such that the participant mix should be similar across treatments.

This is confirmed by an analysis of the demographic questions asked at the end of the experiment, for which results are shown below (detailed results are included in Appendix A4). For example, we find that, across treatments, participants were similar in terms of education level, with no statistically significant variation (see Figure 7.1).

Figure 7.1 Participants' education

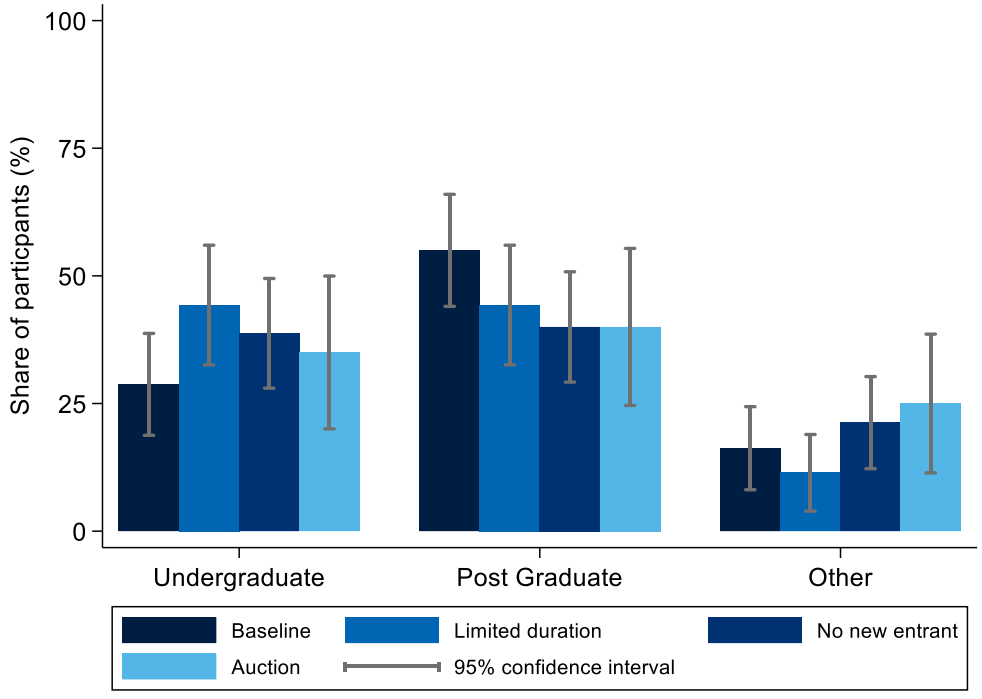

Note: The option 'Other' may have been selected where the participant was studying for a course that is less clearly defined as undergraduate or graduate, such as a Foundation

[87] This is because the experiment data cannot be linked to actual data at the level of granularity of airlines and destinations. In particular: airlines in the experiment cannot be linked to a specific airline; and destinations in the experiment do not capture the full range of destinations currently offered from Heathrow.

Certificate. A confidence interval is a way to measure the precision of an estimate. It is a range of values that contain the true statistical value (i.e. across all samples) with a probability of 95%.

Source: Oxera.

The sample size also needs to be sufficiently large such that any differences between the treatments can be accurately detected. We conducted the experiment multiple times for each treatment with different groups of participants each time. Across the experiment there were 30 different groups of participants, leading to a total of 300 participants.[88] This is shown in Table 7.1.

Table 7.1 Sample size

Treatment	Number of groups	Number of participants
1	8	80
2	8	80
3	8	80
4	6	60
Total	30	300

Source: Oxera.

In order to incentivise participants, they were offered a monetary reward based on their performance. Overall, participants earned between £10 and £15 from the experiment. This is shown in Table 7.2.

Table 7.2 Average payoffs

Treatment	Average payoffs (£)
1	10
2	15
3	12
4	10

Source: Oxera.

7.2 Competition

Competition in the aviation industry can be analysed in various ways, depending on whether it is assessed from the perspective of the airport, passengers, airlines, or public authorities. In this experiment, based on the data we collected, we assess the extent of competition at Heathrow in two ways.

First, we assess the level of competition on each route. Airline markets are often defined at the route level because substitution between origin and destination pairs is limited. For example, if one needs to travel between London and New York, a London to Chicago flight might not be considered as a substitute.[89]

Second, we assess competition at the airport level for all the newly released slots at Heathrow. This provides a measure of competition for slot ownership at Heathrow (irrespective of which routes the slots were used for).

[88] There were some technical IT issues in the lab. In these cases we excluded the group from the sample (i.e. did not use their data for the analysis), leading to a final sample of 30 groups of participants.
[89] However, we recognise that the level of substitution will differ by the segment of the market (e.g. business vs leisure passengers). Also, some long-haul routes from Heathrow may be substitutable with other European hub airports.

7.2.1 Competition at the route level

When analysing competition, one of the most common metrics to consider is market shares. Figure 7.2 shows the average route-level market share at Heathrow for each airline.[90]

In general, we see high average route-level market shares (i.e. mostly above 50%), indicating that airlines tended to focus on routes where they faced little competition. The no new entrant rule and auction treatments tend to reduce the average route-level market shares (which is reflected in the aggregate route-level analysis below).

Figure 7.2 Route-level market share, averaged across routes

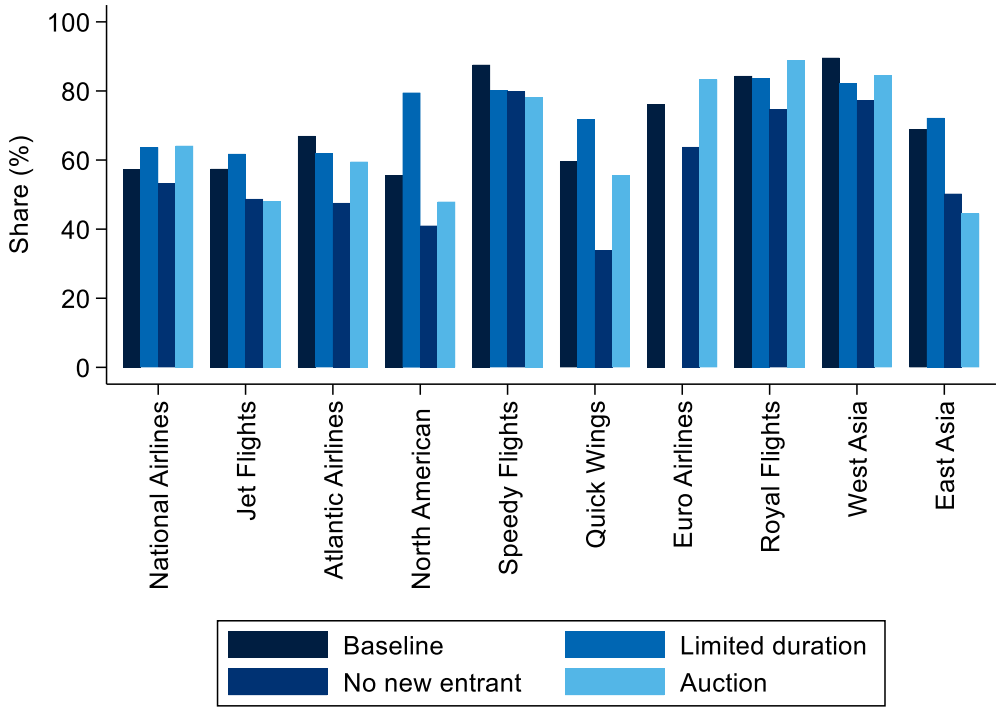

Note: Only those routes on which an airline operated at least one slot are included in the calculation of the average market share. It can therefore be interpreted as the 'average market share on routes that the airline operated'.

Source: Oxera.

Another way to assess the extent of competition in a market is using the Herfindahl–Hirschman Index (HHI). An advantage of the HHI over analysing market shares is that it captures the distribution of market shares across firms. Box 7.1 below describes this measure.

[90] The market share is calculated at the end of a round (i.e. after all bidding and trading stages in a round are completed).

Box 7.1 Herfindahl–Hirschmann Index (HHI)

The HHI is a commonly used indicator of concentration among firms in a market. It accounts for both the number of firms and their respective sizes. The HHI is calculated by summing the squared market shares s_i (in percent) of each firm *i* for the total number of firms *N* active on the market:

$$HHI = \sum_{i=1}^{N} s_i^2$$

The HHI can range from close to zero (in a case of perfect competition) to 10,000 (in the case of a monopoly—i.e. one firm producing 100% of output on the market). In order to assess a merger's impact on the overall market structure, the change in the HHI (also called 'delta') is often calculated by subtracting the pre-merger index from the expected post-merger HHI. The 'delta' value gives an indication of the change in market structure following a merger and is used by the CMA in its merger control decisions. In its merger investigations, the CMA considers a post-merger HHI above 1,000 as concentrated and a post-merger HHI above 2,000 as highly concentrated.

Source: Competition and Markets Authority, Merger Assessment Guidelines, paras 5.3.4–5.3.5.

The average HHI across routes and treatments is shown in Figure 7.3. Overall, we find a lower HHI—i.e. more competition—in treatments 3 and 4 (no new entrant rule, auction) than in the other two treatments.[91]

It is also worth noting that, while the HHI levels are high in absolute terms (i.e. above 6,000), this is likely to be because existing slots are not included. Therefore the focus should be on the impact of each treatment on the HHI, rather than the absolute level.

[91] Pairwise comparison between the no new entrant rule and the baseline treatment shows statistically significant differences at the 5% level using the Mann–Whitney test. Pairwise comparison between the auction and the baseline treatment shows statistically significant differences at the 5.8% level using the Mann–Whitney test. The Mann–Whitney test is a non-parametric test commonly used in experiments. It does not rely on any assumption on the distribution of the data of interest. It is commonly used in analysis of experimental data where the assumptions of parametric tests, such as the t-test, are likely to be invalidated. As a result, while we present confidence intervals as a visual illustration of the degree of variability around the results, our assessment of statistical significance will ultimately be based on the Mann–Whitney test. Therefore, unless stated otherwise, all statistical tests in this analysis will be Mann–Whitney tests. For a full exposition, see Moffatt, P.G. (2015), *Experimetrics: Econometrics for experimental economics*, Macmillan International Higher Education, section 3.5.

Figure 7.3 Average HHI at the route level

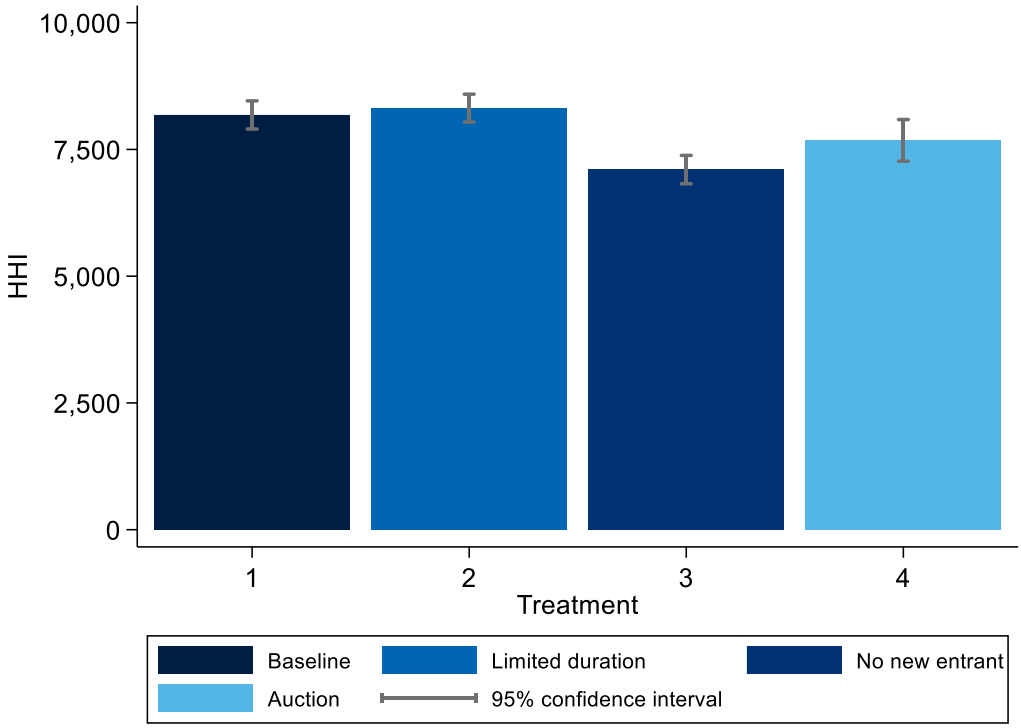

Note: A confidence interval is a way to measure the precision of an estimate. It is a range of values that contain the true statistical value (i.e. across all samples) with a probability of 95%.

Source: Oxera.

Looking at a more disaggregate level, Figure 7.4 shows the average number of airlines flying a given route with the new slots. Across all treatments there is an average of 1.0 airline per route, with the average ranging from 0 (Bern, in treatment 4) to 2.9 (New York, in treatment 4). This corresponds to reality, where there are on average 1.55 airlines per route at Heathrow (see section 5.3).

Figure 7.4 Average number of airlines flying each route

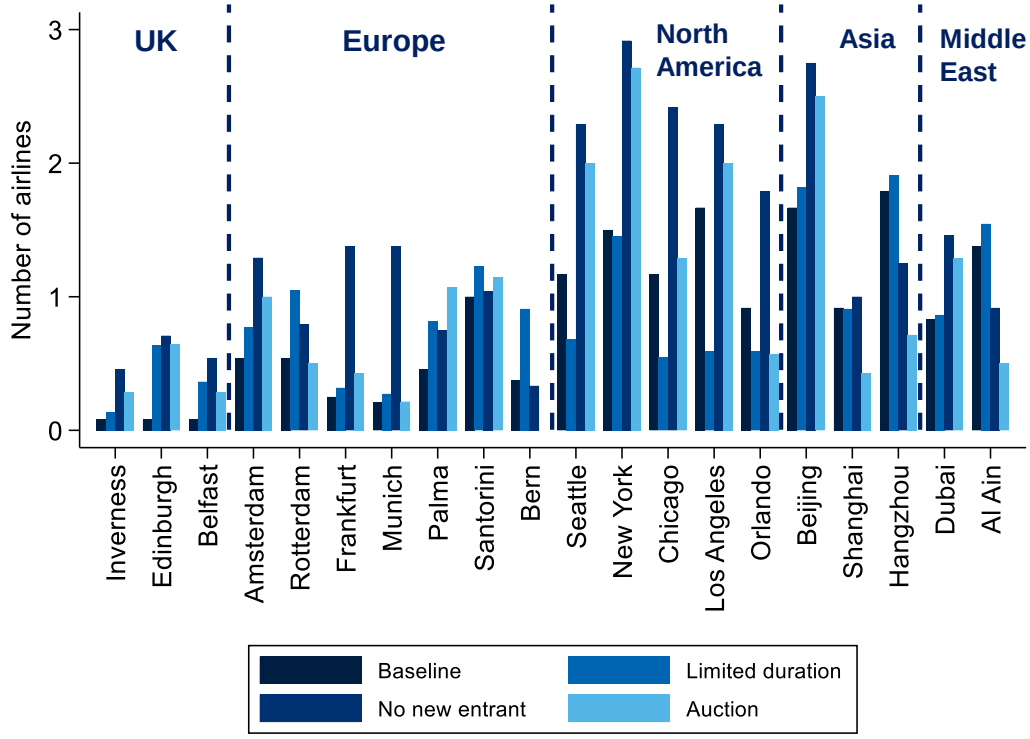

Source: Oxera.

We also observe that the number of airlines on North American routes increased in treatments 3 and 4 (no new entrant rule, auction). This is unsurprising given the larger share of slots won by North American airlines in these treatments (see Figure 7.8).

For routes to the Middle East and Asia, the number of airlines per route in treatments 1 and 2 (baseline, limited duration) was greater to some destinations such as Hangzhou, and smaller to other destinations such as Beijing. We also find that in treatment 1 (baseline), across most sessions, no airline flew to the UK.

More generally, we find that the average number of airlines per route is greater in treatment 3 (no new entrant) than in the other treatments. This can be explained by the fact that there are fewer airlines that do not fly at all. Slot utilisation is also highest in treatment 3 (as explained in section 7.4 below).

To understand these results better, we grouped routes by region and calculated the average route-level HHI. These results are shown in Figure 7.5 and are aligned with those above. We find that competition on North American routes is greater in treatment 3 (no new entrant rule).[92] Furthermore, we observe that in treatment 1 (baseline) competition is weaker on European routes than in treatment 3.[93]

In other words, we find that, on average, removing the new entrant rule redirected traffic towards more profitable routes in North America and Asia.

[92] Results statistically significant at the 5% level using Mann–Whitney tests comparing the no new entrant rule treatment with the baseline and limited duration.
[93] The difference between treatment 1 and treatment 3 is significant at the 5% level using Mann–Whitney tests.

The change in the HHI (by treatment) for the domestic region should be treated with caution, given the very low number of airlines that chose to fly domestic routes.[94]

Figure 7.5 Average route-level HHI at a region level

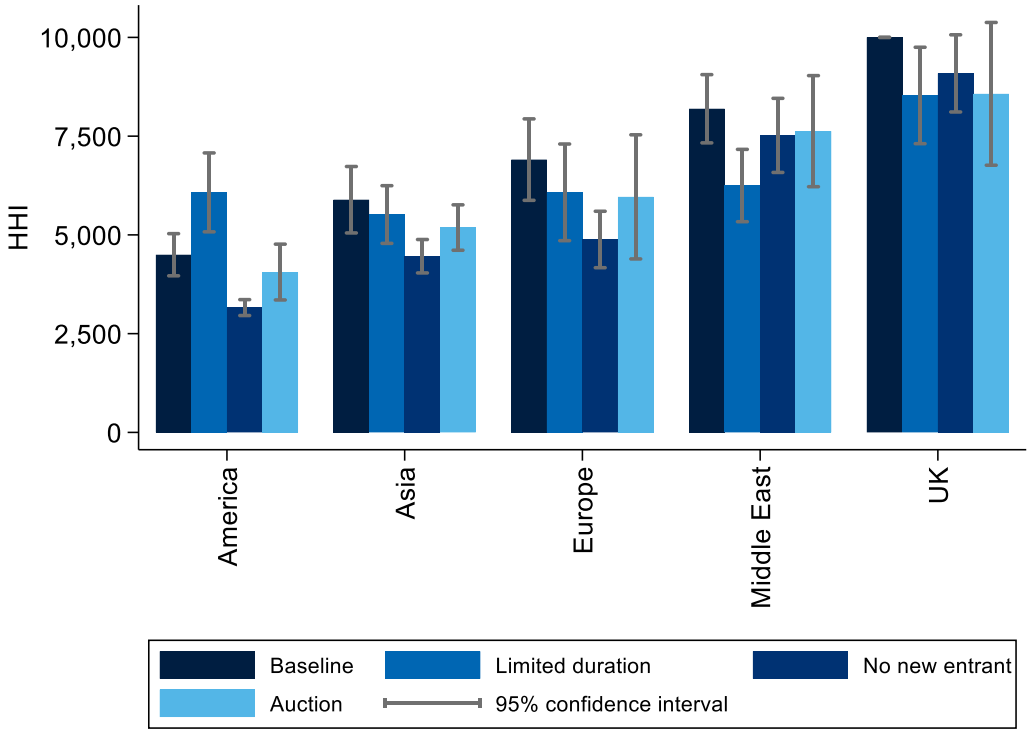

Note: (1) A confidence interval is a way to measure the precision of an estimate. It is a range of values that contain the true statistical value (i.e. across all samples) with a probability of 95%. (2) There is no confidence interval on the average for the UK in the baseline treatment because there is not enough variation for this statistic to be computed.

Source: Oxera.

We also assessed whether the different slot allocation mechanisms affected the level of competition on new routes specifically. As shown in Figure 7.6, route-level HHI is higher in treatments 3 and 4 than in the baseline.[95] This indicates that there is a trade-off between more route-level competition overall and more route-level competition on new routes specifically. (As indicated in Figure 7.13 below, the reduced competition on new routes is likely to be due to the reduced number of slots operated on these routes in treatments 3 and 4.)

[94] In treatment 1, only three participants flew on domestic routes. The three were in separate sessions. This means that, in treatment 1 on domestic routes, we observe only monopolies, explaining the HHI taking the value 10,000. The difference in HHI between treatment 1 and the other treatments is not statistically significant.
[95] The difference between treatments 1 and 3 is significant at the 10% level. The difference between treatments 1 and 4 is significant at the 5% level.

Figure 7.6 Average HHI on new routes

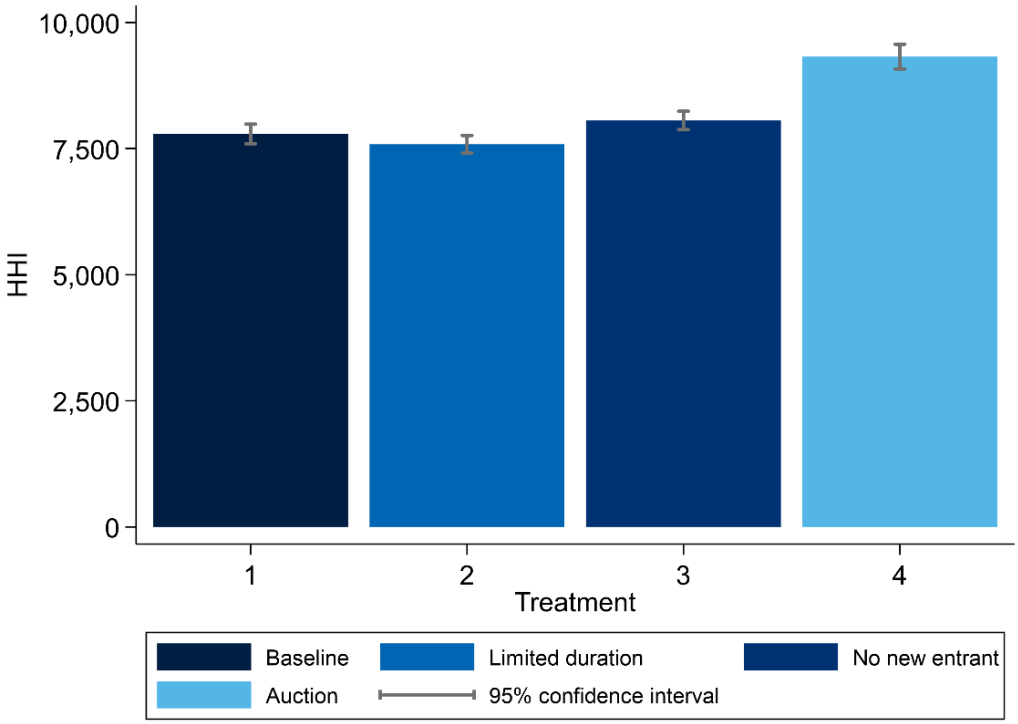

Note: A confidence interval is a way to measure the precision of an estimate. It is a range of values that contain the true statistical value (i.e. across all samples) with a probability of 95%

Source: Oxera.

7.2.2 Competition at the airport level

In this section, we investigate the effect of the different allocation mechanisms on competition for the 48 slots in the experiment.

The HHI for the acquisition of slots at the airport level for each treatment is shown in Figure 7.7. The HHI in the baseline treatment is 2,305. In the other treatments, the average HHI is 1,788 (no new entrant), 3,033 (limited duration) and 2,532 (auction). Only the difference between the baseline treatment and the no new entrant treatment is statistically significant.[96] The largest market share of an airline is also smaller in the no new entrant treatment relative to the baseline.

[96] Statistically significant at the 5% level using Mann–Whitney tests.

Figure 7.7 Average HHI at the airport level

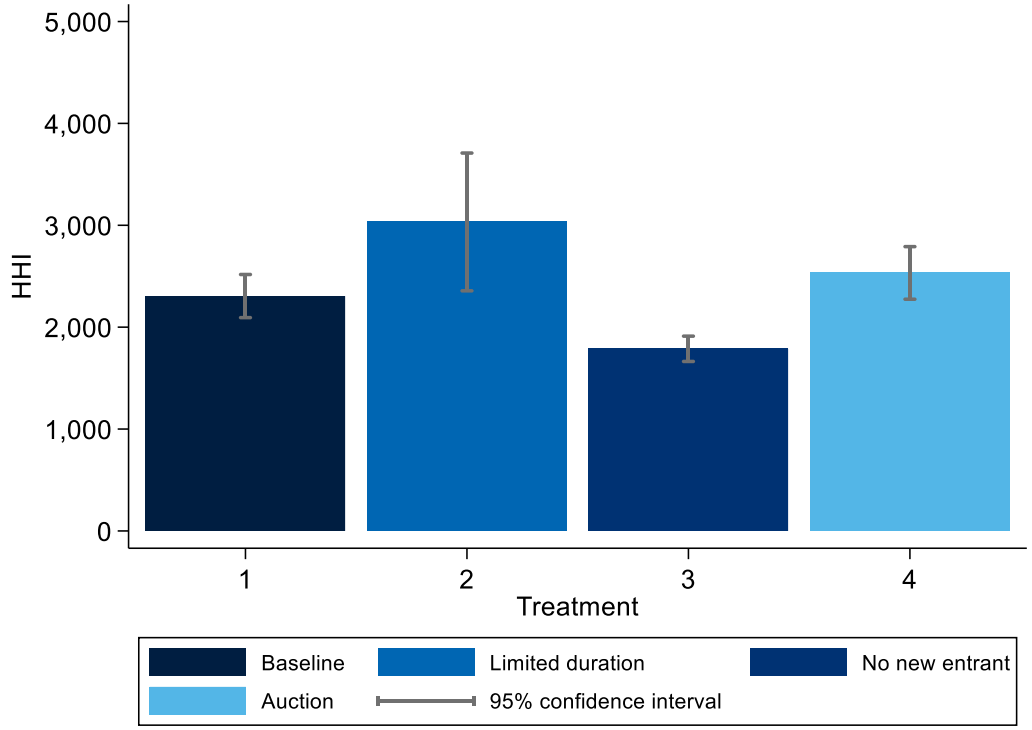

Note: A confidence interval is a way to measure the precision of an estimate. It is a range of values that contain the true statistical value (i.e. across all samples) with a probability of 95%.

Source: Oxera.

In order to disentangle the different effects in these treatments, we consider which airlines were favoured in each treatment. In Figure 7.8 we show the market share by airline type. Four results stand out.

- The North Atlantic airlines had higher market shares in treatments 3 and 4 (no new entrant rule, auction) and lower market shares in treatment 2 (limited duration) than in the baseline.[97] These airlines have a large budget to spend on routes where they can make substantial profits, and therefore they have a strong incentive to bid and trade for slots. In treatment 3 there is no new entrant rule, which means that there are more slots that non-new entrants, such as the North Atlantic airlines, can bid for. They also have large budgets, which means that they are able to acquire a significant number of slots in the auction.
- National Airlines does much better in the auction than in the baseline treatment, for similar reasons to those set out above for the North Atlantic airlines.
- In treatments 3 and 4, East Asia Airlines performed worse than in the baseline. In treatment 3, this can be explained by the fact that East Asia Airlines loses its new entrant status.[98]

[97] All results are statistically significant at the 5% level when comparing each of these three treatments to the baseline.
[98] Result statistically significant at the 5% level comparing each of these two treatments to the baseline. The difference between the baseline and limited duration treatments is not statistically significant.

- In the auction treatment, LCCs had a lower market share than in all other treatments.[99] This is likely to be because of their relatively low budgets.

Figure 7.8 Average market share by airline type at the airport level

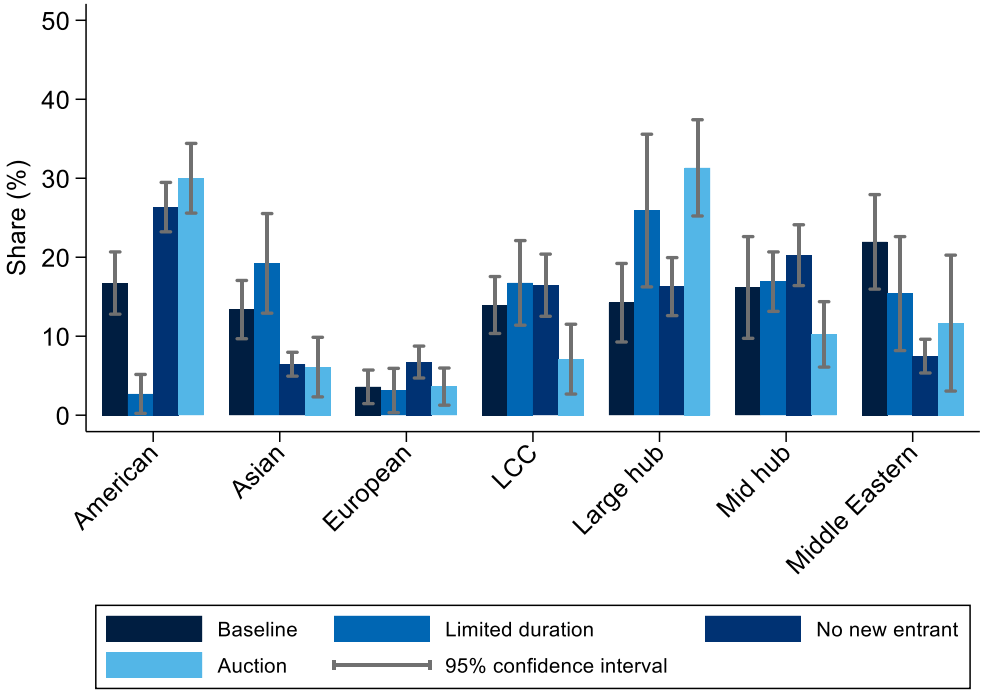

Note: (1) A confidence interval is a way to measure the precision of an estimate. It is a range of values that contain the true statistical value (i.e. across all samples) with a probability of 95%. (2) See Table 5.1, where each of the airline types is described.

Source: Oxera.

We separately look at the effect of the different slot allocation mechanisms on the ability of new entrants to acquire slots. In treatment 2 (limited duration), there is a second slot coordinator allocation instead of the second trading stage, which is present in the baseline, no new entrant and auction treatments. Therefore, new entrants are unable to participate in the trade stage and do not have an opportunity to sell the slots that they win in the bid stage. As a result, we would expect the market shares for new entrants to be higher in treatment 2. In treatment 3, as there are no slots reserved for new entrants, we would expect the share of new entrants to be lower. These results are shown in Figure 7.9. The market share of new entrants increased by 8% in treatment 2, while the market share of new entrants decreased by 4% in treatment 3.

In treatment 4 the share of new entrants also declined, by 13%.[100] While there are still slots reserved for new entrants in the auction, winning general (i.e. non-new entrant) slots would require an airline to out-bid other airlines. As two of the three new entrants have relatively low budgets, it may be difficult for them to win slots against airlines with larger budgets.

[99] Both results are statistically significant at the 5% level comparing to the baseline.
[100] The comparison between the limited duration treatment and the baseline is significant at the 10% level. The comparison between the auction treatment and the baseline is significant at the 5% level. The comparison between the no new entrant treatment and the baseline is not statistically significant.

Figure 7.9 Average new entrant market shares at the airport level

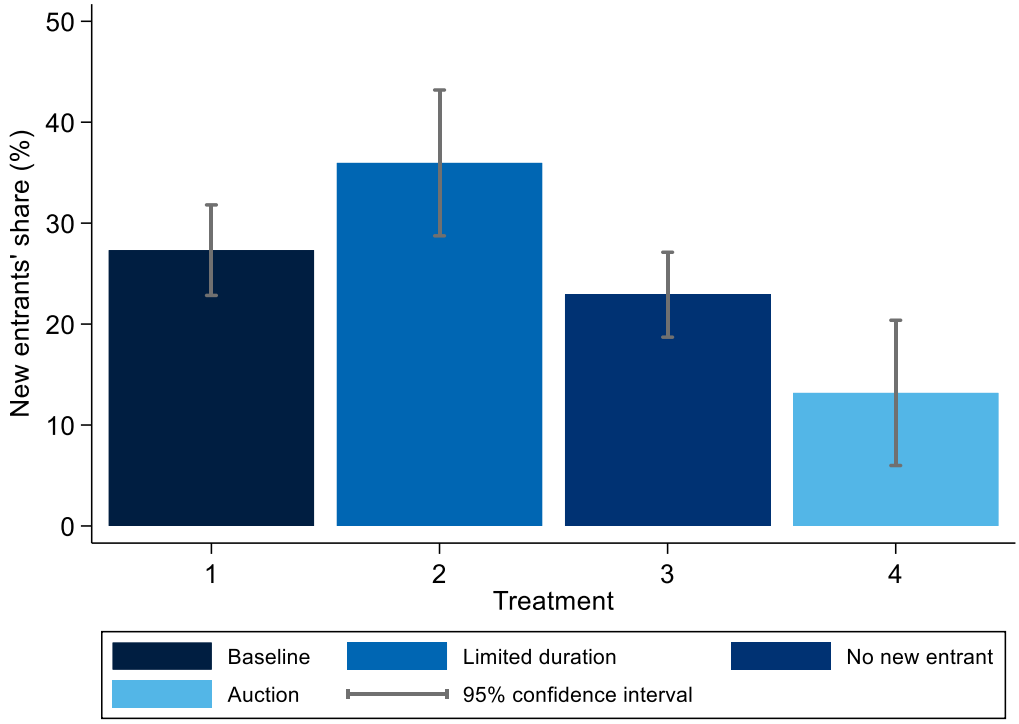

Note: A confidence interval is a way to measure the precision of an estimate. It is a range of values that contain the true statistical value (i.e. across all samples) with a probability of 95%.

Source: Oxera.

7.2.3 Summary of findings on competition

We find that treatment 3 (no new entrant rule) and treatment 4 (auction) increase competition, when looking at competition on routes. We also find that treatment 3 (no new entrant rule) increases competition at the airport level. In treatment 2 (limited duration), the level of competition is similar to the baseline. There is a trade-off between more route-level competition overall and more route-level competition on new routes specifically.

Looking at new entrants, their market share in treatment 4 (auction), while they had the highest market share in treatment 2 (limited duration). This is in part driving the differences in competition between treatments, as in treatment 4 there is more competition on the more profitable routes, which are flown mostly by airlines that are not new entrants.

7.3 Connectivity

Connectivity is the ability for passengers to use Heathrow to travel to a wide range of destinations. Figure 7.10 shows the number of flights between Heathrow and the routes included in the experiment.

Overall, as previously noted, we find that domestic connectivity—i.e. the number of flights between Heathrow and other UK airports—is low. On average, across all treatments in the experiment, there are less than 0.5 slots used for destinations in the UK. The number of connections to other European destinations is also low, with an average of 1.5 slots.

Figure 7.10 Average number of flights by destination

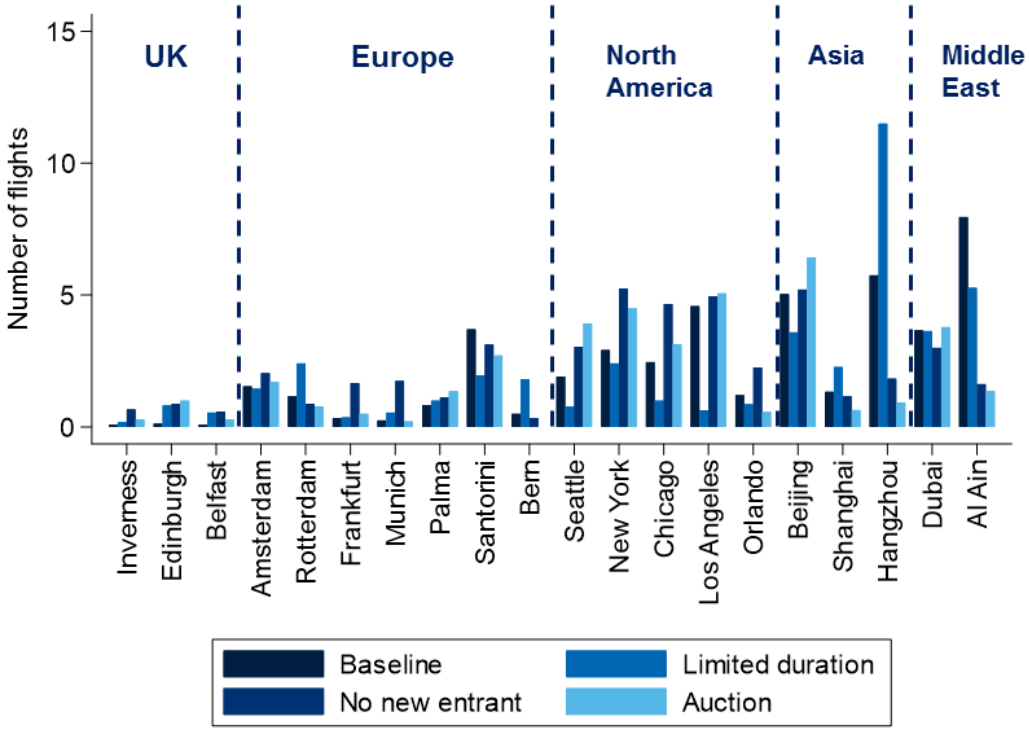

Note: A confidence interval is a way to measure the precision of an estimate. It is a range of values that contain the true statistical value (i.e. across all samples) with a probability of 95%.

Source: Oxera.

Figure 7.11 illustrates the share of long-haul flights across treatments. Overall, we find some variation across treatments, with the average ranging from 72% (no new entrant rule) to 82% (baseline). The difference between the two is statistically significant.[101] This may be a reflection of the fact that, while the North Atlantic airlines have a large share of slots in treatment 3, the share of slots of Asian and Middle Eastern airlines declined. When comparing any other two treatment pairs, the difference is not statistically significant.

[101] At the 5% level using a Mann–Whitney test. The difference between the other treatments are not statistically significant.

Figure 7.11 Average share of long-haul flights

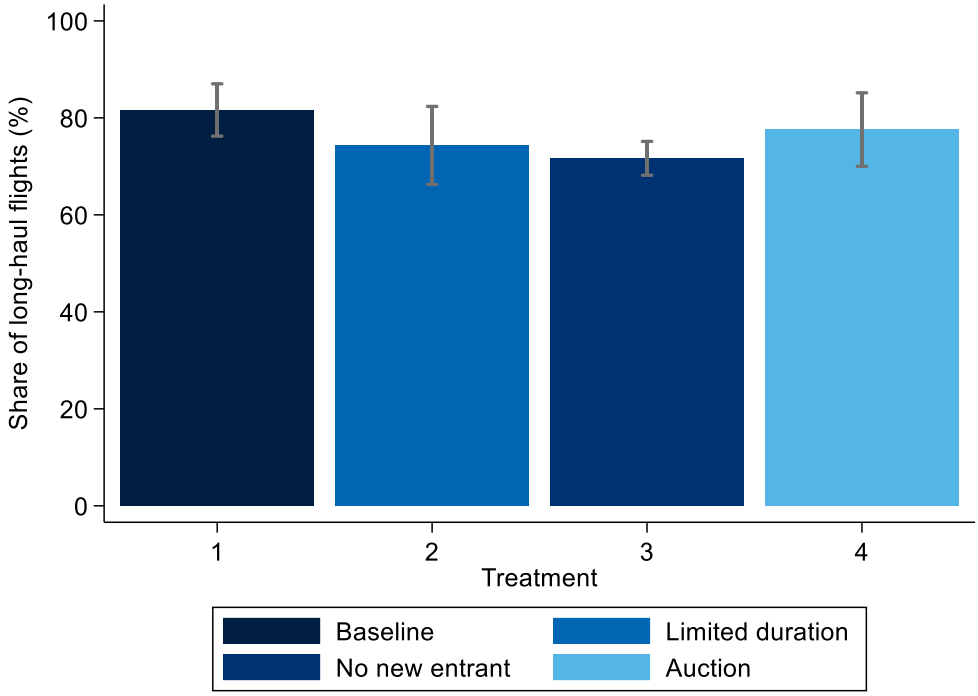

Note: A confidence interval is a way to measure the precision of an estimate. It is a range of values that contain the true statistical value (i.e. across all samples) with a probability of 95%.

Source: Oxera.

Similarly, Figure 7.12 shows the share of international flights. We find that most flights are international, with the baseline treatment having the highest rate (99.4%) and the other treatments a share close to 95%. The difference between the baseline treatment and the other treatments is statistically significant.[102] This suggests that domestic connectivity was higher in treatments 2, 3 and 4 relative to the baseline. However, as discussed above, the results regarding domestic connectivity are based on a small sample and should therefore be treated with caution.

[102] At the 5% level using Mann–Whitney tests for pairwise comparisons.

Figure 7.12 Average share of international flights

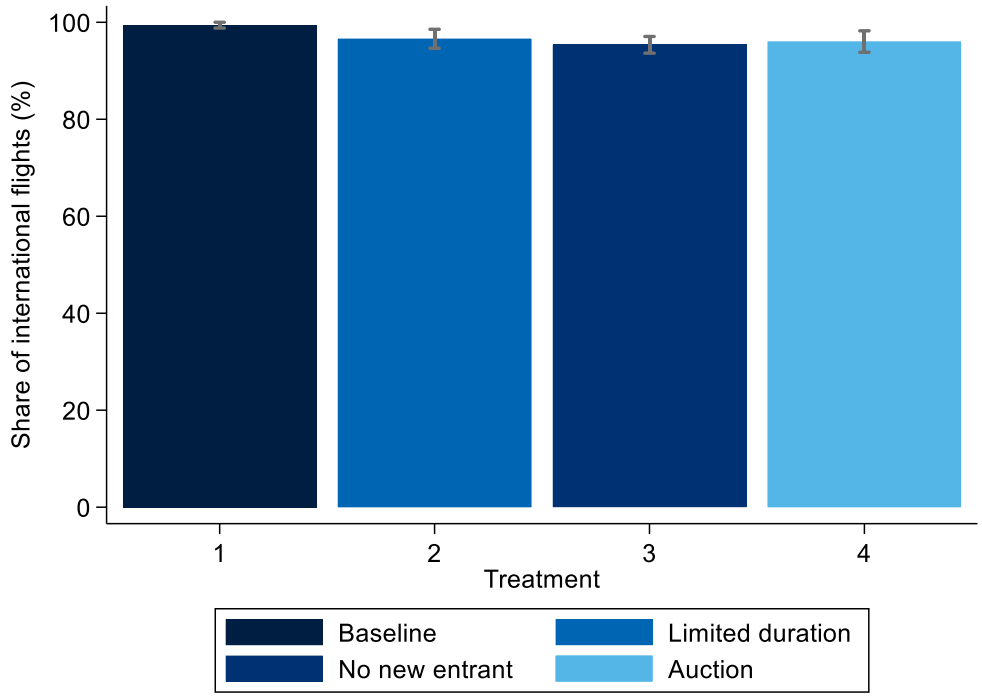

Note: A confidence interval is a way to measure the precision of an estimate. It is a range of values that contain the true statistical value (i.e. across all samples) with a probability of 95%.

Source: Oxera.

Lastly, we assessed the effect of the different allocation mechanisms on the extent of new routes flown. These results are shown in Figure 7.13. We find that the average number of flights on new routes is lower in treatments 3 (no new entrant) and 4 (auction). Regarding treatment 3, this may be because fewer slots held by new entrant airlines resulted in fewer flights on new routes.

While new routes are prioritised by the slot coordinator in treatments 1-3, in treatment 4 airlines do not state their route choice in advance of the slot allocation. Therefore, this may explain why fewer new routes are flown in the auction treatment. See section 7.5 for a discussion of how much airlines altered their route choice.

Figure 7.13 Average number of flights on new routes

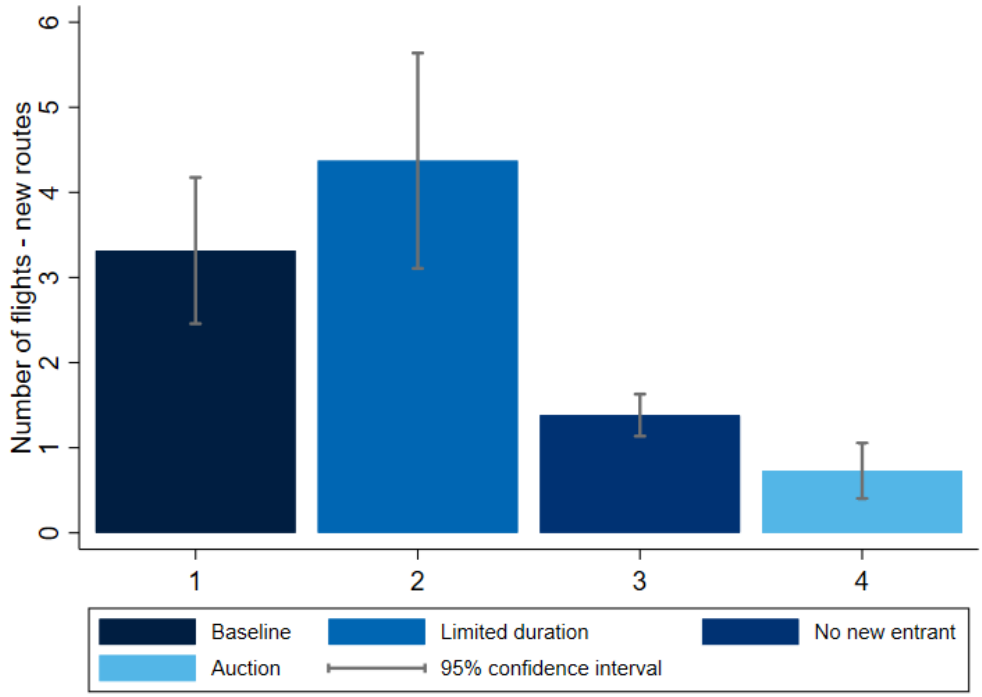

Note: A confidence interval is a way to measure the precision of an estimate. It is a range of values that contain the true statistical value (i.e. across all samples) with a probability of 95%.

Source: Oxera.

7.3.1 Expected impact on connectivity, considering current slot holdings

The experiment results above focus on the slots released with the new runway. However the release of new slots occurs in a context where airlines are already operating at Heathrow. It is therefore useful to combine the results above with information on airlines' existing slot holdings at Heathrow to understand the complete picture of the impact of the increased capacity.

Ideally, we would combine actual data on existing flights with the experiment data. However, the experiment data cannot be linked to actual data at the level of granularity of airlines and destinations. In particular:

- airlines in the experiment cannot be linked to a specific airline in reality;
- destinations in the experiment do not capture the full range of destinations currently offered from Heathrow.

As an alternative, we explore the number of additional slots used to operate flights to a particular region. This also provides some insight into the impact on competition when considering existing slot holdings. For example, if we find that the third runway does not lead to many additional flights to Europe, then it is unlikely that competitive outcomes on European routes with the new runway will be significantly different from the status quo. On the other hand, if there are a number of additional flights to a particular region, then there may be more competition with the new runway (though this will depend on the ultimate destinations selected).

Therefore, we conducted the analysis as follows.[103]

- First, we calculated the current average daily slot pairs to each region, using OAG data. For example, there are 67 daily slot pairs used to operate routes from Heathrow to Asia and 83 daily slot pairs used to operate routes from all London airports to Asia.[104]
- Second, we calculated the number of additional slot pairs to each region from the third runway using the experiment data. Since our experiment contains 48 slots pairs, but the number of slot pairs released by a third runway is around 335, we scale up the 48 slots in the experiment to 335 slots.[105] For example in treatment 1, on average 25% (12/48) of new slots are allocated to Asian routes, so we consider that 25% (84/335) of new slots would be operated on Asian routes.
- Third, we calculated the percentage increase in flights to each region by combining the two steps above. Using the example above, there are 67 existing slots at Heathrow and 83 new ones to Asia, so the percentage increase in slots is 125%.

Figure 7.14 shows the percentage increase in flights to each region when combining the data on existing slots at Heathrow with the experiment data. When considering long-haul flights to North America, Asia and the Middle East, the percentage increase in flights is 90%, 135% and 130% respectively. Conversely, the percentage increase in short-haul flights is much smaller—an average of 21% for European flights and 37% for domestic flights.

[103] Turkey, Iran, Iraq, Israel, Jordan, Kuwait, Lebanon, Oman, Qatar, Saudi Arabia and the United Arab Emirates are defined as Middle Eastern countries. North American countries are Canada, the USA and Mexico. Europe includes Russia (following the NAPDM definition).
[104] London airports include Heathrow, Gatwick, Stansted, Southend, Luton and City.
[105] Analysis of OAG data shows that there are approximately 670 slot pairs. A 50% increase in capacity due to the third runway leads to 335 additional slot pairs.

Figure 7.14 Percentage increase in flights to each region, for Heathrow

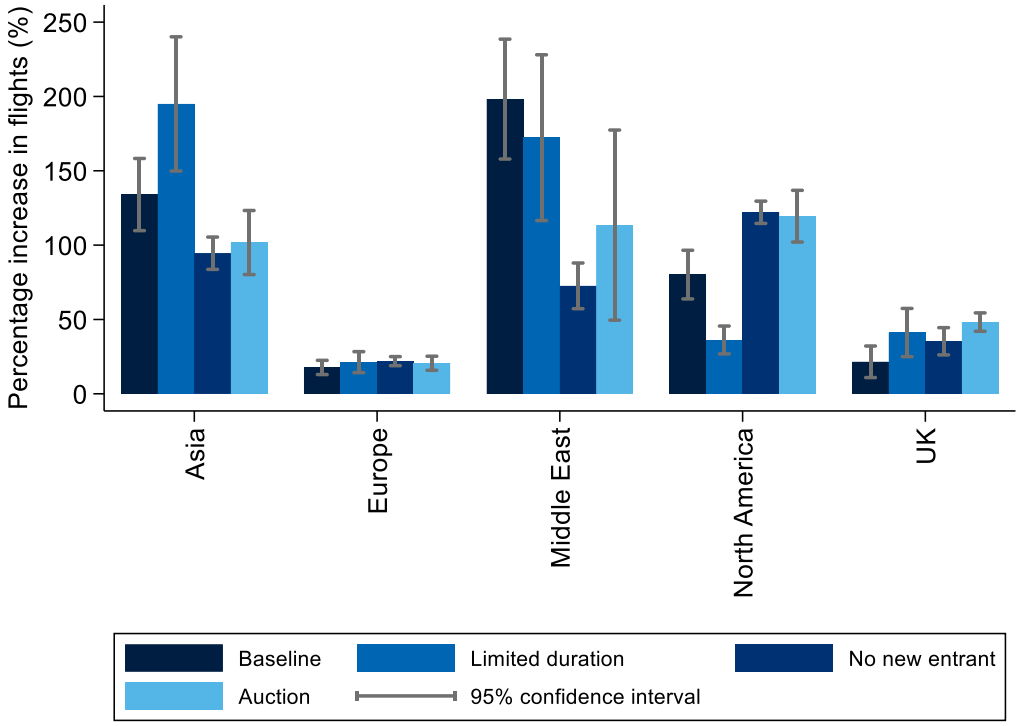

Source: Oxera analysis and OAG data.

In addition to considering the existing capacity at Heathrow, it is relevant to consider the existing capacity at other London airports. Figure 7.15 shows the percentage increase in flights to each region from London as a result of the new runway. As there are many existing domestic and European flights at other London airports, the effect of the additional domestic and European flights from the third runway is small. Indeed, the percentage increase in European and domestic flights is just 7% and 13%, respectively.

Conversely, as there are few existing flights to the Middle East, North America and Asia from other London airports, the effect of the third runway on routes to these regions is significant, with an increase in the number of flights to these regions of 80%, 74% and 105%, respectively. This result is therefore consistent with the result when looking only at Heathrow, with a greater increase in flights on long-haul routes.

Figure 7.15 Percentage increase in flights to each region, for all London airports

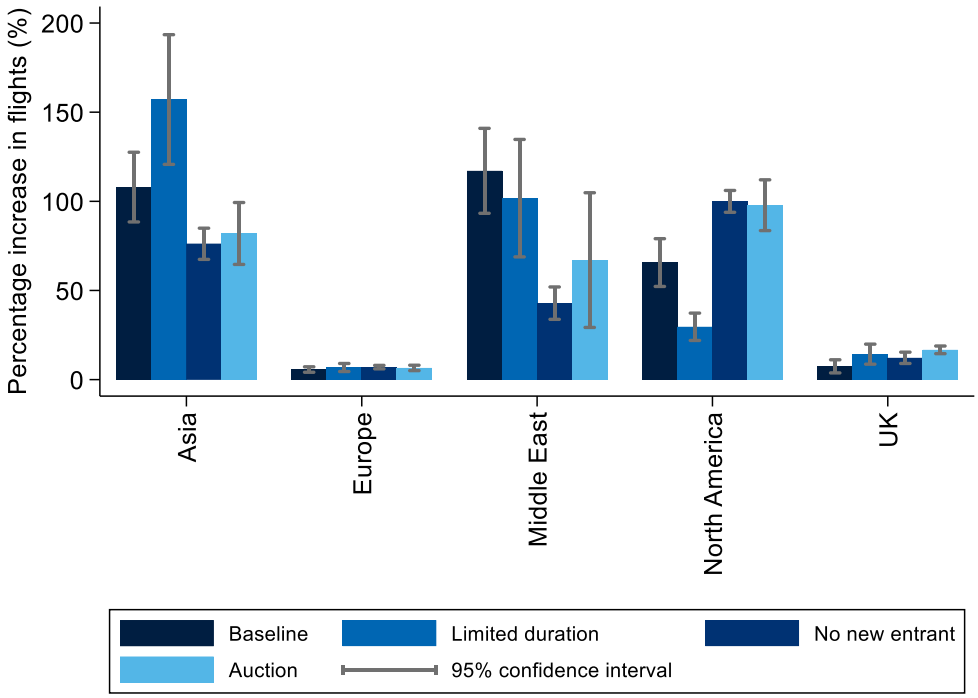

Source: Oxera analysis and OAG data.

7.4 Efficient use of Heathrow capacity

Efficiency is a concept used to assess the use of capacity at an airport. It is often split into two categories.

The first is productive efficiency, which concerns whether the airport infrastructure is fully utilised. We consider productive efficiency by looking at whether there are any unallocated slots and the number of seats per slot.

The second is allocative efficiency, which analyses whether capacity is allocated in a way that maximises social welfare, which is the sum of producer surplus (the benefit to firms) and consumer surplus (the benefit to consumers).

7.4.2 Productive efficiency

Productive efficiency relates to the number of slots flown, with maximum productive efficiency arising when all the slots are flown. Otherwise, it would be possible to obtain a more efficient outcome by allocating the unused slots. We plot the number of slots used across the treatments in Figure 7.16. A line has been added to show the maximum number of slots of 48.

Figure 7.16 Average number of slots used

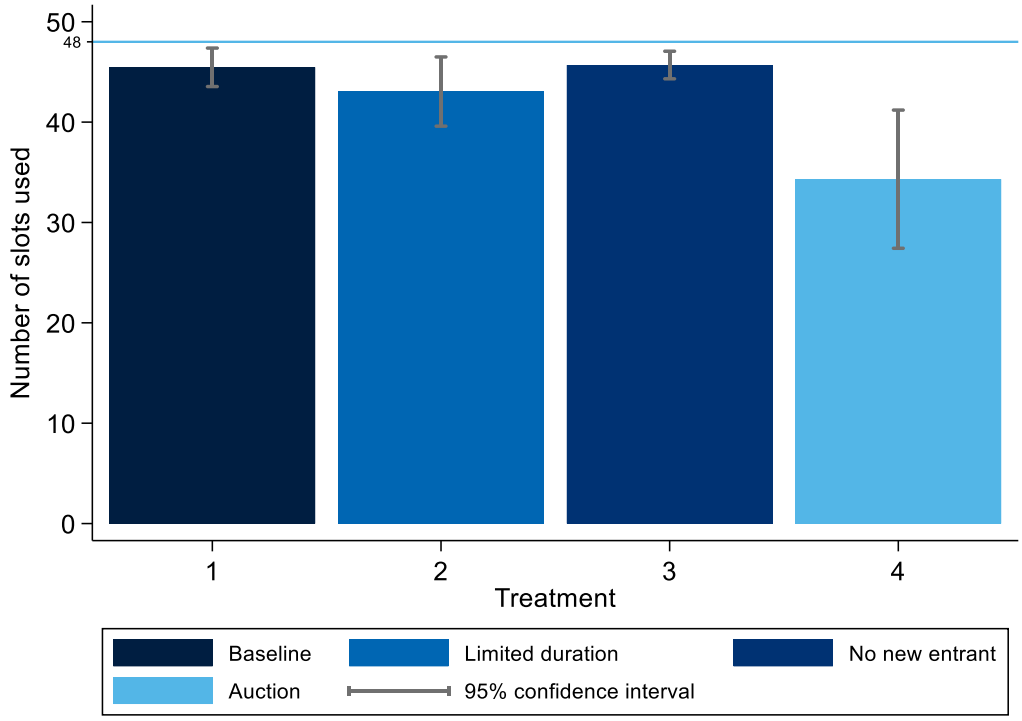

Note: A confidence interval is a way to measure the precision of an estimate. It is a range of values that contain the true statistical value (i.e. across all samples) with a probability of 95%.

Source: Oxera.

The baseline, limited duration and no new entrant treatments show slot usage close to maximum capacity. However, there is under-utilisation of slots in the auction treatment. This is likely to be due, in part, to the particular auction design implemented in the experiment, where the combinatorial auction was a one-shot sealed bid. In the experiment, airlines had one opportunity to make their bids, and the airlines did not make sufficient 'small' package bids (which would have been optimal for the bidder and would have resulted in the remainder of the slots being awarded). In contrast, if an auction were implemented in practice then it may have mechanisms to encourage bidders to make their full range of optimal bids. (See Appendix A3 for further detail.)

Next, we estimate the number of seats and passengers per slot. While this is an intuitive measure of how well a runway is being utilised, it does not account for passenger and airline preferences. For example, a short-haul flight during a peak time for business travellers can still bring significant benefits even if the flight is carrying fewer passengers than a long-haul one. In other words, using seats as a measure of efficiency assumes that all seats are homogenous.

To estimate the number of seats per slot, we consider whether the slot was used for a long-haul or a short-haul flight.[106] This is then used to estimate the average number of seats per flight in the experiment. This is shown in Figure 7.17.

[106] We obtained estimates of the average number of seats for long- and short-haul flights from Heathrow in the summer of 2018 using OAG data. We classified any flight longer than 360 minutes as a long-haul flight, and calculated the number of seats for an average long-haul flight. This was repeated for short-haul flights. A short-haul flight has 161 seats on average, while a long-haul flight has 283 seats.

Figure 7.17 Average number of seats per flight

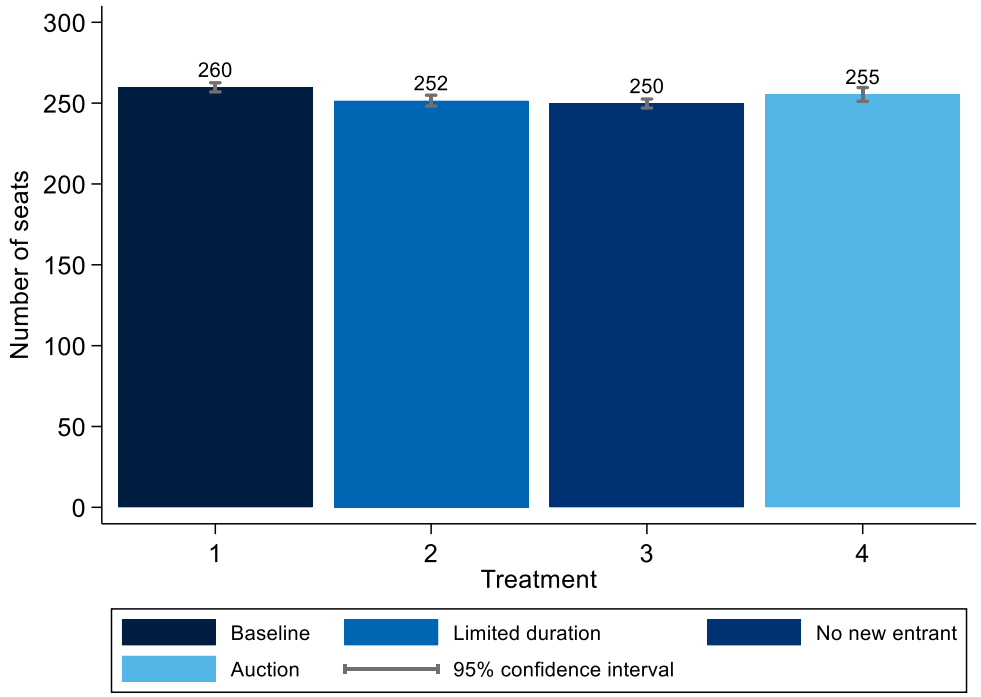

Note: A confidence interval is a way to measure the precision of an estimate. It is a range of values that contain the true statistical value (i.e. across all samples) with a probability of 95%.

Source: Oxera analysis.

The average number of seats per slot is similar between treatments.[107] This is likely to be because the treatments also have a similar proportion of long-haul flights, as shown in Figure 7.18.

[107] However, some results are statistically significant. The difference between the no new entrant and limited duration treatments relative to the baseline are statistically significant at the 5% level using the Mann–Whitney test.

Figure 7.18 Long-haul flights as a proportion of all flights

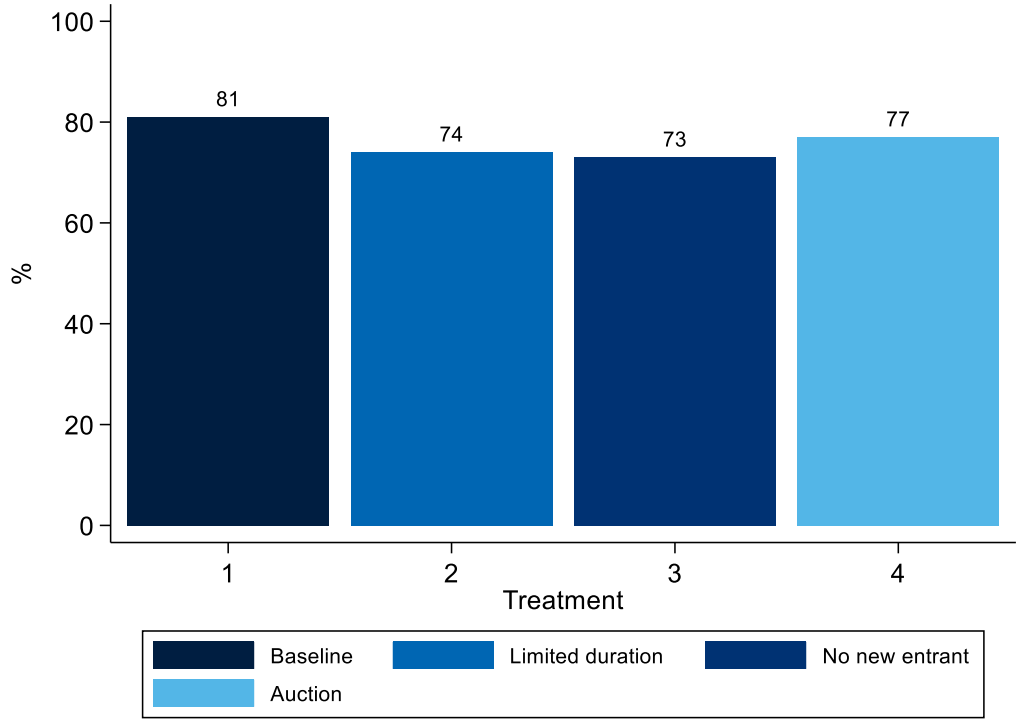

Source: Oxera.

We also estimated the seat-km per slot for each treatment. We obtained the distance between Heathrow and each of the destination airports in the experiment.[108] We then multiplied the distance by the number of seats estimated above to obtain the seat-km per slot. The average by treatment is set out in Figure 7.19.

Relative to the baseline, all treatments have a lower seat-km per flight. However, the difference between the limited duration treatment and the baseline is not statistically significant.[109] The no new entrant and auction treatments have a lower seat-km because the route mix of these flights is more heavily weighted towards North American destinations than towards Asian and Middle Eastern destinations, which are farther away.

[108] Great-circle distance, which is the shortest distance between two points on the surface of the Earth. Where there are multiple airports in a city, we used the largest airport.
[109] Using the Mann–Whitney test at a 5% significance level. The other two treatments show a statistically significant treatment effect relative to the baseline.

Figure 7.19 Average seat-km per flight

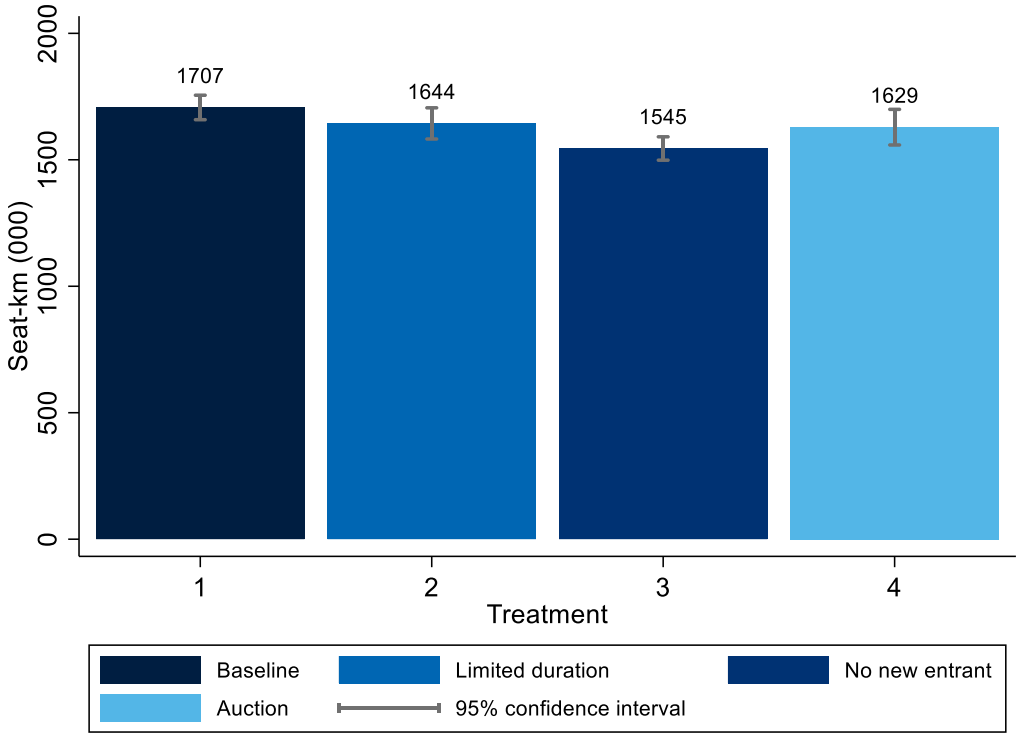

Note: A confidence interval is a way to measure the precision of an estimate. It is a range of values that contain the true statistical value (i.e. across all samples) with a probability of 95%.

Source: Oxera.

7.4.3 Allocative efficiency—producer surplus

Allocative efficiency is made up of benefits to both passengers (known as consumer surplus) and airlines (producer surplus). This section considers producer surplus, while consumer surplus is considered in the following section.

To obtain a measure of producer surplus, we calculated the payoff per slot for each round.[110] This is shown in Figure 7.20.

[110] These are payoffs arising only from operating flights and not payoffs arising from any remaining budget or trading profits.

Figure 7.20 Average payoff per slot

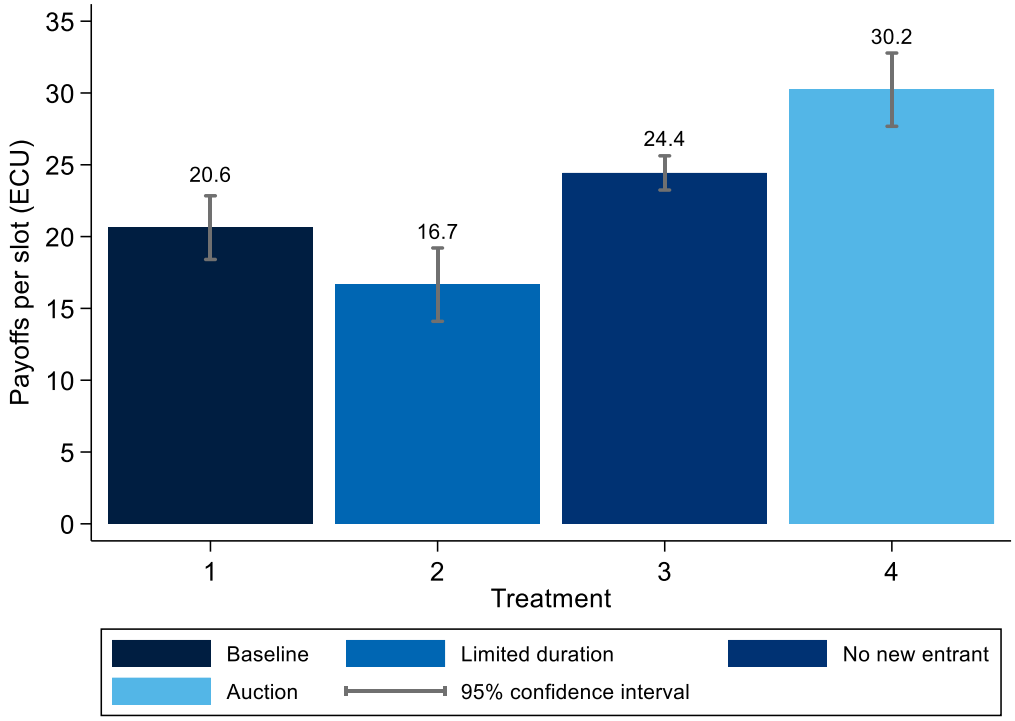

Note: These are the sum of payoffs in stages 1 and 2 divided by the total slots used in stages 1 and 2. A confidence interval is a way to measure the precision of an estimate. It is a range of values that contain the true statistical value (i.e. across all samples) with a probability of 95%.

Source: Oxera.

The limited duration treatment has lower payoffs per slot than the baseline. The no new entrant treatment has higher payoffs per slot than the baseline, while the auction has the highest payoffs per slot of all treatments.[111]

A key conclusion from the above results is that those treatments that include more opportunities for trading are more efficient. If the initial allocation is inefficient, trading allows airlines that value slots highly to obtain them from those that value them less. For example, the limited duration treatment replaces the second stage of trading with another round of slot coordinator allocation, and there are therefore fewer opportunities for trading, leading to a less efficient outcome. The no new entrant treatment also encourages trading, because without new entrant slots all slots can be traded immediately.

Another potential driver of efficiency in treatment 3 is that some slots are no longer being reserved for new entrants. Slots that would otherwise have been new entrant slots can therefore now be allocated to those airlines that value the slot more. Given that new entrants (with the exception of East Asia Airways) tend to have lower slot valuations than other airlines, this effect may be significant.

We note that the auction contains the same opportunities for trading as the baseline, but leads to a more efficient outcome. Indeed, the initial allocation in the baseline treatment is less efficient than the initial allocation in the auction, and trading is insufficient to overcome this gap. This is set out in Figure 7.21 below, which shows the payoffs at the end of the first rather than the second

[111] Relative to the baseline, the other treatments show statistically significant treatment effects at the 5% significance level using the Mann–Whitney test. Furthermore, the auction treatment also has a statistically significant treatment effect relative to the limited duration and no new entrant treatments.

stage.[112] Therefore, even though trading improves efficiency, it is unlikely to be sufficient to fully mitigate inefficiencies that can arise in the current allocation mechanism.

Figure 7.21 Payoff per slot at the end of the first stage

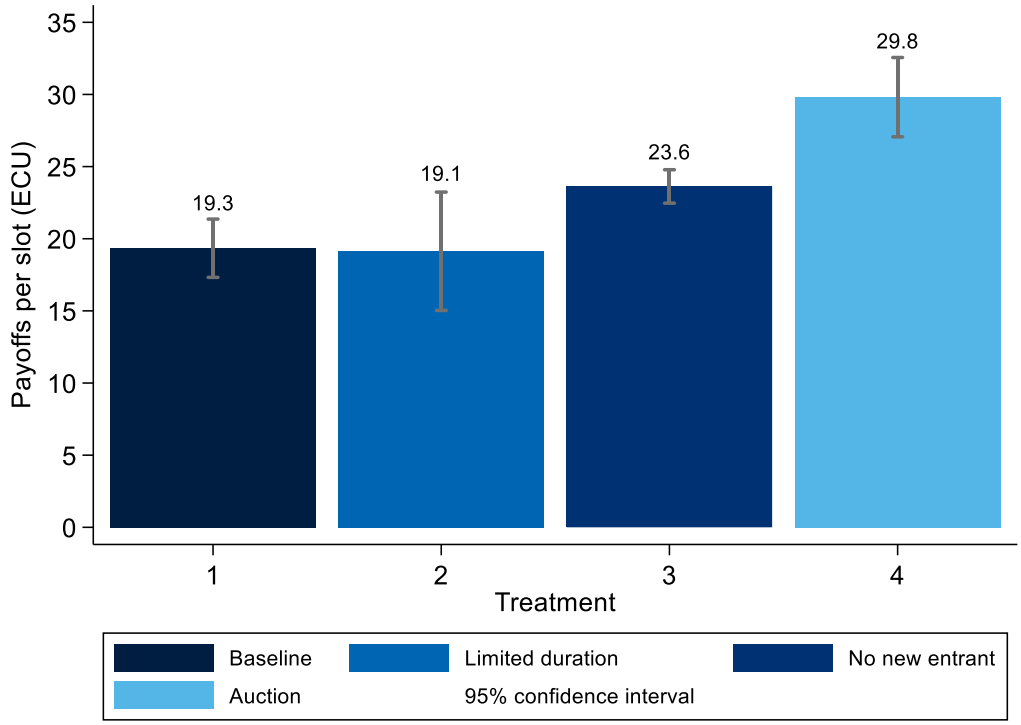

Note: A confidence interval is a way to measure the precision of an estimate. It is a range of values that contain the true statistical value (i.e. across all samples) with a probability of 95%.

Source: Oxera.

At the end of the first stage, the design of the baseline and limited duration treatments is identical, leading to very similar payoffs per slot of ECU 19.3 and ECU 19.1 for treatments 1 and 2 respectively. Since the difference in payoffs between these treatments arises only in the second stage, where trading is allowed in the baseline but not in the limited duration treatment, we can be confident that trading is driving the differences in efficiency.

Furthermore, the results set out in Figure 7.20 suggest that the additional slot coordinator allocation stage has removed the efficiency gains of trading in the first stage for the limited duration treatment. The first stage consists of a slot coordinator allocation followed by a trading stage, leading to a first stage payoff of ECU 19.1 per slot. However, this was followed by a second stage slot coordinator allocation without an opportunity to trade. This allocation effectively 'reset' the benefits of trade in the first stage, leading to an overall lower payoff of ECU 16.7 per slot across both stages.

The average number of trades across treatments is shown in Figure 7.22. The auction treatment has fewer trades, which is likely to be a result of the fact that fewer slots were allocated on average at the bidding stage.

[112] The difference between the baseline and limited duration treatment is not statistically significant at the 5% level using the Mann–Whitney test.

Figure 7.22 Average number of trades in each stage

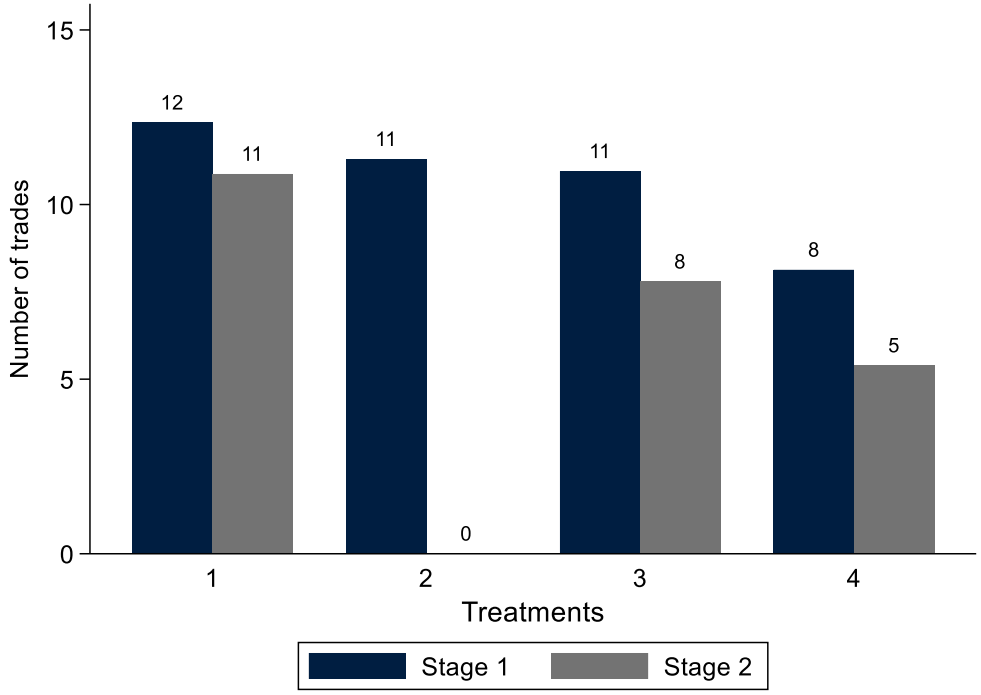

Note: The number of trades in the second stage of treatment 2 is zero because the limited duration treatment has no second trading stage.

Source: Oxera.

7.4.4 Allocative efficiency—consumer surplus

Different slot allocation mechanisms may affect the distribution of slot ownership and therefore the ability of airlines to meet passenger demand. If slots are allocated in a way that encourages competition or leads to routes being offered that are most highly demanded by passengers, we would expect consumer surplus to be higher.

In order to calculate consumer surplus, we would need to know the value that a passenger has placed on a flight (also known as the reservation price) and the price of the flight. For example, if a passenger values a flight to Amsterdam at £100 and pays a price of £55, the consumer surplus is £45. However, in our experiment we have not considered the reservation price for passengers.

In order to determine consumer surplus, we have therefore assumed that consumer surplus is derived entirely from payoffs lost by airlines due to competition and volume effects in the experiment. Therefore, this measure of consumer surplus consists solely of the transfer of surplus from producers to consumers. This should lead to a lower bound on the level of consumer surplus because it assumes that, in the absence of competition, consumer surplus is zero.[113] Further details are provided in Appendix A5.

The consumer surplus outcomes for each treatment are shown in Figure 7.23. The magnitude of the consumer surplus is smaller than the producer surplus, because we are using a conservative measure of consumer surplus. The difference between treatments is qualitatively similar to the difference between airline payoffs, with the auction and no new entrant treatment being more

[113] For airlines to be able to extract all consumer surplus, they would need to undertake first-price discrimination—i.e. the airlines would have to know each passenger's reservation price for their flight and be able to charge exactly the reservation price for that passenger.

efficient than the baseline, while the limited duration treatment is less efficient than the baseline.[114] These results are also consistent with the results in section 7.2.1, where the auction and no new entrant treatment have a lower HHI than the baseline. These results provide some reassurance that the higher allocative efficiency as measured by airline payoffs is not achieved at the expense of benefits to consumers.

Figure 7.23 Consumer surplus per slot by treatment, averaged over all slots

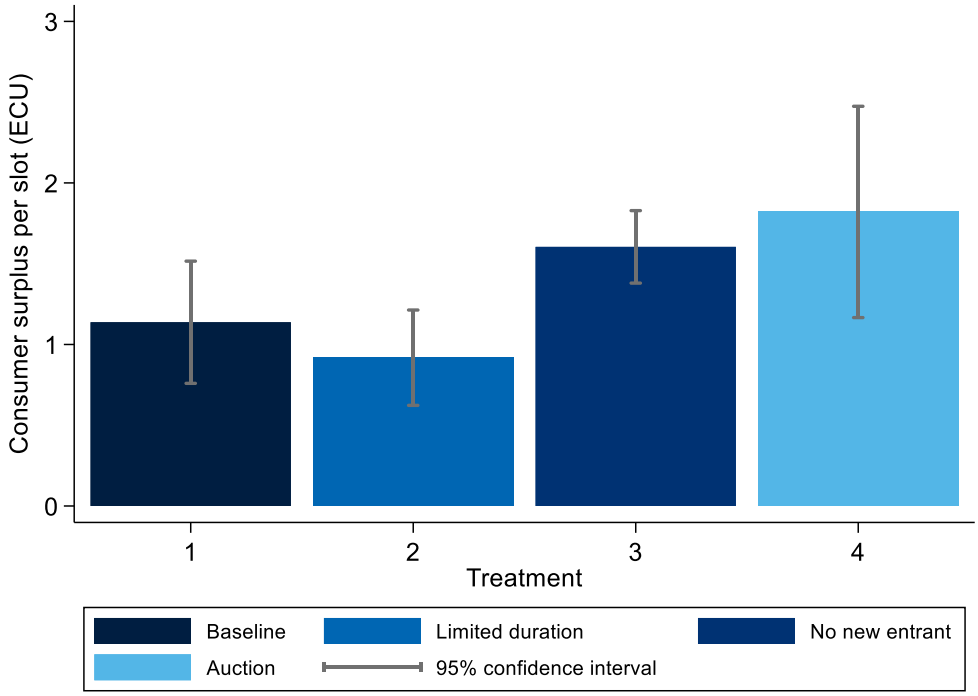

Note: A confidence interval is a way to measure the precision of an estimate. It is a range of values that contain the true statistical value (i.e. across all samples) with a probability of 95%.

Source: Oxera.

7.4.5 Allocative efficiency—total surplus

Ensuring an efficient use of Heathrow capacity is one of the government's objectives, and economic theory shows that an outcome is efficient when the benefits to both passengers and airlines are maximised.

The experiment results suggest that an auction, relative to other treatments, leads to the highest consumer and producer surplus, and therefore maximises total surplus. The no new entrant rule also leads to higher total surplus relative to the baseline, as it allows slots that would otherwise be allocated to new entrants to be allocated more efficiently. Conversely, the limited duration treatment leads to a relatively inefficient outcome due to fewer opportunities for trading.

7.5 Trading

In this section, we explore the outcomes of trading. In particular, we consider which airlines in the experiment are net sellers and which are net buyers of slots. The net number of slots bought, which is the total number of slots bought

[114] A Mann–Whitney test at the 5% significance level shows a statistically significant treatment effect for the auction and no new entrant treatments relative to the baseline, while it does not show a statistically significant difference between the limited duration treatment and the baseline treatment.

in both stages 1 and 2 subtracted by the number of total slots sold in stages 1 and 2, was calculated for each airline in each round, and averaged across all rounds in a treatment. This is shown in Figure 7.24.

Figure 7.24 Net number of slots bought by each airline

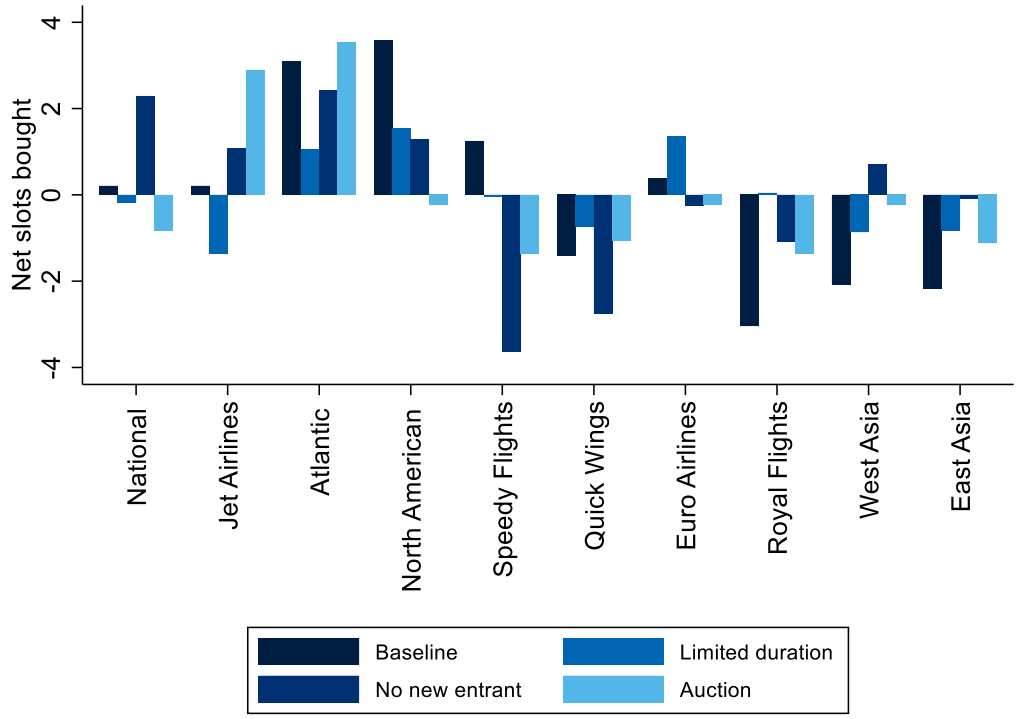

Source: Oxera.

We observe that the airlines that tend to be net buyers are North American and Atlantic, while the net sellers of slots are the low-cost and Middle Eastern/Asian carriers. There are also some airlines that are net buyers in some treatments and net sellers in other treatments (such as Euro Airlines and Jet Airlines). Given this variation, it is difficult to draw clear conclusions, but it appears that airlines that can obtain only a relatively low payoff from their routes, or that can profitably operate only a few routes, are more likely to be net sellers of slots. An implication of this result is that, even if an airline places a relatively low value on slots, it still has an incentive to bid for slots from the slot coordinator because it will be able to sell them to other airlines.

7.5.1 Impact of trading on competition

Trading may affect competition as airlines change their slot holdings, and routes flown may also change as a result. In order to investigate the effect of trade on competition, we calculated the change in HHI at a route level between the first and second stages of each round, and compared this across treatments.[115]

As shown in Figure 7.25, we find that there is not a clear relationship between trading and competition. Trading reduces route-level HHI in the baseline but increases HHI in the no new entrant and auction treatments. The HHI also increases in the limited duration treatment, but this is due to the second bidding stage as opposed to trading. Only the auction results are statistically

[115] This assessed the impact of the second trade stage on route-level competition. It would not be possible to assess the impact of the first trade stage on route-level competition, as participants did not reveal their route choice until after the first trade stage.

significant, but even here the confidence intervals are large (and thus the result should be treated with caution).

Figure 7.25 Effect of trade on competition: change in HHI at a route level between stages 1 and 2

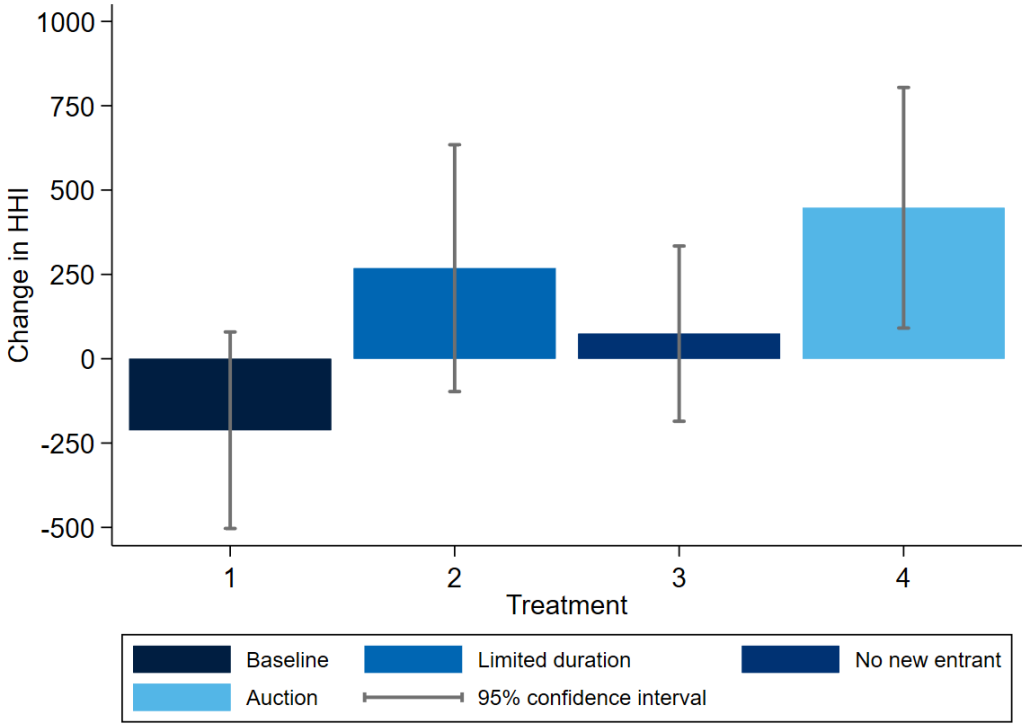

Note: Treatment 2 does not have a second trade stage—the impact shown is due to the second bidding stage. A confidence interval is a way to measure the precision of an estimate. It is a range of values that contain the true statistical value (i.e. across all samples) with a probability of 95%.

Source: Oxera.

7.6 Bidding behaviour

Bids to the slot coordinator (rather than to an auctioneer) are a key part of the slot allocation process in treatments 1, 2 and 3. In treatments 1, 2 and 3, participants are able to acquire slots for free, which may lead to two types of behaviour:

- **altering the route choice after winning slots**. When submitting their bids for slots, participants state a route that they intend to fly. This influences the decision of the slot coordinator. However, airlines can declare that they will fly on certain routes to try and increase their chances of being awarded a slot, but then fly a different route that will be more profitable;
- **direct reselling**. Some airlines may exploit the priority rules (e.g. those for new entrants) to obtain some slots that they can then resell to airlines with greater valuation for these slots.

To understand the effect of the different mechanisms on bidding behaviour, Table 7.3 shows the proportion of airlines altering route choice after winning slots, and direct reselling. Overall, we find that bidding behaviour is similar across all treatments. Altering route choice after winning slots is widespread

(affecting around two-thirds of bids), and we do not find evidence of significant differences across treatments.[116]

Similarly, we find that, on average, 21% of participants directly resold slots that they acquired. Only the difference between treatment 1 and treatment 2 is statistically significant at a 10% level.

Table 7.3 Share of false declaration and direct reselling

	Treatment 1	Treatment 2	Treatment 3	Treatment 4
Altering route choice after winning slots (percentage of allocated slots)	66%	66%	64%	n.a.
Direct reselling (percentage of allocated slots)	24%	19%	20%	21%

Note: In treatment 4 airlines do not have to declare which route they intend to fly before submitting their bids in the auction.

Source: Oxera.

We note that these behaviours are less likely to occur in reality than in the experiment. This is because, in reality, airlines engage in repeated interactions with the slot coordinator and are likely to want to avoid harming their relationships with the slot coordinator and other airlines.

7.7 Summary of results

In summary, we find that the no new entrant rule and auction treatments increase route-level competition. The auction is the most effective way of increasing allocative efficiency, but reduces the share of slots held by new entrants. While trading increases efficiency, it does not fully mitigate the effects of an inefficient starting allocation.

In more detail, we find the following.

- Both treatment 3 (no new entrant rule) and treatment 4 (auction) increase route-level competition (the primary measure of competition considered). However, treatment 4 reduces the share of slots held by new entrant airlines. In treatment 2 (limited duration), the level of competition is similar to in treatment 1 (baseline), but the market share of new entrants is the highest.
- Given the mix of routes operated by airlines, the effect of competition from the release of new slots is greatest on long-haul routes.
- We also analyse allocative efficiency (i.e. whether capacity is allocated in a way that maximises social welfare) as the sum of producer surplus and consumer surplus. Trading improves allocative efficiency, although it does not fully mitigate inefficiencies that can arise in the current allocation mechanism. While auction and the baseline provide the same opportunities for trading, the auction results in higher allocative efficiency. This is because the auction allocates slots to airlines that are best placed to meet passenger demand, and therefore have the largest incentive to bid for slots. Treatment 3 (no new entrant rule) also leads to higher allocative efficiency than treatment 1 (baseline), as it allows slots that would otherwise be allocated to new entrants to be allocated more efficiently—although it is still less efficient

[116] We cannot reject the null hypothesis setting significance at a 5% threshold.

than the auction treatment. Treatment 2 (limited duration) leads to a relatively inefficient outcome as there are fewer opportunities for trading.

- We also look at productive efficiency (i.e. whether the airport infrastructure is fully utilised) by considering the share of long-haul flights and the average seat-km per flight. Productive efficiency is lowest in treatment 3 (no new entrant rule) due to a change in the mix of airlines, with fewer Middle Eastern and Asian airlines acquiring slots. On the other hand, treatment 3 (no new entrant rule) increases the number of domestic flights.
- Airlines alter their route choice after winning slots where there is an administrative allocation mechanism (treatments 1, 2 and 3), but such behaviour may be less common in reality as airlines are likely to want to avoid harming their relationships with the slot coordinator.

A1 Stakeholder engagement and data sources

A1.1 Stakeholder engagement

In addition to a number of industry experts, we engaged with the following stakeholders (in alphabetical order).

- Airport Coordination Limited (ACL)
- All Nippon Airways (ANA)
- American Airlines
- Avianca
- British Airways
- Civil Aviation Authority (CAA)
- Competition and Markets Authority (CMA)
- Delta Air Lines
- Deutsche Lufthansa
- FedEx
- Finnair
- Heathrow Airport Limited
- Hong Kong Express Airways
- International Air Transport Association (IATA)
- KLM Royal Dutch Airlines
- National Air Traffic Services (NATS)
- Norwegian Air Shuttle
- Qantas Airways
- TUI fly
- Turkish Airlines
- United Airlines
- Virgin Atlantic Airways
- Wizz Air

A1.2 Data sources

Our data analysis and desk-based research focused on a number of key sources:

- publicly available DfT demand traffic forecasts;
- insights from the DfT fleet mix model;
- OAG schedule data at Heathrow, other London airports and other key European airports;

- publicly available ACL data;
- publicly available responses to the Airports Commission;
- publicly available annual and financial reports from airlines;
- publicly available fleet order books;
- other publicly available sources of information and reports on slot allocation;
- RDC Aviation data on the profitability of routes flown from Heathrow and other London airports, by airline.

A2 Experiment screenshots and key information sheets

This appendix presents screenshots of the experiment and the hard copy instructions.

A2.1 Treatment 1: baseline ('control')

Figure A2.1 Welcome screen

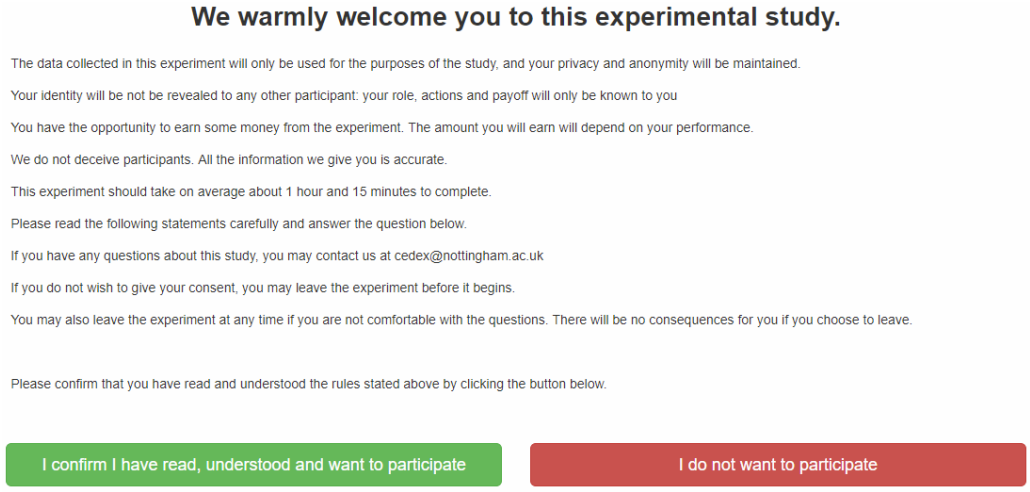

Figure A2.2 Thanks for your participation

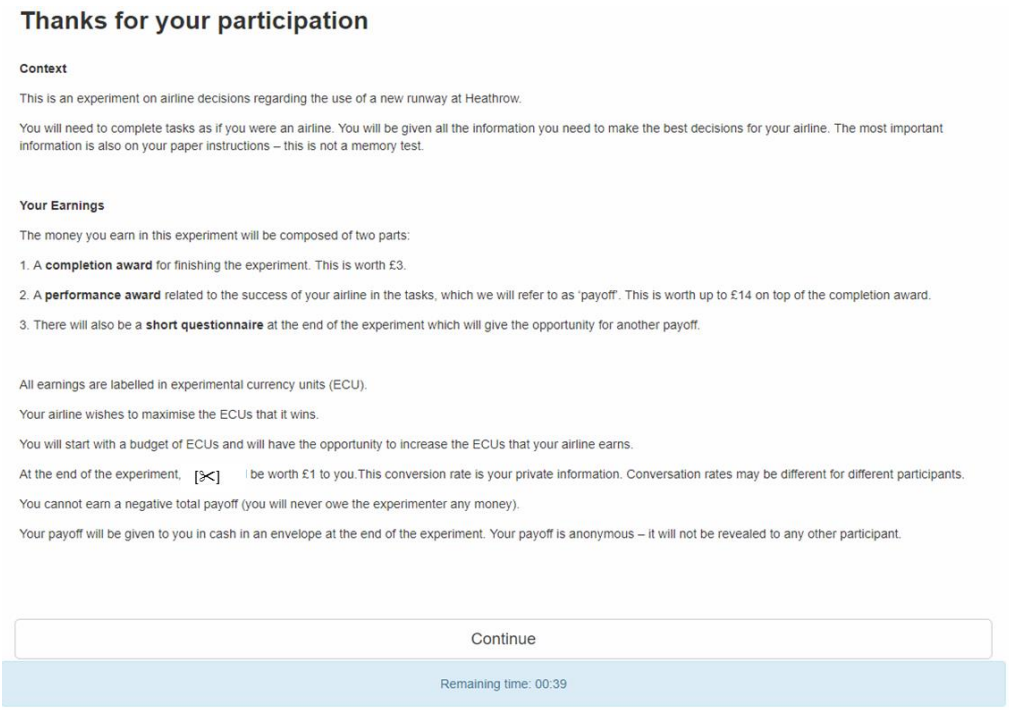

Figure A2.3 Slots

Slots

In order for an airline to fly to and from an airport, it must hold runway slots.

Each slot gives the airline the right to use the airport (for a departure or arrival) at a given time.

Slots are awarded to airlines in pairs, so that airlines can arrive and depart from an airport.

Slots are allocated to airlines by the computer. We will explain how the computer makes this decision.

While some of the airlines already fly from Heathrow, this study is about how the new slots will be used.

There are 48 new slots in total, at five times of day. This will be clear to you when you come to make decisions.

Half the slots are 'new entrant slots'. We will explain what this means.

	6am-8am	8am-10am	10am-2pm	2pm-6pm	6pm-10pm
New entrant slots	3	3	6	6	6
General slots	3	3	6	6	6
Total	6	6	12	12	12

→ 48 slots in total

Airlines

There are 10 airlines, each played by one participant. Each airline has a different budget and earns different payoffs.

There are 20 possible routes that airlines can fly from Heathrow, but some airlines cannot fly to all 20 routes. You will be told where you can fly.

You are: **National Airlines**

- **National Airlines**, which uses Heathrow as its 'hub airport', flying from Heathrow all over the world.
- **Jet Flights**, which aims to make Heathrow its 'hub airport', flying from Heathrow all over the world.
- **Atlantic Airlines**, a USA-based airline, flying from Heathrow to the USA.
- **North American Airways**, a USA-based airline, flying from Heathrow to the USA.
- **Quick Wings**, a low-cost airline, flying from Heathrow to the USA.
- **Speedy Flights**, a low-cost airline, flying from Heathrow to other parts of the UK and Europe.
- **Euro Airlines**, a German airline, flying from Heathrow to Germany.
- **Royal Flights**, a Dutch airline, flying from Heathrow to the Netherlands.
- **West Asia Airways**, a United Arab Emirates airline, flying from Heathrow to the United Arab Emirates .
- **East Asia Airways**, a Chinese airline, flying from Heathrow to China.

Please write the number of the computer you are:

Continue

Remaining time: 02:54

Figure A2.4 Rounds

Rounds

There will be 3 identical rounds of the same scenario.

Each round is separate: at the start of each round, your budget goes back to where it started (your payoff from one round does not affect other rounds).

One of the 3 rounds will be randomly selected to determine your performance award at the end of the experiment.

There will first be a practice round, which will not be used to determine your performance award. The practice round is otherwise identical to the other 3 rounds.

Each round contains a number of stages, where you will bid for slots, trade slots between airlines, and decide which route to fly with the slots you win.

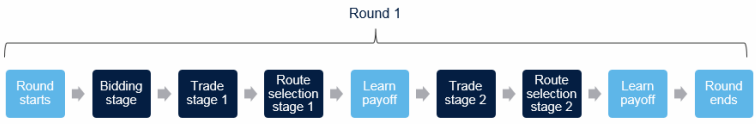

You will always be told what stage you are in, and what decisions you can make.

On the next screen we show you screenshots of what each stage looks like.

Continue

Remaining time: 00:51

Figure A2.5 Screenshots

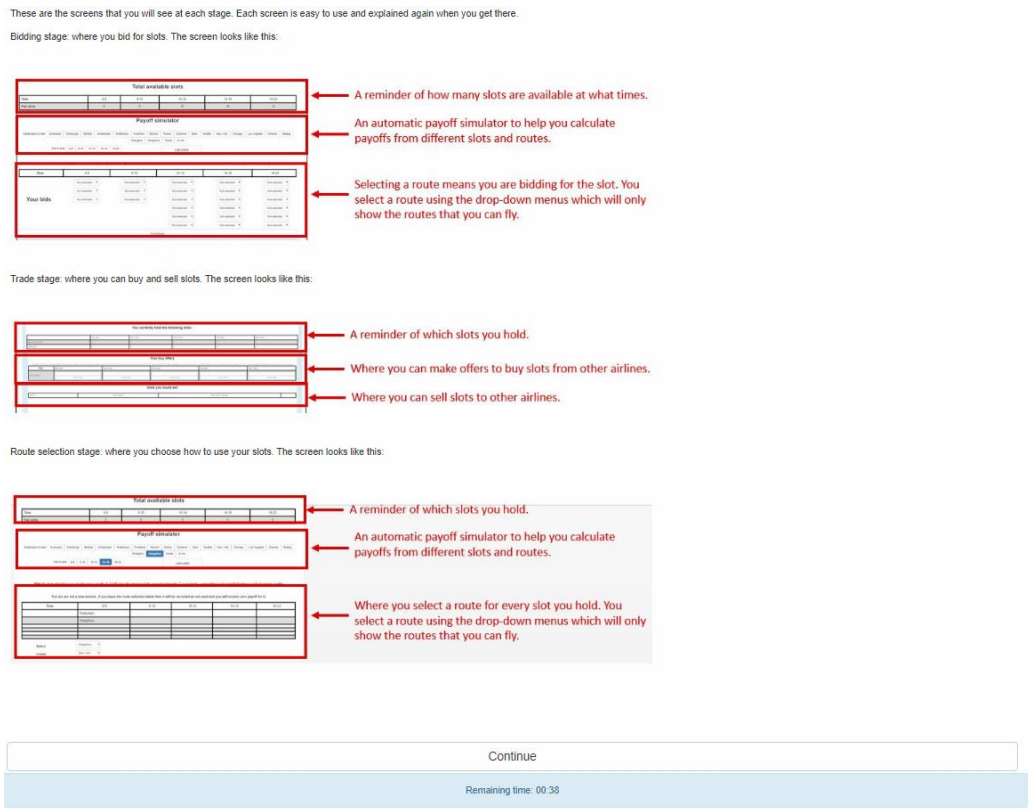

Figure A2.6 Your airline

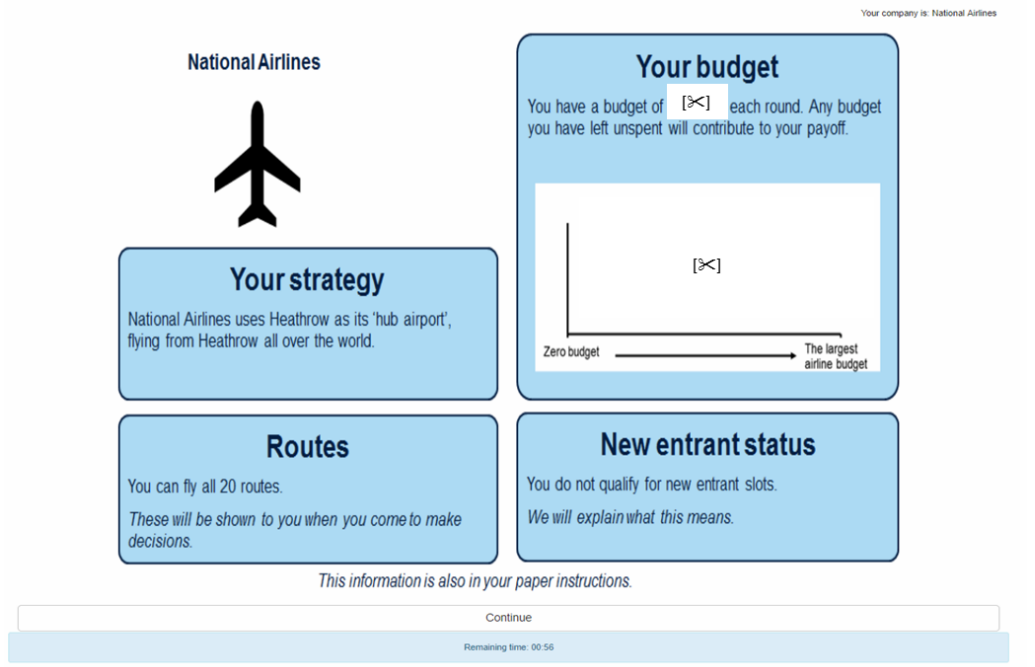

Figure A2.7 Your airline's payoffs

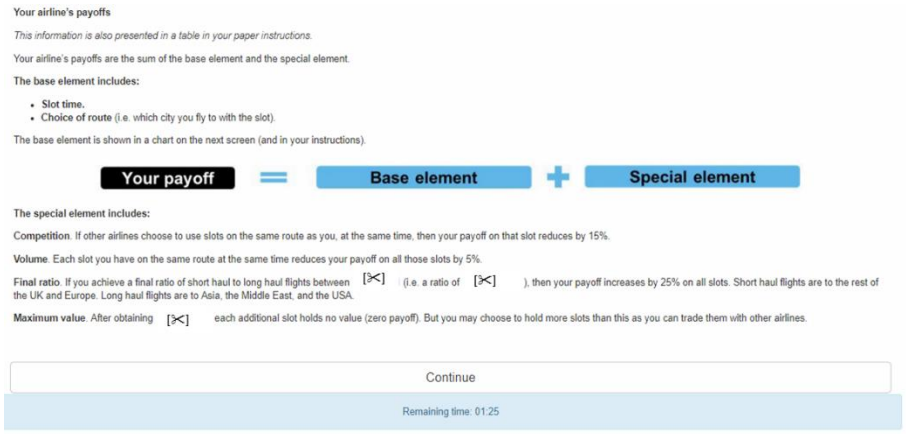

Figure A2.8 The base element of your payoffs

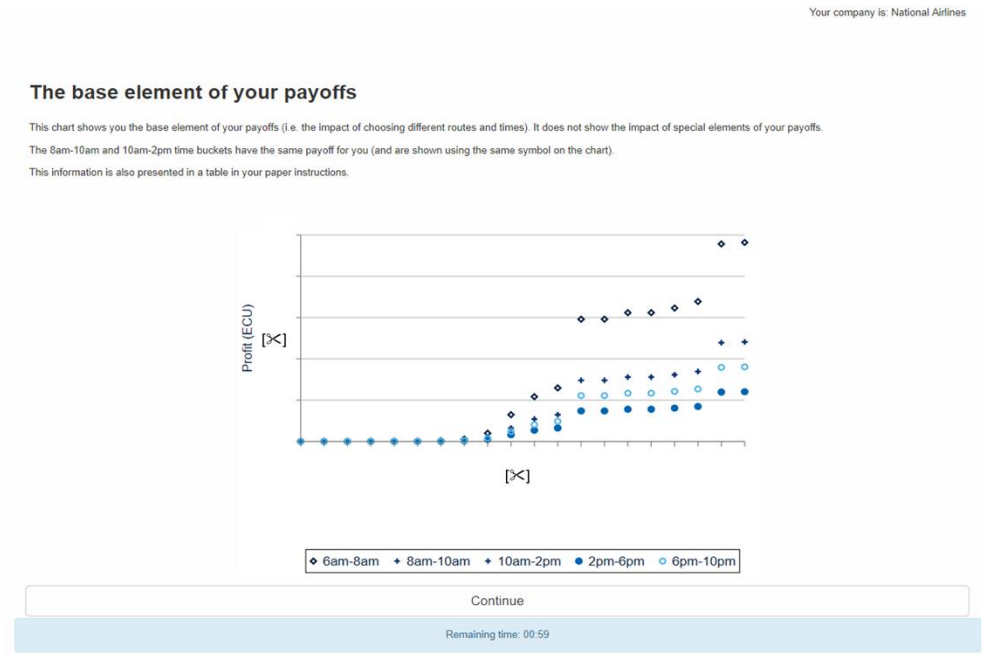

Figure A2.9 The computer (part 1)

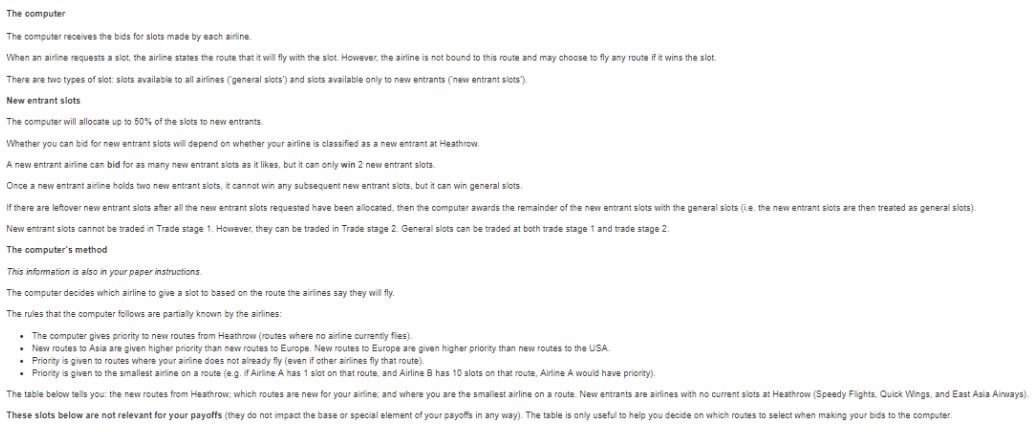

Figure A2.10 The computer (part 2)

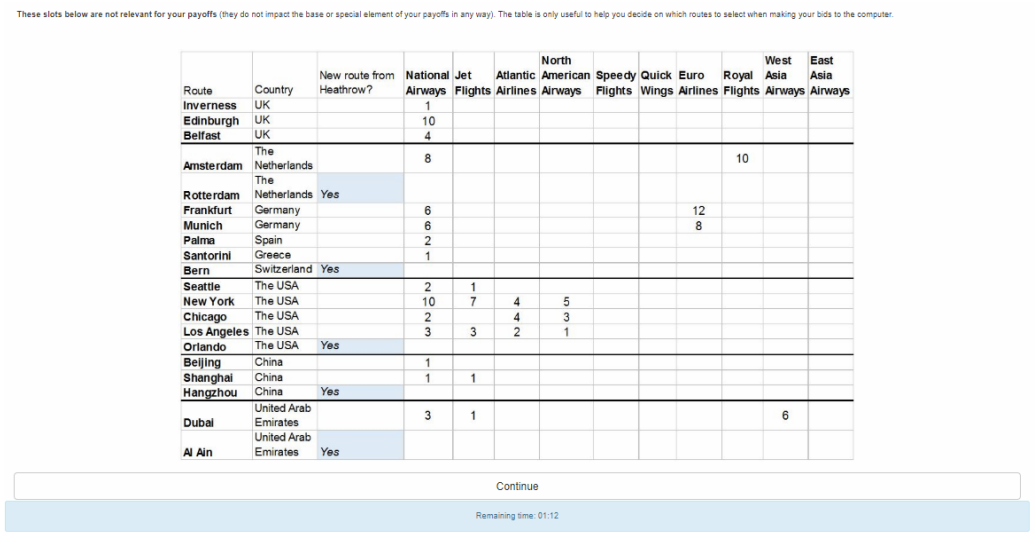

These slots below are not relevant for your payoffs (they do not impact the base or special element of your payoffs in any way). The table is only useful to help you decide on which routes to select when making your bids to the computer.

Route	Country	New route from Heathrow?	National Airways	Jet Flights	Atlantic Airlines	North American Airways	Speedy Flights	Quick Wings	Euro Airlines	Royal Flights	West Asia Airways	East Asia Airways
Inverness	UK		1									
Edinburgh	UK		10									
Belfast	UK		4									
Amsterdam	The Netherlands		8							10		
Rotterdam	The Netherlands	Yes										
Frankfurt	Germany		6						12			
Munich	Germany		6						8			
Palma	Spain		2									
Santorini	Greece		1									
Bern	Switzerland	Yes										
Seattle	The USA		2	1								
New York	The USA		10	7	4	5						
Chicago	The USA		2		4	3						
Los Angeles	The USA		3	3	2	1						
Orlando	The USA	Yes										
Beijing	China		1									
Shanghai	China		1	1								
Hangzhou	China	Yes										
Dubai	United Arab Emirates		3	1							6	
Al Ain	United Arab Emirates	Yes										

Continue

Remaining time: 01:12

Figure A2.11 Test questions

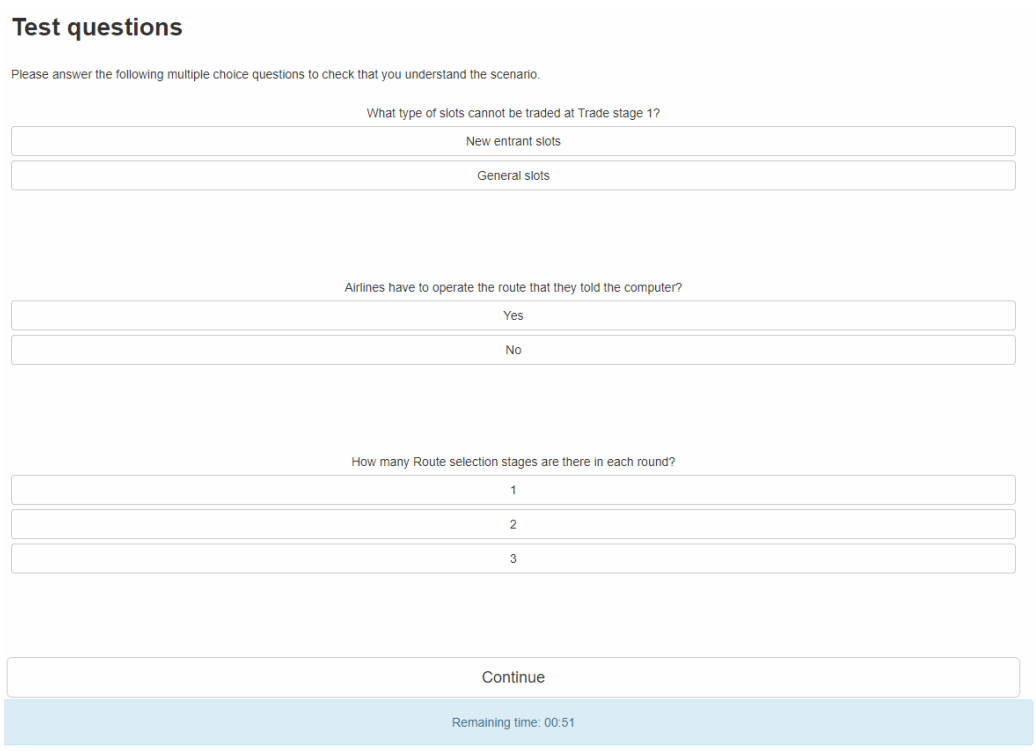

Test questions

Please answer the following multiple choice questions to check that you understand the scenario.

What type of slots cannot be traded at Trade stage 1?

New entrant slots

General slots

Airlines have to operate the route that they told the computer?

Yes

No

How many Route selection stages are there in each round?

1

2

3

Continue

Remaining time: 00:51

Figure A2.12 Waiting for other airlines

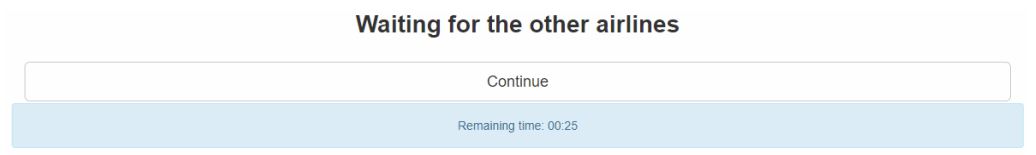

Waiting for the other airlines

Continue

Remaining time: 00:25

Figure A2.13 Bidding stage 1

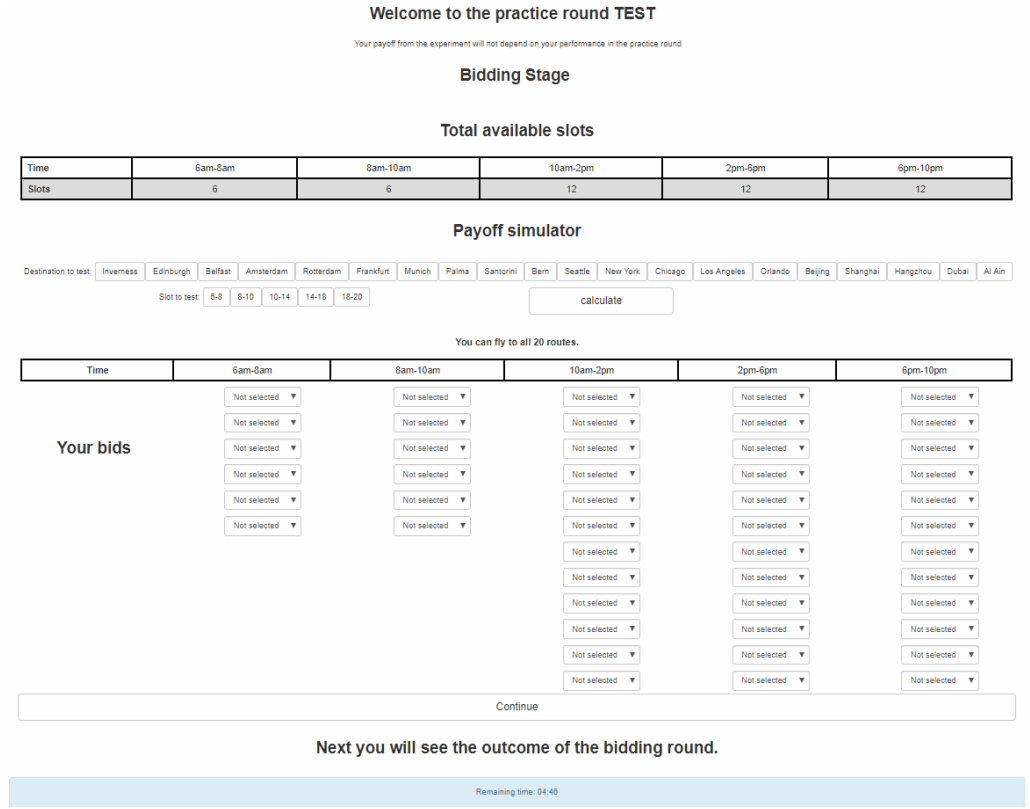

Figure A2.14 Your slots

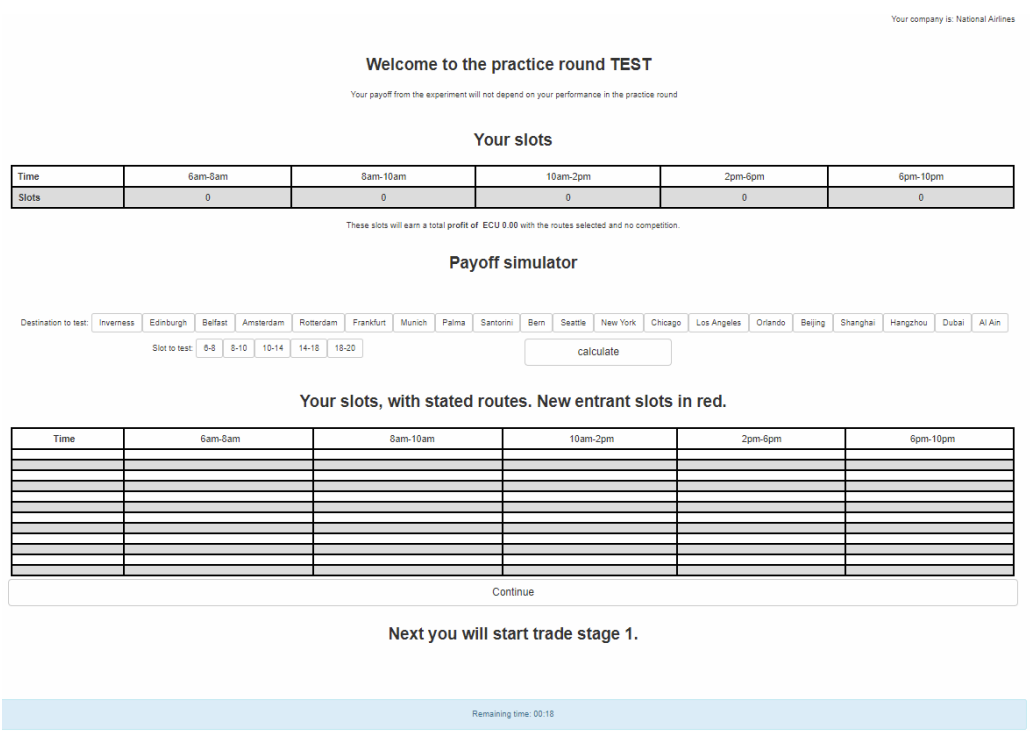

Figure A2.15 Trade stage 1

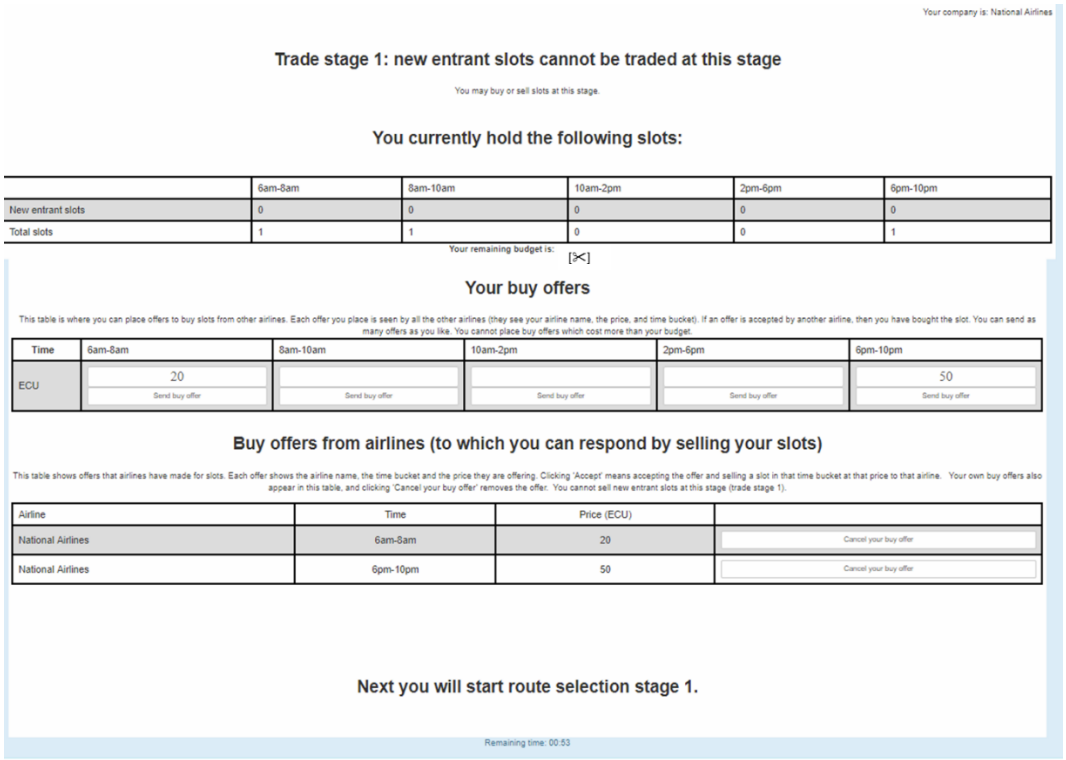

Figure A2.16 Route selection stage 1

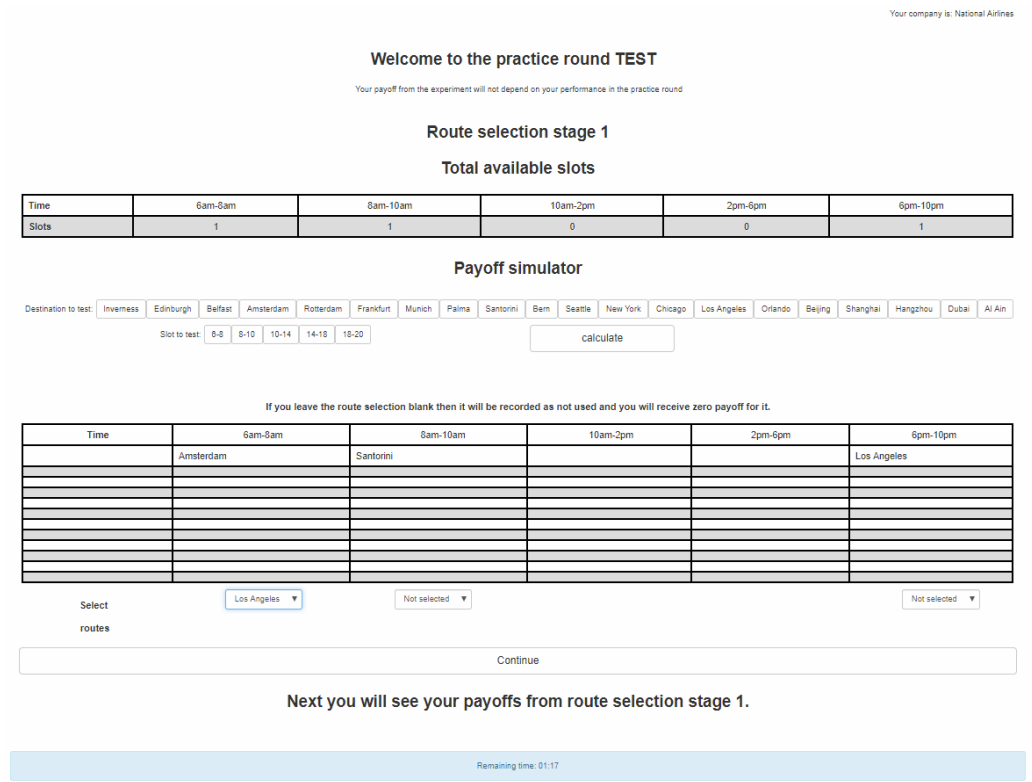

Figure A2.17 Your payoff after route selection stage 1

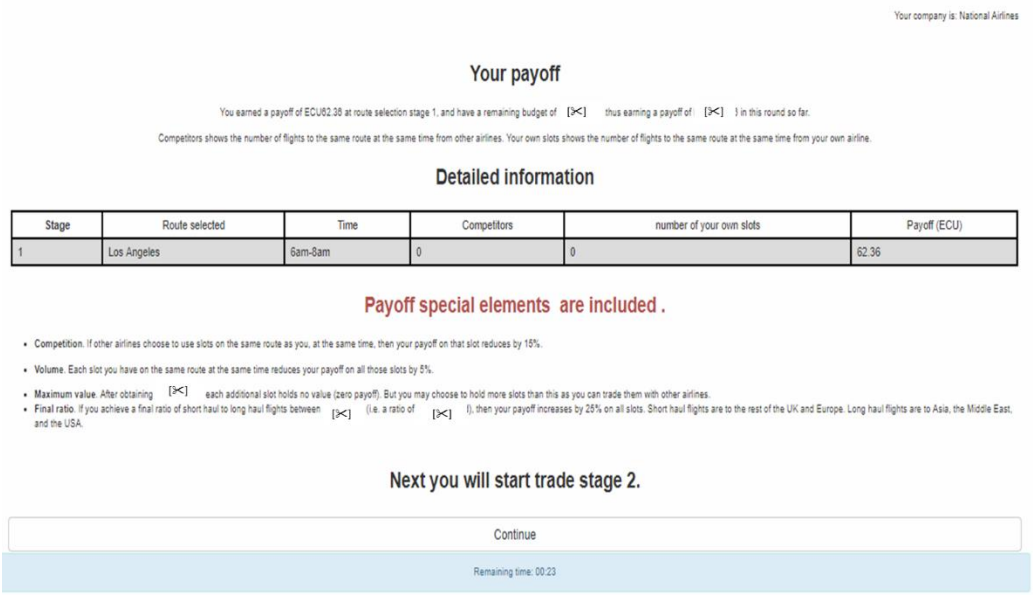

Your company is: National Airlines

Your payoff

You earned a payoff of ECU62.36 at route selection stage 1, and have a remaining budget of [✂] thus earning a payoff of [✂] ł in this round so far.

Competitors shows the number of flights to the same route at the same time from other airlines. Your own slots shows the number of flights to the same route at the same time from your own airline.

Detailed information

Stage	Route selected	Time	Competitors	number of your own slots	Payoff (ECU)
1	Los Angeles	6am-8am	0	0	62.36

Payoff special elements are included .

- **Competition**. If other airlines choose to use slots on the same route as you, at the same time, then your payoff on that slot reduces by 15%.
- **Volume**. Each slot you have on the same route at the same time reduces your payoff on all those slots by 5%.
- **Maximum value**. After obtaining [✂] each additional slot holds no value (zero payoff). But you may choose to hold more slots than this as you can trade them with other airlines.
- **Final ratio**. If you achieve a final ratio of short haul to long haul flights between [✂] (i.e. a ratio of [✂] l), then your payoff increases by 25% on all slots. Short haul flights are to the rest of the UK and Europe. Long haul flights are to Asia, the Middle East, and the USA.

Next you will start trade stage 2.

Continue

Remaining time: 00:23

Figure A2.18 Trade stage 2

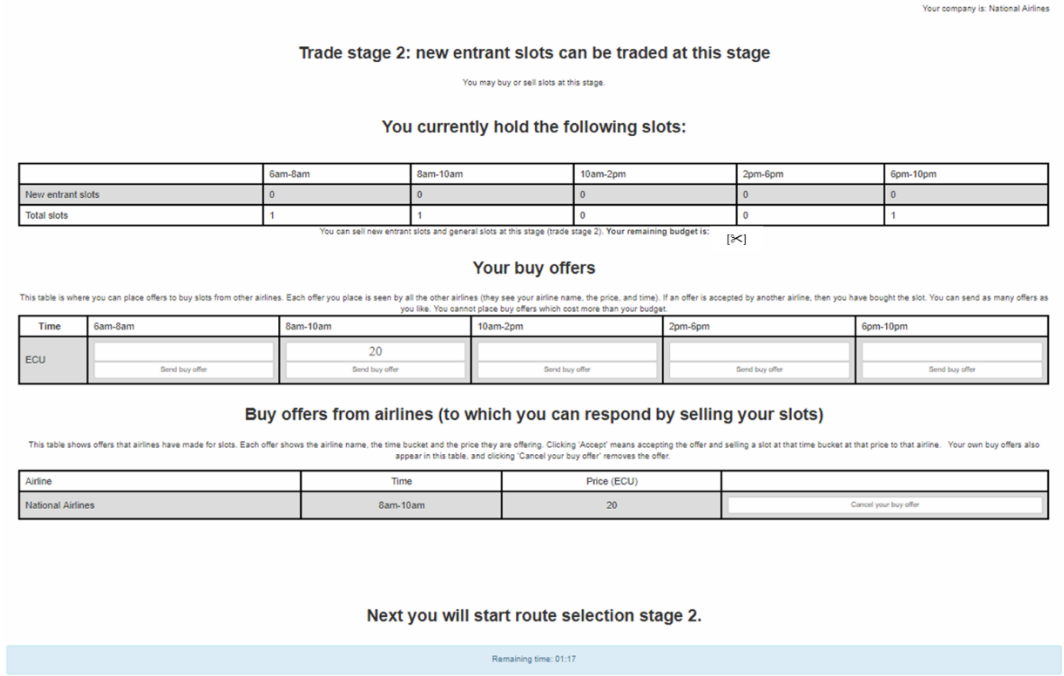

Your company is: National Airlines

Trade stage 2: new entrant slots can be traded at this stage

You may buy or sell slots at this stage.

You currently hold the following slots:

	6am-8am	8am-10am	10am-2pm	2pm-6pm	6pm-10pm
New entrant slots	0	0	0	0	0
Total slots	1	1	0	0	1

You can sell new entrant slots and general slots at this stage (trade stage 2). **Your remaining budget is:** [✂]

Your buy offers

This table is where you can place offers to buy slots from other airlines. Each offer you place is seen by all the other airlines (they see your airline name, the price, and time). If an offer is accepted by another airline, then you have bought the slot. You can send as many offers as you like. You cannot place buy offers which cost more than your budget.

Time	6am-8am	8am-10am	10am-2pm	2pm-6pm	6pm-10pm
ECU	Send buy offer	20 Send buy offer	Send buy offer	Send buy offer	Send buy offer

Buy offers from airlines (to which you can respond by selling your slots)

This table shows offers that airlines have made for slots. Each offer shows the airline name, the time bucket and the price they are offering. Clicking 'Accept' means accepting the offer and selling a slot at that time bucket at that price to that airline. Your own buy offers also appear in this table, and clicking 'Cancel your buy offer' removes the offer.

Airline	Time	Price (ECU)	
National Airlines	8am-10am	20	Cancel your buy offer

Next you will start route selection stage 2.

Remaining time: 01:17

Figure A2.19 Route selection stage 2

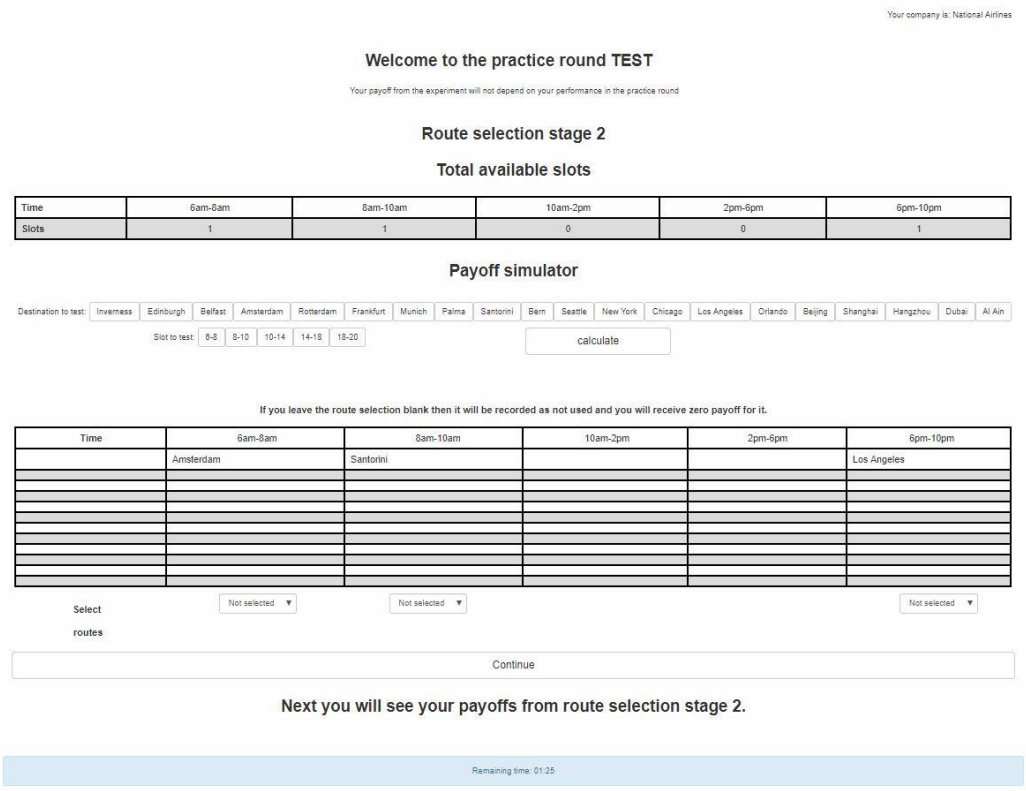

Figure A2.20 Your payoff after route selection stage 2

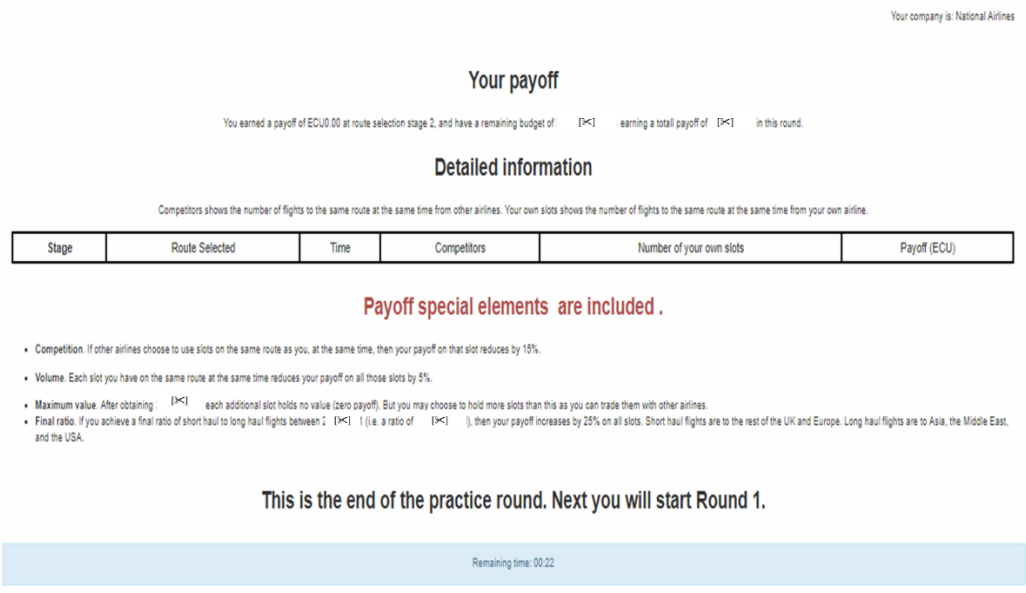

Figure A2.21 Incentivised questions

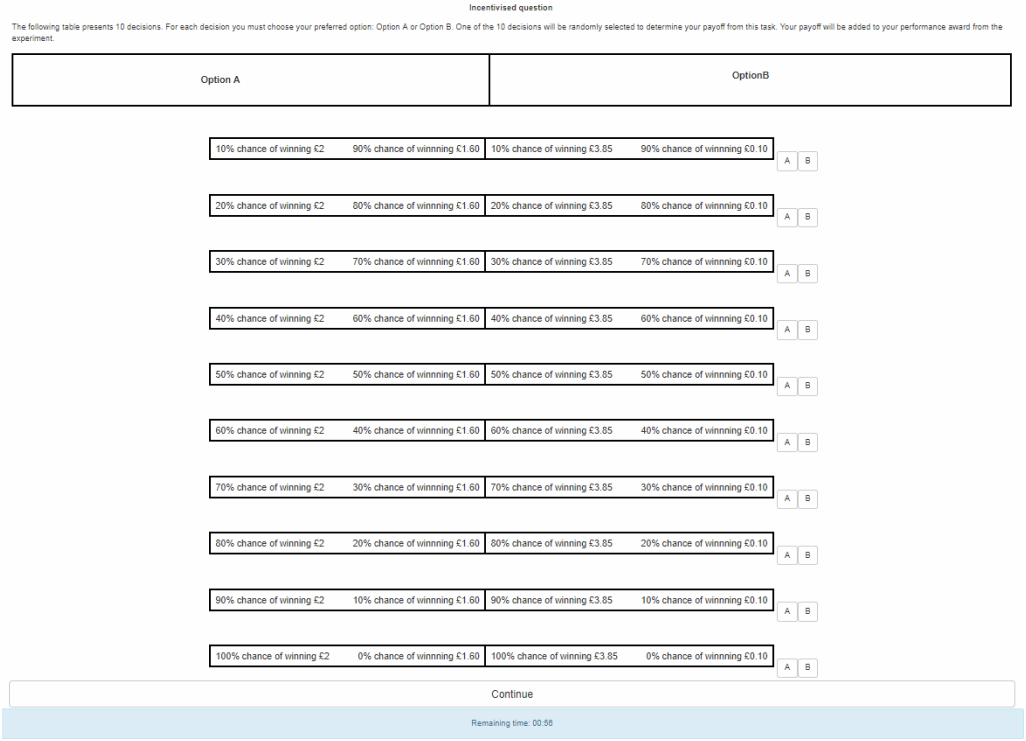

Figure A2.22 Your total earnings

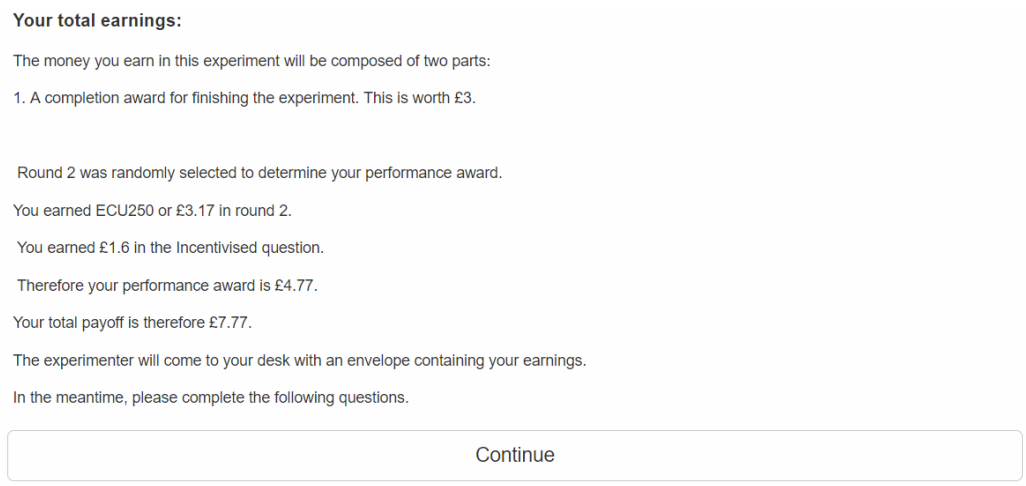

Figure A2.23 Questionnaire

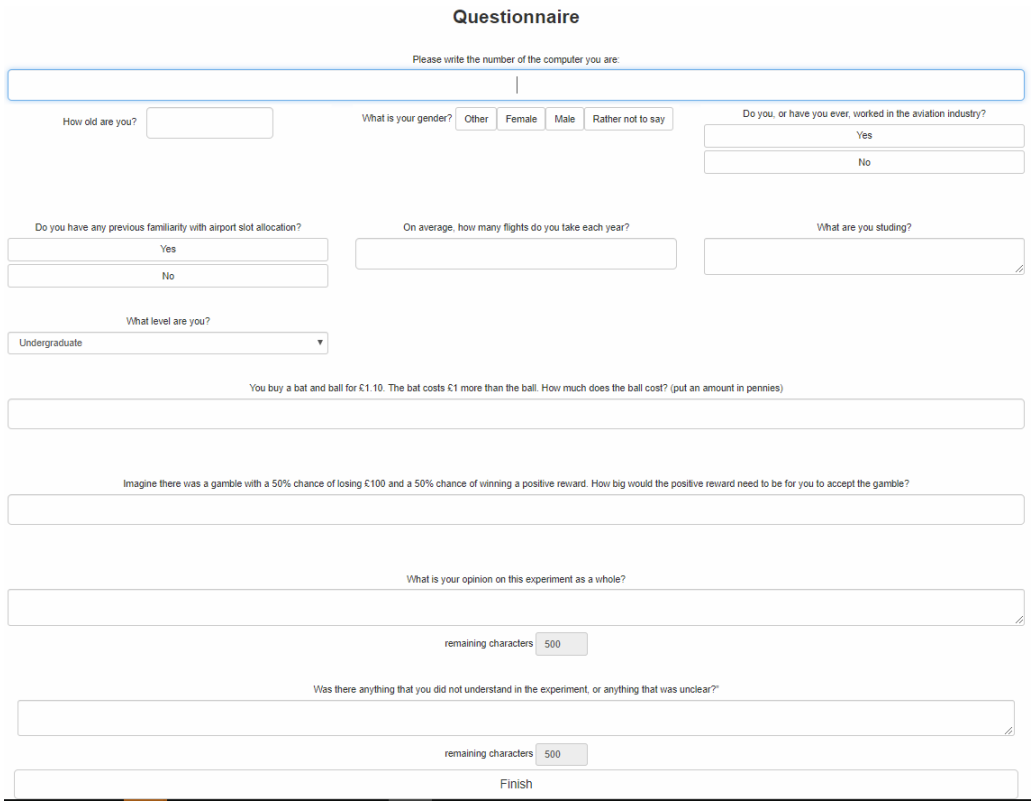

Figure A2.24 Thank you for taking part in this experiment

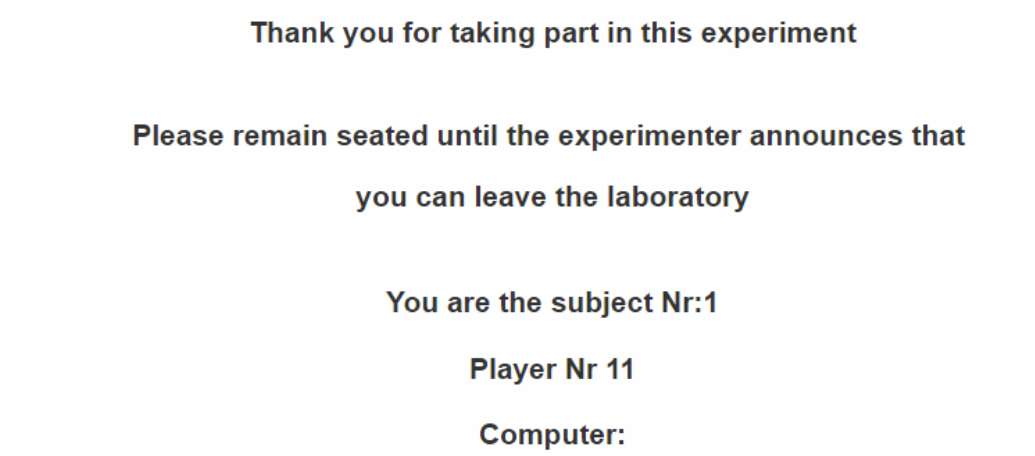

A2.1.1 Key information sheet

Figure A2.25 Key information sheet, treatment 1, National Airlines (page 1)

Key information

Welcome to the experiment! This sheet summarises key information that is important for the experiment so you always have it at hand. Full instructions are displayed on your screen.

Your airline

National Airlines

Your budget

You have a budget of [✂] each round. Any budget you have left unspent will contribute to your payoff.

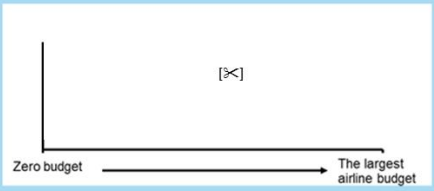

Your strategy

National Airlines uses Heathrow as its 'hub airport', flying from Heathrow all over the world.

Routes

You can fly all 20 routes.

These will be shown to you when you come to make decisions.

New entrant status

You do not qualify for new entrant slots.

We will explain what this means.

Yourpayoff

Your payoff = **Base element** + **Special element**

The base element includes:

- **Slot time**.
- **Choice of route** (i.e. which city you fly to with the slot).

The special element is different for each airline. Yours includes:

- **Competition**. If other airlines choose to use slots on the same route as you, at the same time, then your payoff on that slot reduces by 15%.
- **Volume**. Each slot you have on the same route at the same time reduces your payoff on all those slots by 5%.
- **Final ratio**. If you achieve a final ratio of short haul to long haul flights between [✂] (i.e. a ratio of [✂]), then your payoff increases by 25% on all slots. Short haul flights are to the rest of the UK and Europe. Long haul flights are to Asia, the Middle East, and the USA.
- **Maximum value**. After obtaining [✂], each additional slot holds no value (zero payoff). But you may choose to hold more slots than this as you can trade them with other airlines.

Figure A2.26 Key information sheet, treatment 1, National Airlines (page 2)

The base element of your payoff

This chart and table show you the base element of your payoffs (i.e. the impact of choosing different routes and times). It does not show the impact of special elements of your payoffs.

The chart and table show the same information.

The 8am-10am and 10am-2pm time buckets have the same payoff for you (and are shown using the same symbol on the chart).

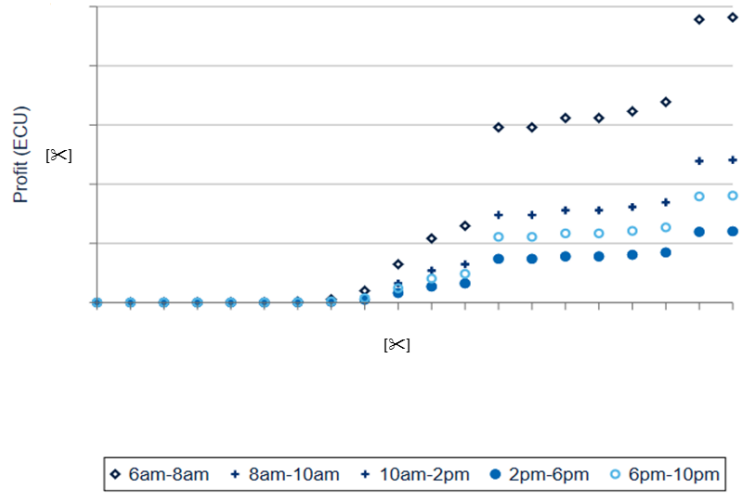

[✂]

Each round

Each round contains a number of stages, where you will bid for slots, trade slots between airlines, and decide which route to fly with the slots you win.

Figure A2.27 Key information sheet, treatment 1, National Airlines (page 3)

The computer

The computer decides which airline to give a slot to based on the route the airlines say they will fly.

The rules that the computer follows are partially known by the airlines:

- The computer gives priority to new routes from Heathrow (routes where no airline currently flies).
- New routes to Asia are given higher priority than new routes to Europe. New routes to Europe are given higher priority than new routes to the USA.
- Priority is given to routes your airline is not already flying to (but other airlines may fly to).
- Priority is given to the smallest airline on a route (e.g. if Airline A has 1 slot on that route, and Airline B has 10 slots on that route, Airline A would have priority).

The table below tells you: the new routes from Heathrow; which routes are new for your airline; and where you are the smallest airline on a route. New entrants are airlines with no current slots at Heathrow (Speedy Flights, Quick Wings, and East Asia Airways).

These slots below are not relevant for your payoffs (they do not impact the base or special element of your payoffs in any way). The table is only useful to help you decide on which routes to select when making your bids to the computer.

Number of slots to each route from Heathrow, by airline

Route	Country	New route from Heathrow?	**National Airways**	**Jet Flights**	**Atlantic Airlines**	**North American Airways**	**Speedy Flights**	**Quick Wings**	**Euro Airlines**	**Royal Flights**	**West Asia Airways**	**East Asia Airways**
Inverness	UK		1									
Edinburgh	UK		10									
Belfast	UK		4									
Amsterdam	The Netherlands		8							10		
Rotterdam	The Netherlands	*Yes*										
Frankfurt	Germany		6						12			
Munich	Germany		6						8			
Palma	Spain		2									
Santorini	Greece		1									
Bern	Switzerland	*Yes*										
Seattle	The USA		2	1								
New York	The USA		10	7	4	5						
Chicago	The USA		2		4	3						
Los Angeles	The USA		3	3	2	1						
Orlando	The USA	*Yes*										
Beijing	China		1									
Shanghai	China		1	1								
Hangzhou	China	*Yes*										
Dubai	United Arab Emirates		3	1							6	
Al Ain	United Arab Emirates	*Yes*										

A2.2 Treatment 2: limited duration rights

This section presents screenshots for the screens that differ between Treatment 2 and Treatment 1.

Figure A2.28 Rounds

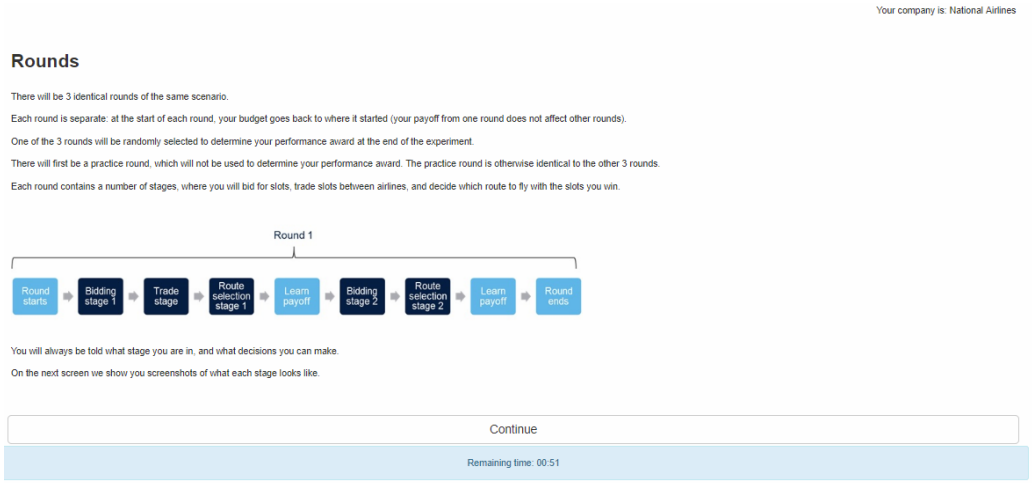

Figure A2.29 Bidding stage 2

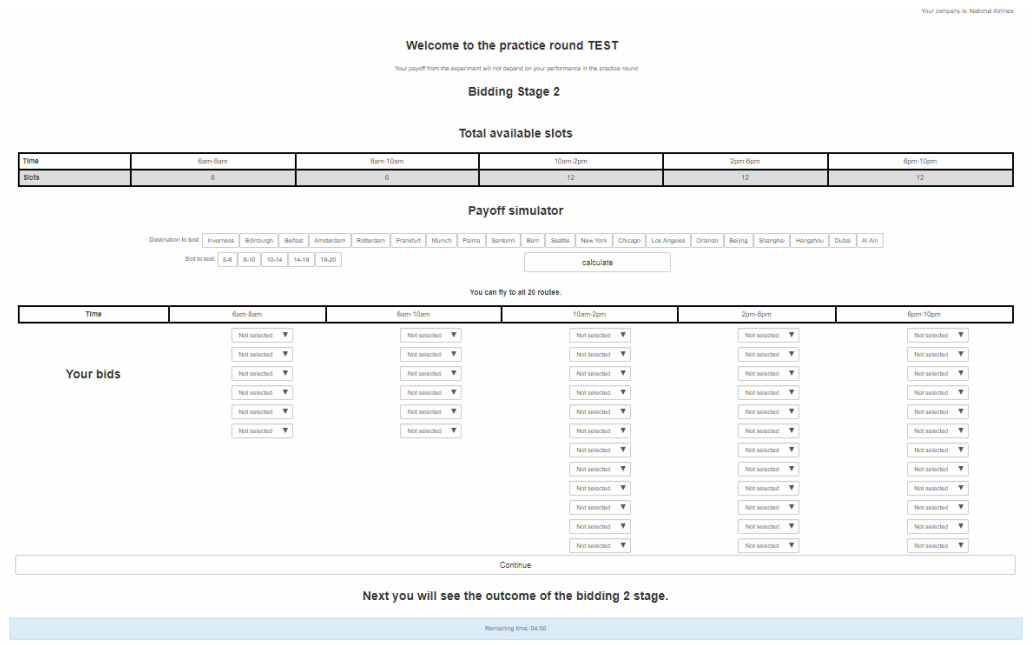

A2.2.1 Key information sheet

Figure A2.30 Key information sheet, treatment 2, National Airlines (page 1)

Key information

Welcome to the experiment! This sheet summarises key information that is important for the experiment so you always have it at hand. Full instructions are displayed on your screen.

Your airline

National Airlines

Your strategy

National Airlines uses Heathrow as its 'hub airport', flying from Heathrow all over the world.

Your budget

You have a budget of [✂] each round. Any budget you have left unspent will contribute to your payoff.

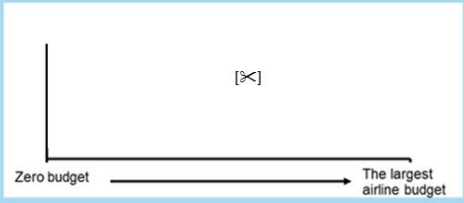

Routes

You can fly all 20 routes.

These will be shown to you when you come to make decisions.

New entrant status

You do not qualify for new entrant slots.

We will explain what this means.

Yourpayoff

Your payoff = **Base element** + **Special element**

The base element includes:

- **Slot time**.
- **Choice of route** (i.e. which city you fly to with the slot).

The special element is different for each airline. Yours includes:

- **Competition**. If other airlines choose to use slots on the same route as you, at the same time, then your payoff on that slot reduces by 15%.
- **Volume**. Each slot you have on the same route at the same time reduces your payoff on all those slots by 5%.
- **Final ratio**. If you achieve a final ratio of short haul to long haul flights between [✂] (i.e. a ratio of [✂]), then your payoff increases by 25% on all slots. Short haul flights are to the rest of the UK and Europe. Long haul flights are to Asia, the Middle East, and the USA.
- **Maximum value**. After obtaining [✂], each additional slot holds no value (zero payoff). But you may choose to hold more slots than this as you can trade them with other airlines.

Figure A2.31 Key information sheet, treatment 2, National Airlines (page 2)

The base element of your payoff

This chart and table show you the base element of your payoffs (i.e. the impact of choosing different routes and times). It does not show the impact of special elements of your payoffs.

The chart and table show the same information.

The 8am-10am and 10am-2pm time buckets have the same payoff for you (and are shown using the same symbol on the chart).

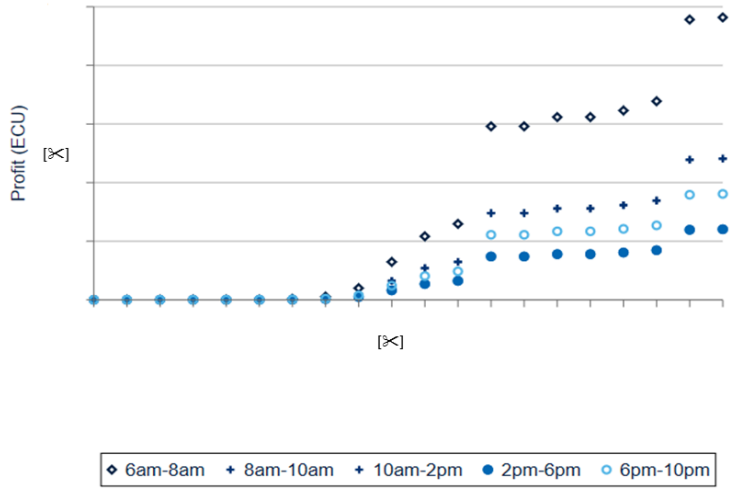

[✂]

Each round

Each round contains a number of stages, where you will bid for slots, trade slots between airlines, and decide which route to fly with the slots you win.

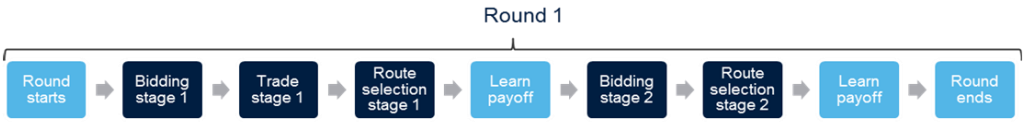

Figure A2.32 Key information sheet, treatment 2, National Airlines (page 3)

The computer

The computer decides which airline to give a slot to based on the route the airlines say they will fly.

The rules that the computer follows are partially known by the airlines:

- The computer gives priority to new routes from Heathrow (routes where no airline currently flies).
- New routes to Asia are given higher priority than new routes to Europe. New routes to Europe are given higher priority than new routes to the USA.
- Priority is given to routes your airline is not already flying to (but other airlines may fly to).
- Priority is given to the smallest airline on a route (e.g. if Airline A has 1 slot on that route, and Airline B has 10 slots on that route, Airline A would have priority).

The table below tells you: the new routes from Heathrow; which routes are new for your airline; and where you are the smallest airline on a route. New entrants are airlines with no current slots at Heathrow (Speedy Flights, Quick Wings, and East Asia Airways).

These slots below are not relevant for your payoffs (they do not impact the base or special element of your payoffs in any way). The table is only useful to help you decide on which routes to fly when making your bids to the computer.

Number of slots to each route from Heathrow, by airline

Route	Country	New route from Heathrow?	**National Airways**	**Jet Flights**	**Atlantic Airlines**	**North American Airways**	**Speedy Flights**	**Quick Wings**	**Euro Airlines**	**Royal Flights**	**West Asia Airways**	**East Asia Airways**
Inverness	UK		1									
Edinburgh	UK		10									
Belfast	UK		4									
Amsterdam	The Netherlands		8							10		
Rotterdam	The Netherlands	*Yes*										
Frankfurt	Germany		6						12			
Munich	Germany		6						8			
Palma	Spain		2									
Santorini	Greece		1									
Bern	Switzerland	*Yes*										
Seattle	The USA		2	1								
New York	The USA		10	7	4	5						
Chicago	The USA		2		4	3						
Los Angeles	The USA		3	3	2	1						
Orlando	The USA	*Yes*										
Beijing	China		1									
Shanghai	China		1	1								
Hangzhou	China	*Yes*										
Dubai	United Arab Emirates		3	1							6	
Al Ain	United Arab Emirates	*Yes*										

A2.3 Treatment 3: no new entrant rule

This section presents screenshots for the screens that differ between Treatment 3 and Treatment 1.

Figure A2.33 Slots

Slots

In order for an airline to fly to and from an airport, it must hold runway slots.

Each slot gives the airline the right to use the airport (for a departure or arrival) at a given time.

Slots are awarded to airlines in pairs, so that airlines can arrive and depart from an airport.

Slots are allocated to airlines by the computer. We will explain how the computer makes this decision.

While some of the airlines already fly from Heathrow, this study is about how the new slots will be used.

There are 48 new slots in total, at five times of day. This will be clear to you when you come to make decisions.

	6am-8am	8am-10am	10am-2pm	2pm-6pm	6pm-10pm
slots	6	6	12	12	12

→ 48 slots in total

Airlines

There are 10 airlines, each played by one participant. Each airline has a different budget and earns different payoffs.

There are 20 possible routes that airlines can fly from Heathrow, but some airlines cannot fly to all 20 routes. You will be told where you can fly.

You are: **National Airlines**

- **National Airlines**, which uses Heathrow as its 'hub airport', flying from Heathrow all over the world.
- **Jet Flights**, which aims to make Heathrow its 'hub airport', flying from Heathrow all over the world.
- **Atlantic Airlines**, a USA-based airline, flying from Heathrow to the USA.
- **North American Airways**, a USA-based airline, flying from Heathrow to the USA.
- **Quick Wings**, a low-cost airline, flying from Heathrow to the USA.
- **Speedy Flights**, a low-cost airline, flying from Heathrow to other parts of the UK and Europe.
- **Euro Airlines**, a German airline, flying from Heathrow to Germany.
- **Royal Flights**, a Dutch airline, flying from Heathrow to the Netherlands.
- **West Asia Airways**, a United Arab Emirates airline, flying from Heathrow to the United Arab Emirates .
- **East Asia Airways**, a Chinese airline, flying from Heathrow to China.

Please write the number of the computer you are:

Continue

Remaining time: 02:59

Figure A2.34 Your airline

Your company is: National Airlines

National Airlines

Your strategy

National Airlines uses Heathrow as its 'hub airport', flying from Heathrow all over the world.

Your budget

You have a budget of [✂] each round. Any budget you have left unspent will contribute to your payoff.

[✂]

Zero budget → The largest airline budget

Routes

You can fly all 20 routes.

These will be shown to you when you come to make decisions.

New entrant status

You do not qualify for new entrant slots.

We will explain what this means.

This information is also in your paper instructions.

Continue

Remaining time: 00:56

Figure A2.35 The computer

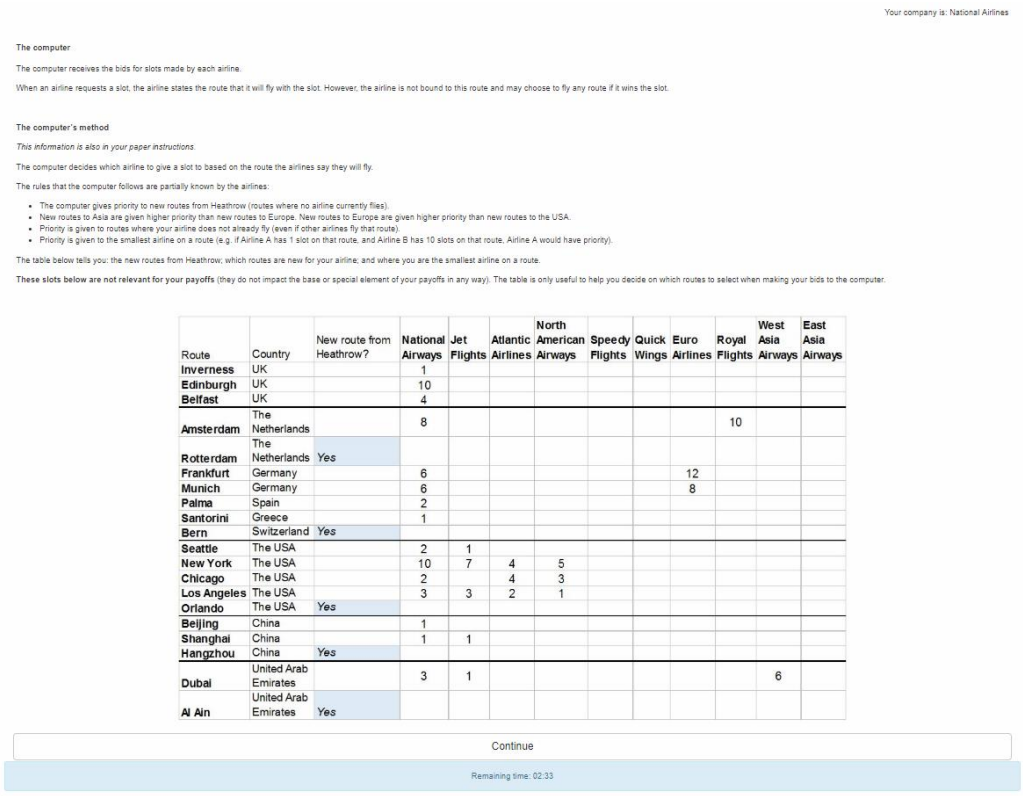

Your company is: National Airlines

The computer

The computer receives the bids for slots made by each airline.

When an airline requests a slot, the airline states the route that it will fly with the slot. However, the airline is not bound to this route and may choose to fly any route if it wins the slot.

The computer's method

This information is also in your paper instructions.

The computer decides which airline to give a slot to based on the route the airlines say they will fly.

The rules that the computer follows are partially known by the airlines:

- The computer gives priority to new routes from Heathrow (routes where no airline currently flies).
- New routes to Asia are given higher priority than new routes to Europe. New routes to Europe are given higher priority than new routes to the USA.
- Priority is given to routes where your airline does not already fly (even if other airlines fly that route).
- Priority is given to the smallest airline on a route (e.g. if Airline A has 1 slot on that route, and Airline B has 10 slots on that route, Airline A would have priority).

The table below tells you: the new routes from Heathrow; which routes are new for your airline; and where you are the smallest airline on a route.

These slots below are not relevant for your payoffs (they do not impact the base or special element of your payoffs in any way). The table is only useful to help you decide on which routes to select when making your bids to the computer.

Route	Country	New route from Heathrow?	National Airways	Jet Flights	Atlantic Airlines	North American Airways	Speedy Flights	Quick Wings	Euro Airlines	Royal Flights	West Asia Airways	East Asia Airways
Inverness	UK		1									
Edinburgh	UK		10									
Belfast	UK		4									
Amsterdam	The Netherlands		8							10		
Rotterdam	The Netherlands	*Yes*										
Frankfurt	Germany		6						12			
Munich	Germany		6						8			
Palma	Spain		2									
Santorini	Greece		1									
Bern	Switzerland	*Yes*										
Seattle	The USA		2	1								
New York	The USA		10	7	4	5						
Chicago	The USA		2		4	3						
Los Angeles	The USA		3	3	2	1						
Orlando	The USA	*Yes*										
Beijing	China		1									
Shanghai	China		1	1								
Hangzhou	China	*Yes*										
Dubai	United Arab Emirates		3	1							6	
Al Ain	United Arab Emirates	*Yes*										

Continue

Remaining time: 02:33

Figure A2.36 Test questions

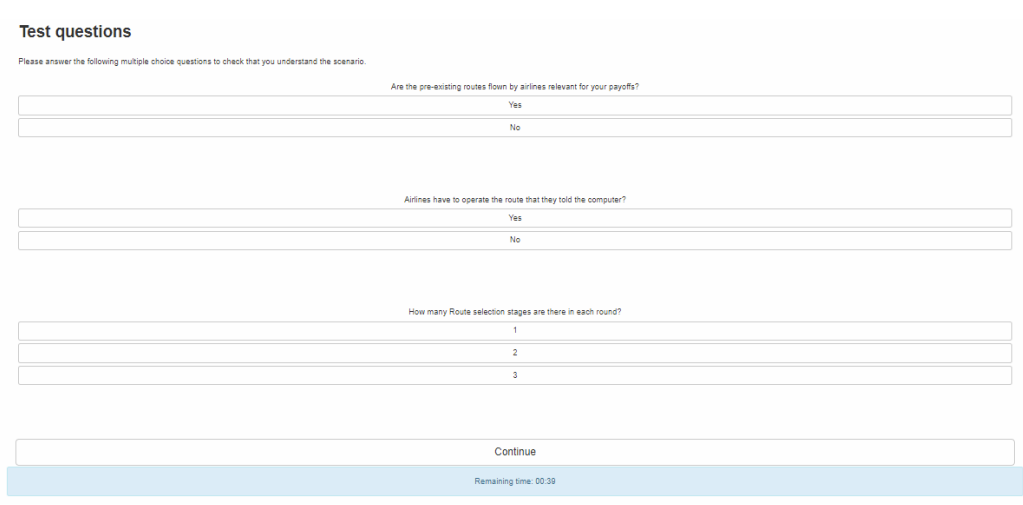

Test questions

Please answer the following multiple choice questions to check that you understand the scenario.

Are the pre-existing routes flown by airlines relevant for your payoffs?

Yes

No

Airlines have to operate the route that they told the computer?

Yes

No

How many Route selection stages are there in each round?

1

2

3

Continue

Remaining time: 00:39

Figure A2.37 Trade stage 1

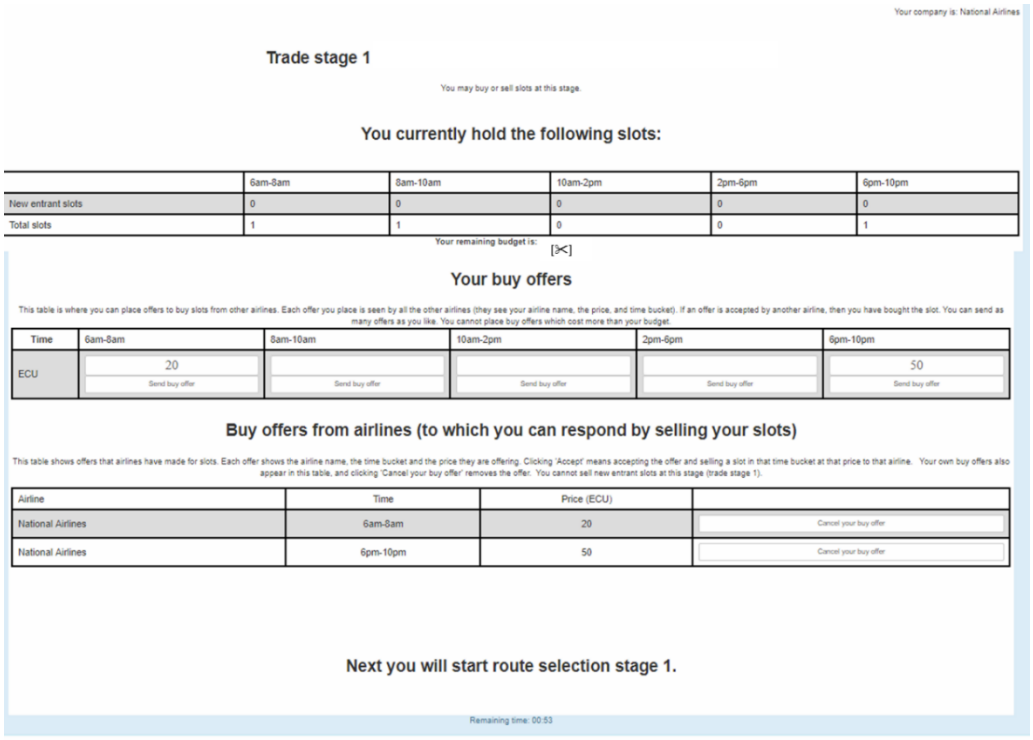

A2.3.1 Key information sheet

Figure A2.38 Key information sheet, treatment 3, National Airlines (page 1)

Key information

Welcome to the experiment! This sheet summarises key information that is important for the experiment so you always have it at hand. Full instructions are displayed on your screen.

Your airline

National Airlines

Your strategy

National Airlines uses Heathrow as its 'hub airport', flying from Heathrow all over the world.

Routes

You can fly all 20 routes.

These will be shown to you when you come to make decisions.

Your budget

You have a budget of [✂] each round. Any budget you have left unspent will contribute to your payoff.

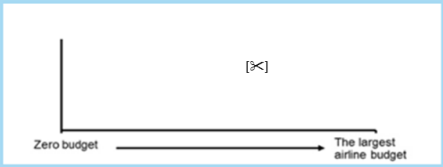

Yourpayoff

Your payoff = **Base element** + **Special element**

The base element includes:

- **Slot time**.
- **Choice of route** (i.e. which city you fly to with the slot).

The special element is different for each airline. Yours includes:

- **Competition**. If other airlines choose to use slots on the same route as you, at the same time, ther your payoff on that slot reduces by 15%.
- **Volume**. Each slot you have on the same route at the same time reduces your payoff on all those slots by 5%.
- **Final ratio**. If you achieve a final ratio of short haul to long haul flights between [✂] (i.e. a ratio of [✂] then your payoff increases by 25% on all slots. Short haul flights are to the rest of the UK and Europe. Long haul flights are to Asia, the Middle East, and the USA.
- **Maximum value**. After obtaining [✂], each additional slot holds no value (zero payoff). But you may choose to hold more slots than this as you can trade them with other airlines.

Figure A2.39 Key information sheet, treatment 3, National Airlines (page 2)

The base element of your payoff

This chart and table show you the base element of your payoffs (i.e. the impact of choosing different routes and times). It does not show the impact of special elements of your payoffs.

The chart and table show the same information.

The 8am-10am and 10am-2pm time buckets have the same payoff for you (and are shown using the same symbol on the chart).

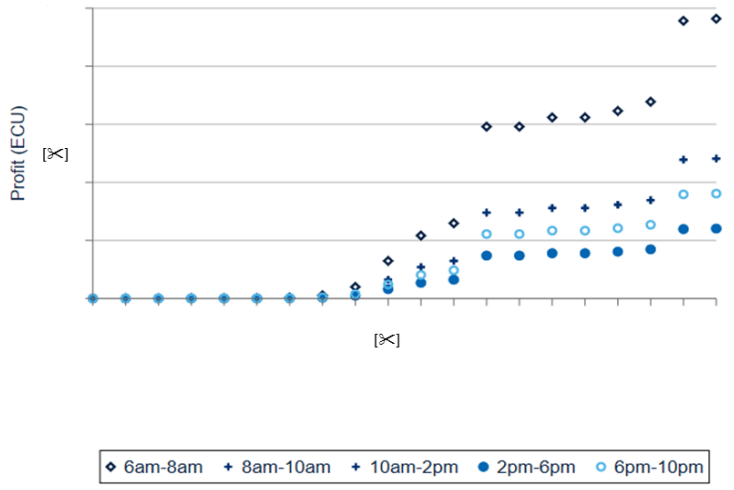

[✂]

Each round

Each round contains a number of stages, where you will bid for slots, trade slots between airlines, and decide which route to fly with the slots you win.

Figure A2.40 Key information sheet, treatment 3, National Airlines (page 3)

The computer

The computer decides which airline to give a slot to based on the route the airlines say they will fly.

The rules that the computer follows are partially known by the airlines:

- The computer gives priority to new routes from Heathrow (routes where no airline currently flies).
- New routes to Asia are given higher priority than new routes to Europe. New routes to Europe are given higher priority than new routes to the USA.
- Priority is given to routes your airline is not already flying to (but other airlines may fly to).
- Priority is given to the smallest airline on a route (e.g. if Airline A has 1 slot on that route, and Airline B has 10 slots on that route, Airline A would have priority).

The table below tells you: the new routes from Heathrow; which routes are new for your airline; and where you are the smallest airline on a route.

These slots below are not relevant for your payoffs (they do not impact the base or special element of your payoffs in any way). The table is only useful to help you decide on which routes to fly when making your bids to the computer.

Number of slots to each route from Heathrow, by airline

Route	Country	New route from Heathrow?	**National Airways**	**Jet Flights**	**Atlantic Airlines**	**North American Airways**	**Speedy Flights**	**Quick Wings**	**Euro Airlines**	**Royal Flights**	**West Asia Airways**	**East Asia Airways**
Inverness	UK		1									
Edinburgh	UK		10									
Belfast	UK		4									
Amsterdam	The Netherlands		8							10		
Rotterdam	The Netherlands	*Yes*										
Frankfurt	Germany		6						12			
Munich	Germany		6						8			
Palma	Spain		2									
Santorini	Greece		1									
Bern	Switzerland	*Yes*										
Seattle	The USA		2	1								
New York	The USA		10	7	4	5						
Chicago	The USA		2		4	3						
Los Angeles	The USA		3	3	2	1						
Orlando	The USA	*Yes*										
Beijing	China		1									
Shanghai	China		1	1								
Hangzhou	China	*Yes*										
Dubai	United Arab Emirates		3	1							6	
Al Ain	United Arab Emirates	*Yes*										

A2.4 Treatment 4: auction

This section presents screenshots for the screens that differ between Treatment 4 and Treatment 1.

Figure A2.41 Rounds

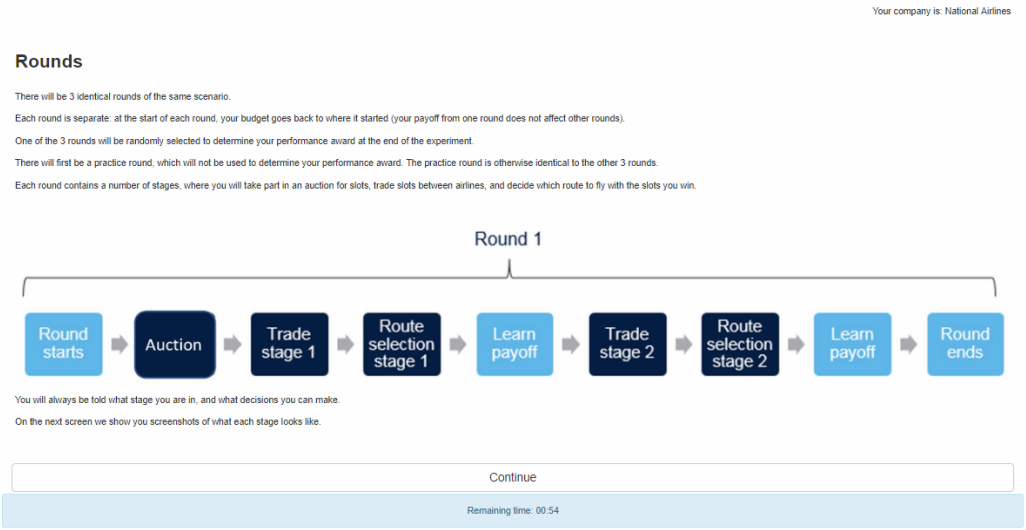

Figure A2.42 The computer (part 1)

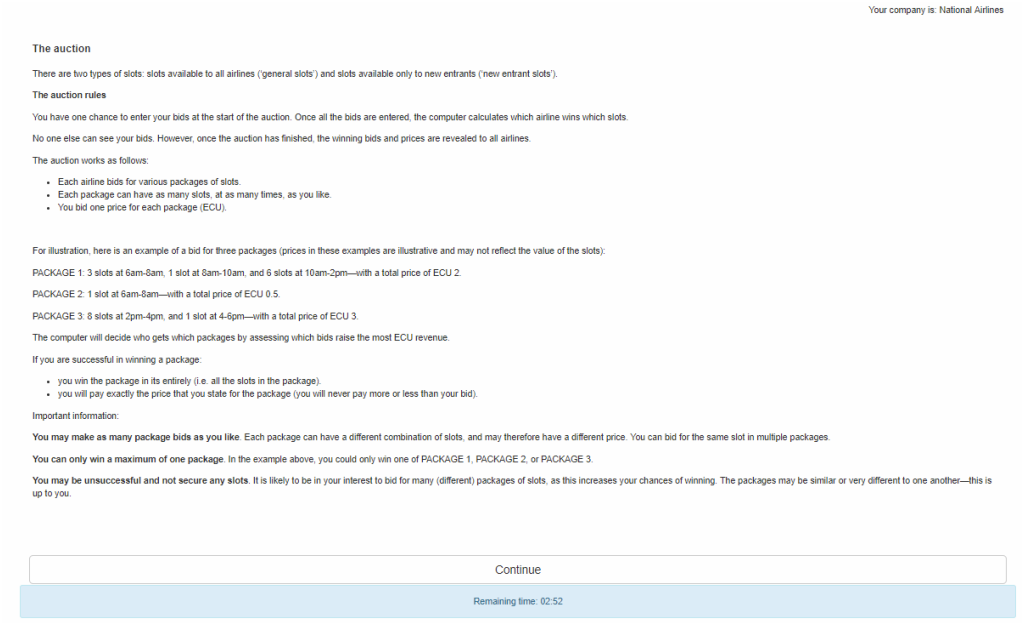

Figure A2.43 The computer (part 2)

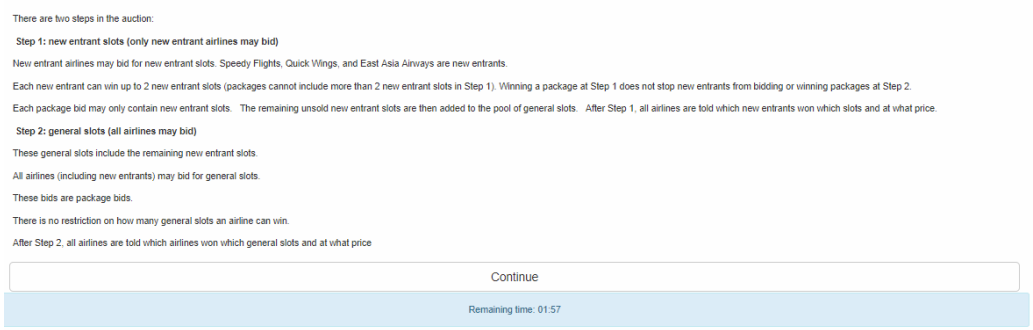

Figure A2.44 Test questions

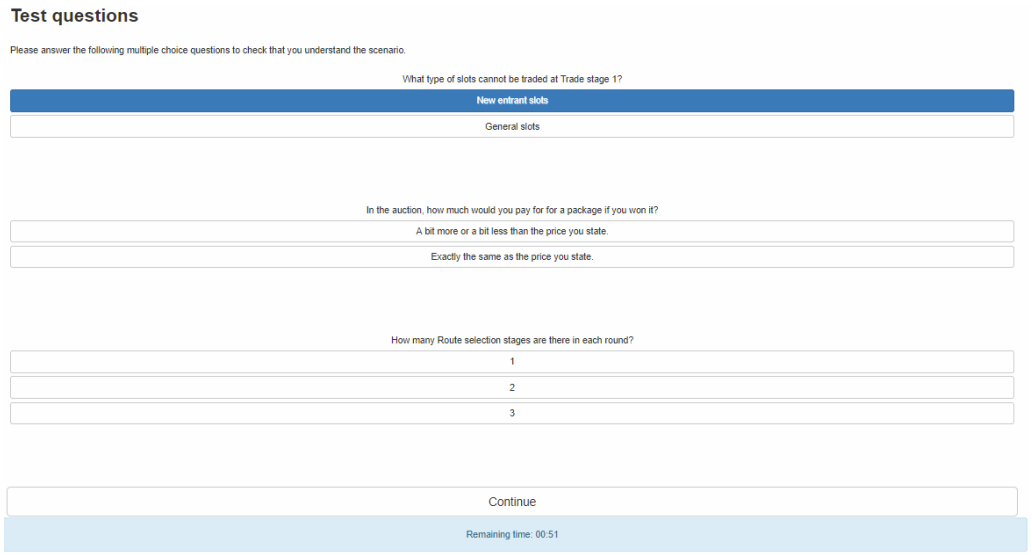

Figure A2.45 Bidding for new entrant slots

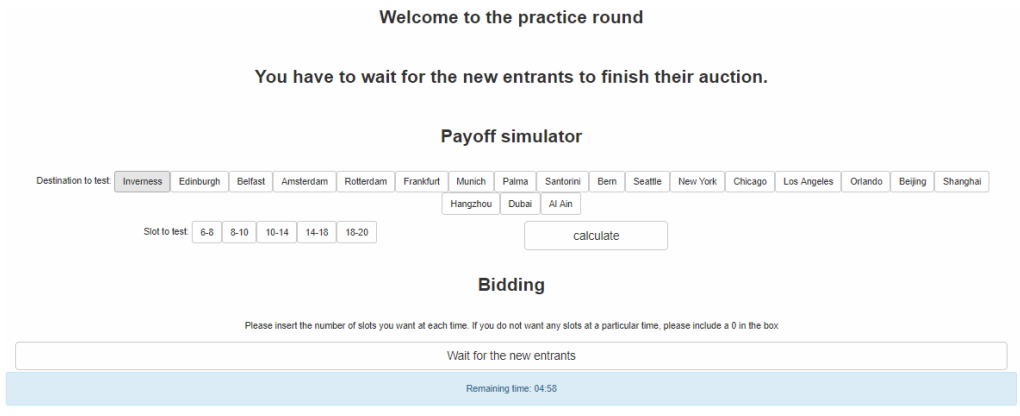

Figure A2.46 Results from bidding for new entrant slots

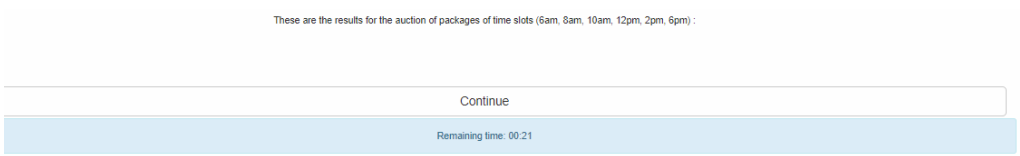

Figure A2.47 Bidding for general slots

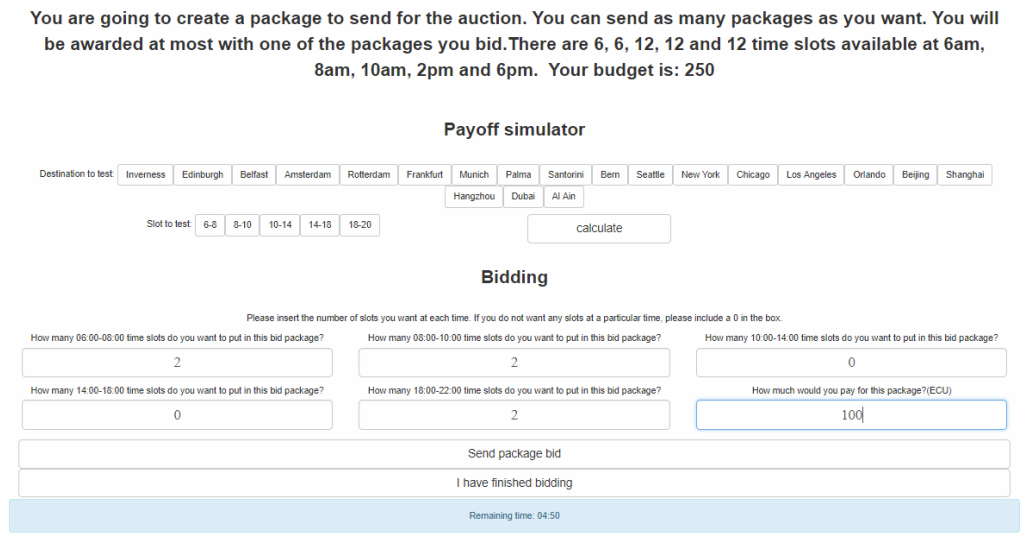

Figure A2.48 Results from bidding for general slots

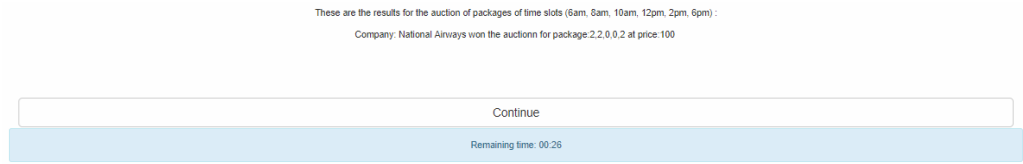

A2.4.1 Key information sheet

Figure A2.49 Key information sheet, treatment 4, National Airlines (page 1)

Key information

Welcome to the experiment! This sheet summarises key information that is important for the experiment so you always have it at hand. Full instructions are displayed on your screen.

Your airline

National Airlines

Your budget

You have a budget of [✂] each round. Any budget you have left unspent will contribute to your payoff.

[✂]

Zero budget → The largest airline budget

Your strategy

National Airlines uses Heathrow as its 'hub airport', flying from Heathrow all over the world.

Routes

You can fly all 20 routes.

These will be shown to you when you come to make decisions.

New entrant status

You do not qualify for new entrant slots.

We will explain what this means.

Yourpayoff

Your payoff = **Base element** + **Special element**

The base element includes:

- **Slot time**.
- **Choice of route** (i.e. which city you fly to with the slot).

The special element is different for each airline. Yours includes:

- **Competition**. If other airlines choose to use slots on the same route as you, at the same time, then your payoff on that slot reduces by 15%.
- **Volume**. Each slot you have on the same route at the same time reduces your payoff on all those slots by 5%.
- **Final ratio**. If you achieve a final ratio of short haul to long haul flights between [✂] (i.e. a ratio of [✂]), then your payoff increases by 25% on all slots. Short haul flights are to the rest of the UK and Europe. Long haul flights are to Asia, the Middle East, and the USA.
- **Maximum value**. After obtaining [✂], each additional slot holds no value (zero payoff). But you may choose to hold more slots than this as you can trade them with other airlines.

Figure A2.50 Key information sheet, treatment 4, National Airlines (page 2)

The base element of your payoff

This chart and table show you the base element of your payoffs (i.e. the impact of choosing different routes and times). It does not show the impact of special elements of your payoffs.

The chart and table show the same information.

The 8am-10am and 10am-2pm time buckets have the same payoff for you (and are shown using the same symbol on the chart).

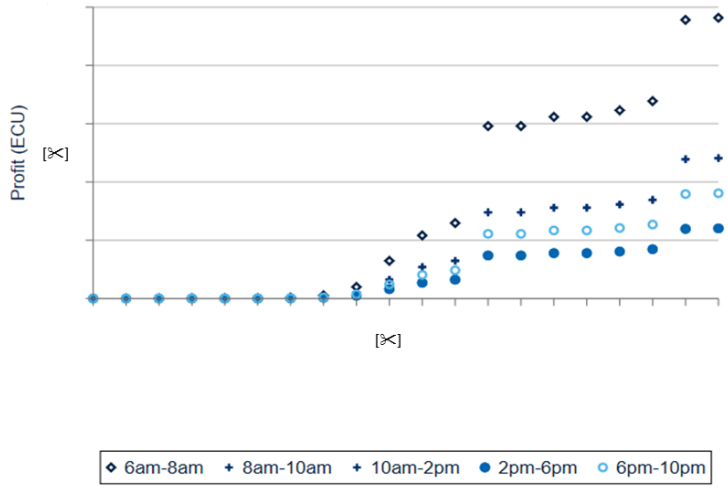

[✂]

Each round

Each round contains a number of stages, where you will bid for slots, trade slots between airlines, and decide which route to fly with the slots you win.

Figure A2.51 Key information sheet, treatment 4, National Airlines (page 3)

The auction

There are two types of slots: slots available to all airlines ('general slots') and slots available only to new entrants ('new entrant slots').

The auction rules

You have one chance to enter your bids at the start of the auction. Once all the bids are entered, the computer calculates which airline wins which slots.

No one else can see your bids. However, once the auction has finished, the winning bids and prices are revealed to all airlines.

The auction works as follows:

- Each airline bids for various packages of slots.
- Each package can have as many slots, at as many times, as you like.
- You bid one price for each package (ECU).

For illustration, here is an example of a bid for three packages (*prices in these examples are illustrative and may not reflect the value of the slots)*:

- PACKAGE 1: 3 slots at 6am-8am, 1 slot at 8am-10am, and 6 slots at 10am-2pm—with a total price of ECU 2.
- PACKAGE 2: 1 slot at 6am-8am—with a total price of ECU 0.5.
- PACKAGE 3: 8 slots at 2pm-4pm, and 1 slot at 4-6pm—with a total price of ECU 3.

The computer will decide who gets which packages by assessing which bids raise the most ECU revenue.

If you are successful in winning a package:

- you win the package in its entirely (i.e. all the slots in the package).
- you will pay exactly the price that you state for the package (you will never pay more or less than your bid).

Important information:

- **You may make as many package bids as you like.** Each package can have a different combination of slots, and may therefore have a different price. You can be for the same slot in multiple packages.
- **You can only win a maximum of one package**. In the example above, you could only win one of PACKAGE 1, PACKAGE 2, or PACKAGE 3.
- **You may be unsuccessful and not secure any slots**. It is likely to be in your interest to bid for many (different) packages of slots, as this increases your chances of winning. The packages may be similar or very different to one another—this is up to you.

Figure A2.52 Key information sheet, treatment 4, National Airlines (page 4)

There are two steps in the auction:

Step 1: new entrant slots (only new entrant airlines may bid)

- New entrant airlines may bid for new entrant slots. Speedy Flights, Quick Wings, and East Asia Airways are new entrants.
- Each new entrant can win up to 2 new entrant slots (packages cannot include more than 2 new entrant slots in Step 1). Winning a package at Step 1 does not stop new entrants from bidding or winning packages at Step 2.
- Each package bid may only contain new entrant slots.
- The remaining unsold new entrant slots are then added to the pool of general slots.
- After Step 1, all airlines are told which new entrants won which slots and at what price.

Step 2: general slots (all airlines may bid)

- These general slots include the remaining new entrant slots.
- All airlines (including new entrants) may bid for general slots.
- These bids are package bids.
- There is no restriction on how many general slots an airline can win.
- After Step 2, all airlines are told which airlines won which general slots and at what price.

A3 Literature review

A3.1 Review of the literature on slot allocation mechanisms

In this review we explore airport slot allocation mechanisms that have been considered in the literature. A number of these sources focus on how an auction would be applied in this context.

It is clear that any auction of slots would need to be for packages of slots, as opposed to (solely) for single departures or arrivals or even single slot pairs. As we discuss below, this is due to the complementarity of desired slots for each airline in operating its particular business model. This need for a 'combinatorial' auction with 'package bidding' is also reflected in the theoretical, experimental and practical literature that we have reviewed, as follows.

- Rassenti et al. (1982)—in a seminal laboratory experiment with paid participants, a sealed-bid combinatorial auction was developed for the allocation of airport slots to competing airlines. Airlines submitted their various contingent package bids for slots, with winning bids determined through an algorithm. Secondary trading was also permitted.[117]

- Ball et al. (2007)—in the run-up to the potential introduction of slot reform to relieve congestion at LaGuardia airport in New York, the authors discuss the design of two experimental simulations undertaken with industry professionals. The first simulation allowed participants to solve congestion using various administrative measures and congestion pricing. The second simulation used a combinatorial slot auction. Based on their research, the authors recommended using a second-price ascending clock combinatorial auction (with an activity rule to discourage snipe bidding). Secondary trading would also be permitted.[118]

- Steer Davies Gleave (2011)—in a report for the European Commission, the authors discuss the eventual auction design proposed by the Federal Aviation Administration (FAA) for La Guardia (and two other New York airports).[119] In the end, the proposals for an auction were not implemented.

    - Auctions were to cover existing and new airport capacity. During the auction, bidders would need to submit up to 2,000 mutually exclusive bids at each of the three airports.

    - A single-round sealed bid auction was to be used. This meant that, in one confidential submission (communication between bidders was prohibited), bidders would need to fully record their preferences for all potential packages of interest.

    - Based on the bids submitted, an algorithm would select winning bidders according to (feasible) packages that would maximise auctioneer revenue.

[117] Rassenti, S.J., Smith, V.L. and Bulfin, R.L. (1982), 'A Combinatorial Auction Mechanism for Airport Time Slot Allocation', *The Bell Journal of Economics*, **13**:2, pp. 402–417.
[118] Ball, M.O., Ausubel, L.M., Berardino, F., Cramton, P., Donohue, G., Hansen, M. and Hoffman, K. (2007), 'Market-Based Alternatives for Managing Congestion at New York's LaGuardia Airport', article presented at Airneth Annual Conference, April, Digital Repository at the University of Maryland.
[119] Steer Davies Gleave (2011), 'Impact assessment of revisions to Regulation 95/93', March.

- However, payments made by winning bidders would be determined using a second-placed price approach. In most instances, this would mean that winning bidders would pay less than their actual bid.

- The use of the second-placed price sealed-bid combinatorial auction was aimed at encouraging truthful bidding.

- Pertuiset and Santos (2014)—in this theoretical (as opposed to experimental) paper the authors propose a sealed-bid second-price combinatorial auction for a portion (10%) of existing slots across Europe. They state that this would be efficient, incentive-compatible, understandable and implementable. The authors envisage that 10% of slots would be auctioned per year (and that the current administrative system would be phased out within ten years).[120]

- Herranz et al. (2016)—in a pan-European research project, the authors summarise their simulation research of alternative slot allocation mechanisms. This employ agent-based modelling (ABM). In contrast to a laboratory experiment, ABM is a computational-based approach that simulates the actions and interactions of autonomous agents.[121]

  - The ABM platform allows for the assessment of various airport slot allocation mechanisms in different scenarios. This includes primary allocation based on administrative procedures and various forms of combinatorial auction—ascending, descending and 'Walrasian'. It also allows for alternative secondary trading arrangements.

  - The 'Walrasian' auction format (also known as a combinatorial price-setting iterative auction) is the main area of focus. Here, the auctioneer first announces prices. Given the package bids obtained, and using an algorithm, the auctioneer then modifies prices iteratively as a function of demand and supply.

  - The results of the ABM experiments using the combinatorial price-setting auction are compared with those obtained using the current administrative slot allocation mechanism. The authors conclude that the auction approach leads to a more efficient use of airport capacity, increasing the number of flights, the number of passengers, and total social welfare.

- Araúzo et al. (2018)—the authors use an ABM to compare the results of their proposed iterative combinatorial price-setting auction design with the administrative approach. They conclude that the auction design can provide a way to solve the airport allocation problem in an economically optimal or best-possible way.[122]

[120] Pertuiset, T. and Santos, G. (2014), 'Primary auction of slots at European Airports', *Research in Transportation Economics*, 45, pp. 66–71.
[121] Herranz, R., Toribo, D., Alsina, N., Garrigó, L., Poza, D., Pessenti, R. and Castelli, L. (2016), 'Agent-Based Simulation of Airport Slot Allocation Mechanisms: Analysis of Results', ACCESS (AppliCation of agent-based Computational Economics to Strategic Slot Allocation), Working Paper 6, February.
[122] Araúzo, J.A., Villafáñez, F.A., García, D.P., Pajares, J. and Pavón, J. (2018), 'Agent Based Modelling and Simulation of an Auction Market for Airport Slots Allocation', in Bajo, J. et al. (eds), *Highlights of Practical Applications of Agents, Multi-Agent Systems, and Complexity: The PAAMS Collection*, PAAMS 2018, Communications in Computer and Information Science, 887.

A3.2 Review of the literature on auction design

A3.2.1 Why auctions?

Auctions are applicable to situations in which there is a contested resource—i.e. where demand exceeds supply—and are an alternative to an administrative approach to resource allocation.

Auctions involve competition between bidders to determine who wins items and how much they pay for them. There is an extensive theoretical literature on auctions, and on their advantages over administrative ('direct allocation') approaches.

- Asymmetric information—an allocator may not know the value of individual lots to users or the value of different packages of lots, whereas auctions reveal 'private information'.
- Common value uncertainty—if no one bidder knows the true value of an item due to uncertainty over future market conditions, the auction process can lead to a process of price discovery.
- Allocative efficiency—total societal welfare is maximised, as lots are assigned to the highest-value user.
- Dynamic efficiency—auctions may act as a portal for entrants to gain a foothold in a market, delivering efficiency benefits over time.

However, whether positive outcomes will be obtained from an auction depends on the sector concerned and many interrelated factors. If an auction is chosen, its precise design can be important. Three issues are worth bearing in mind in the context of the current laboratory experiment.

- There is a wealth of theoretical literature on auction design, which also indicates in which circumstances one auction design might be preferable over another. This guided our experiment design.
- While there is much evidence on spectrum auctions and their appropriate design, there is less evidence of appropriate airport slot auction design (although we reviewed the available evidence above).
- Any auction design needs to be well understood and practical—both in the real world and within an experimental setting.

On this last point, what the experimental environment needs to do is mimic essential features of a real-world design that theory and the evidence base indicate would be appropriate, while taking into account the practical limitations of the experimental environment and the need for participants to quickly understand the mechanism. In the experiment, our chosen design may be one that is used in a real-world setting, but equally an alternative might be adopted in practice.

A3.2.2 Setting the scene: single unit auctions

The idea behind auction design is that the auctioneer maximises revenue through allocating units to bidders with the highest willingness to pay, subject to certain constraints. This revenue maximisation in the current context (airport slots) is a means to an end, rather than an end in itself.

The choice of auction design is simpler when there are single units for sale (that are independent from other lots from the buyers' perspective)—for example, the sale of a stand-alone piece of land at a real estate auction. In this case the auctioneer will need to choose between a sealed-bid (one-shot) format (in which rival bidders are not allowed to communicate with one another) and an open-auction format (such as an ascending clock auction).

The open format has advantages where bidders face 'common value uncertainty'. Here, there are common unknowns across bidders—such as the future development potential of a piece of land. Given this uncertainty, bidders may be wary of bidding too much for an item—known as the 'winner's curse'—which may in turn hamper competition. The multiple rounds of the ascending auction mean that bidders can iteratively refine their valuations for the item based on the observed behaviour of others—a process known as 'price discovery'.

However, this openness comes with a price. Open auctions may be prone to 'snipe bidding'—where certain bidders hide their truthful valuations and bid only at a very late stage in the process. This impedes the price discovery process. Minimum bidding 'activity rules' throughout the auction may then be required to counter sniping behaviour, but this can add complexity.

The sealed-bid format is simpler to administer than a multi-round auction. Given its one-shot nature, and the lack of communication between bidders, it also eliminates snipe bidding. In a sealed-bid auction, 'winner determination' is based on the highest bid submitted. Alternative 'payment rules' are then also possible within the sealed-bid format. Under a 'first-price' rule the winning bidder pays what he/she bid. Under a 'second-price' rule the winning bidder pays a price determined by the losing bids (here, the second-highest price).

For 'incentive compatibility'—or to obtain truthful bidding in advance of winner determination—a second-price rule is generally preferred. This sealed-bid second-price approach is known as a 'Vickery auction'.[123]

In contrast to the sealed-bid approach, a clock auction is a type of open auction in which the price rises continuously while bidders gradually quit the auction. However, for single-unit auctions the clock auction has the same outcome as the Vickery (second-price sealed-bid) auction—in terms of the winning bidder and the price paid.

In the single unit case, therefore, the choice between a one-shot sealed-bid auction and an open multiple-round ascending auction may simply come down to an assessment of the extent to which there is common value uncertainty and practical considerations around strategic bidding behaviour (such as sniping).

A3.2.3 Combinatorial auctions: design issues

However, as has been noted at the beginning of this section, there are both complementarities and substitutability between lots.

- Complementarities—the value placed on two individual slots, 1 and 2, by an airline may be greater than the value it places individually on the two slots, due to its business model. This property is called valuation 'super-additivity'.

[123] Vickrey, W. (1961), 'Counterspeculation, Auctions, and Competitive Sealed Tenders', *The Journal of Finance*, **16**:1, pp. 8–37.

- Substitutability—airlines may wish to win Package A of slots (its preferred bundle 1+2+3) or Package B (as a second best 4+5), but not both.

The complementary valuations across individual slots mean that, under a single-unit format, airlines will be subject to the 'exposure problem' (Rothkopf et al. (1998)).[124] Bidders face the risk that they may end up winning only a subset of their desired slots and that they may end up paying too much for them. The consequence is that airlines will tend to undervalue individual slots in anticipation of this.

Combinatorial auctions, in which bidders can express preferences for packages of slots, can help to solve the exposure problem. In addition, 'bidding languages' can be adopted within combinatorial auctions, such that the substitutability of valuations between (say) first- and second-best packages can be taken into account.

Therefore, combinatorial auctions are appropriate for airline slot auctions , whereby bidders can bid for packages of slots to more fully express their underlying preferences, and bidders are shielded from the exposure problem.

Over recent years there has been a growing literature on the theory and practice of combinatorial auction design. This has been summarised by Cramton et al. (2006)[125] and more recently by Bichler (2018).[126]

Perhaps unsurprisingly, the issues around auction design are more complex than in the single-unit case. Our key design decisions are as follows:

- auction format—ascending clock (open-bid) versus one-off sealed-bid;
- packaging—including fixed versus flexible packages;
- bidding language—including taking account of both complementarity and substitutability;
- winner determination—using an appropriate approximate algorithm;
- payments—first-price versus second-price rule.

In terms of the appropriate auction format, our review of the airport slot literature (see above) indicates the following two general alternatives.

- Ascending clock combinatorial auction—this is akin to that proposed by Ball et al. (2007) for airport slots in the USA. Prices are declared by the auctioneer for individual slots. Package bids can be made by bidders for combinations of slots. Prices increase incrementally. An activity rule seeks to limit bid sniping.
- Second-price sealed-bid combinatorial auction—bidders express, in one shot, their full preferences over all feasible bundles that suit their needs. This is known as a Vickrey–Clarke–Groves (VCG) auction and is an extension of the unit-based Vickery auction.[127] It was used by Rassenti et al.

[124] Rothkopf, M., Pekeč, A. and Ronald, M. (1998), 'Computationally Manageable Combinational Auctions', *Management Science*, **44**:8, pp. 1131–1147.
[125] Cramton, P., Shoham, Y. and Steinberg, R. (eds) (2006), *Combinatorial Auctions*, MIT Press.
[126] Bichler, M. (2018), *A Linear Programming Approach to Auctions and Matching*, Cambridge University Press.
[127] Vickrey, W. (1961), 'Counterspeculation, Auctions, and Competitive Sealed Tenders', *The Journal of Finance*, **16**:1, pp. 8–37; Clarke, E. (1971), 'Multipart Pricing of Public Goods', *Public Choice*, **11**:1, pp. 17–33; Groves, T. (1973), 'Incentives in Teams', *Econometrica*, **41**:4, pp. 617–631.

(1982) and was also the method initially favoured by the FAA for slot auctions in the USA.

Within our experimental environment the sealed-bid method has the advantage of taking less time in the laboratory—it would involve the submission of all requests by each airline in one shot. An ascending-clock multi-round approach would be more time-consuming to administer. The sealed-bid method would also eliminate any potential bid-sniping issues that could be present in an open auction. Finally, in the case of airlines, common value uncertainty may be less relevant as each airline will have a good idea of its business model and the value of different types of slot. We note that, in the experiment, there was no common value uncertainty as each airline was given a complete payoff schedule.

There is therefore no strong reason in theory to adopt an open auction within the experiment, whereas there as a strong practical case for adopting a sealed-bid approach. In the real world, however, either might be used.

Turning now to package design, a choice needs to be made between 'fixed' and 'flexible' packages. In the former the auctioneer pre-determines the bundles of slots that are grouped together. In the latter the airlines are free to determine this. The latter approach seems more reasonable, as different airlines will have different preferences of how to bundle slots together to meet their needs (e.g. long-haul versus short-haul carriers). We also note that stakeholders emphasised the advantage of 'flexible' packages.

The next issue is the bidding language. This refers to how bidders can express their preferences over bundles within the auction environment—which bundles, at which prices, and with what caveats. For example, thinking both about complementarity and substitutability, an airline may want to obtain:

- a package A of morning slots (1, 2, 3) at price p_1; or
- a package B of evening slots (4, 5) at price p_2; but
- not both A and B.

The way to implement this would be to use the XOR bidding language. In notation terms, for example:

- A(1, 2, 3) at p_1; *XOR* B(4, 5) at p_2

Such bids would need to made such that bidders do not exceed their budget (and the XOR format helps in this regard).

Ideally, bidders would fully express their preferences over all *m* slots that are available. However, this would be overly complex as there would then be an exponential number (2^m - 1) of slot packages to be evaluated. For example, four slots would mean 15 combinations to evaluate; ten slots would mean 1,023 combinations; and 20 slots would mean over 1m combinations. This would require a lot of effort on the part of bidders (although real-life airlines might use an algorithm to do this).

In practice, participants will only be able to express a preference over a more limited number of slot bundles. There will therefore be 'missing bids'. However, this may not be too problematic. It is likely that airlines will have slot preferences that are restricted by their business plans and intended routes (see section 3).

Winner determination is the next issue. This is more complicated in the case of combinatorial auctions than in single-unit auctions. The objective of the auctioneer is to accept bid combinations that maximise total revenue. However, there will be a myriad of bids across overlapping packages, making this exercise non-trivial. Moreover, with the exponential number of packages to be evaluated (save for the missing bid problem discussed above), the problem of revenue maximisation is 'N-P hard'. This means that determination of the winner(s) to maximise revenue is not mathematically tractable.

An algorithm therefore needs to be used that, while an approximation, arrives at a solution to the 'winner determination problem' (WDP). However, the latest theory shows that the VCG (second-price sealed-bid) combinatorial auction leads to truthful bidding *if and only if* an exact solution can be obtained to the WDP.

There is a way around this. In a seminal study, Lehmann et al. (2002) showed that using a simple 'greedy algorithm' to approximate a VCG can be tractable, solving the WDP while leading to truthful bidding.[128] However, this requires a restriction to be imposed: that bidders are 'single-minded'—i.e. each bidder has only one package in mind, which they either do or do not get.

In the current context, this would mean an airline wanting either a particular bundle or nothing at all. However, as alluded to above, airlines are likely to be more 'multi-minded'—hence the XOR bidding language.

Thus, to solve the WDP, any approximate algorithm we use will not fully replicate the VCG outcome. This also means that there are few additional efficiency benefits of using a second-price over a first-price auction. We also note that, within the experimental environment, communicating the second-price rule to participants for bundles of products may be challenging and more time-consuming. This does not, however, rule out using a second-price auction in a real-world setting.

A3.2.4 Auction design in the experiment

Given the literature review and stakeholder engagement, the auction design in the experiment was as follows:

- auction format—one-off sealed-bid combinatorial auction;
- packaging—flexible packaging (airline decision);
- bidding language—XOR, taking account of both complementarity and substitutability;
- winner determination—using an approximate algorithm;
- payments—first-price rule.

[128] Lehmann, D., O'Callaghan, L. and Shoham, Y. (2002), 'Truth revelation in approximately efficient combinatorial auctions', *Journal of the ACM*, **49**:5, pp. 577–602.

A4 Comparison of demographics across treatments

A4.1 Age

Figure A4.1 Participants' age

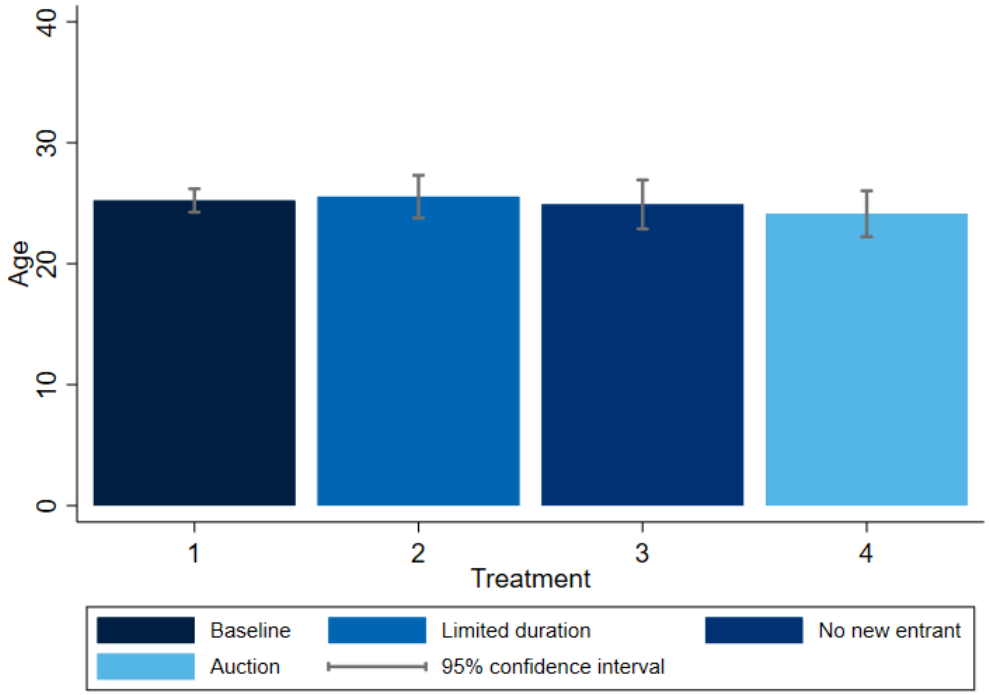

Note: A confidence interval is a way to measure the precision of an estimate. It is a range of values that contain the true statistical value (i.e. across all samples) with a probability of 95%.

Source: Oxera.

A4.2 Flights

Figure A4.2 Average number of flights taken by participants

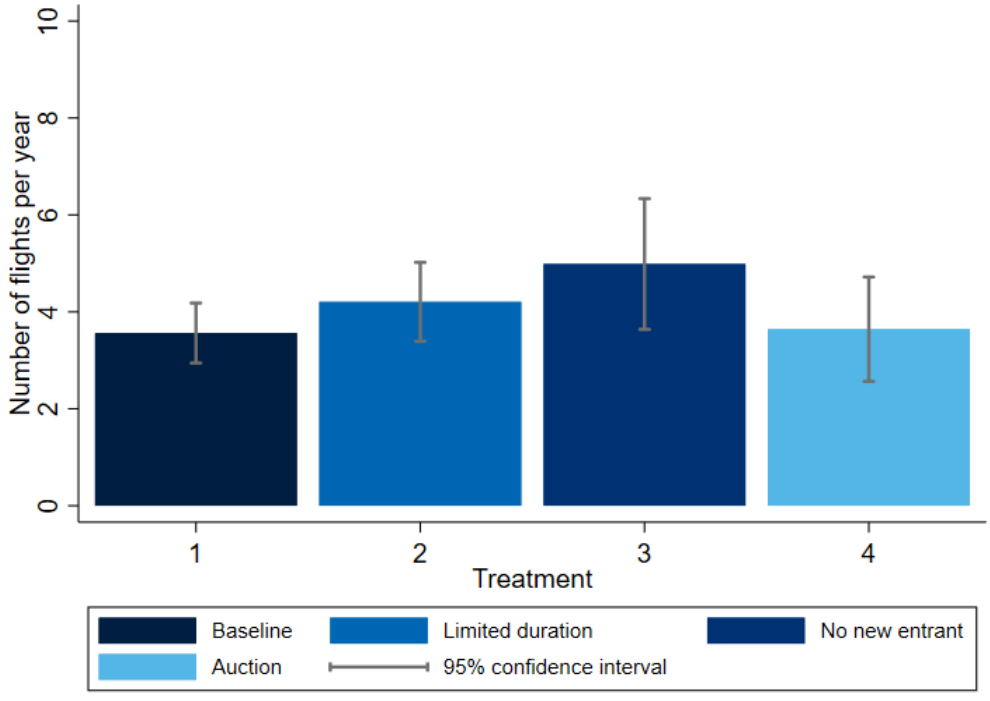

Note: A confidence interval is a way to measure the precision of an estimate. It is a range of values that contain the true statistical value (i.e. across all samples) with a probability of 95%.

Source: Oxera.

A4.3 Gender

Figure A4.3 Participants' gender (percentage of female participants)

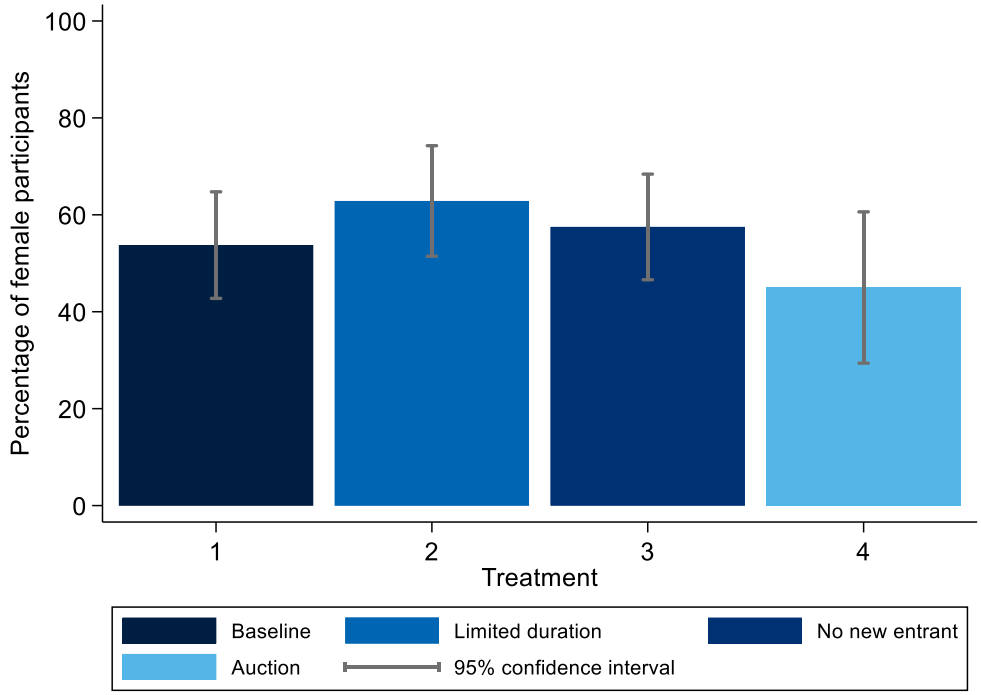

Note: A confidence interval is a way to measure the precision of an estimate. It is a range of values that contain the true statistical value (i.e. across all samples) with a probability of 95%.

Source: Oxera.

A5 Consumer surplus

Intuitively, we would expect an intensifying of competition to result in more benefits for passengers. Airlines have to compete to attract customers, which may involve lower prices or higher quality services. Consequently, competition tends to reduce the total benefits for airlines. Our measure of consumer surplus is driven by the assumption that increases in consumer surplus are driven wholly by decreases in producer surplus. This is conservative because standard economic theory suggests increases in consumer surplus due to competition also come from the elimination of deadweight losses—passengers who were not able to travel are not able to travel at prices that are above their reservation price. We provide a mathematical exposition of our methodology below.

Consider an airline A offering a single early morning flight. Passengers place a value v_1 on the flight. The airline charges a price that depends on the number of competitor flights n on the same route at the same time, earning a payoff $p_A(n)$. Since it is currently the only airline offering the early morning flight, it earns a payoff of $p_A(1)$. Then, the consumer surplus is given by:

$$v_1 - p_A(1) \geq 0$$

The consumer surplus cannot be negative because if it is, it means that the price is so high that passengers would simply not buy the ticket. Given that in the experiment we do not know the true value placed by consumers on the flight v_1, we make the conservative assumption that $v_1 = p(1)$.

Now, let us suppose that another airline B offers the same flight at the same time. Due to the effect of competition, airline A is forced to lower its price, leading to a fall in payoffs from $p_A(1)$ to $p_A(2)$. Consumer surplus on the flight offered by A is now:

$$v_1 - p_A(2) = p_A(1) - p_A(2)$$

which is equal to the loss in payoffs due to competition.

Now, let us consider airline B. It is likely that its passengers place a lower value $v_2 < v_1$ compared to the passengers on airline A. This reflects the possibility that airline B may be less well established, and as a result of being a newer entrant can only capture a segment of the market with a lower valuation on flights. As a result, the payoffs it can earn is also lower at $p_B(2)$.Then, the consumer surplus on the flight offered by airline B is:

$$v_2 - p_B(2) \geq 0$$

Again, we do not know what v_2 is, and so we make the conservative assumption that $v_2 = p_B(1)$.

The above analysis can be readily extended to a route and time with more than two competitors.

www.oxera.com

oxera
compelling economics