# Airline profitability: conceptual framework and review of stakeholder submissions

Report to the Civil Aviation Authority

July 2026

This report has been prepared for the Civil Aviation Authority (CAA) in connection with providing a report on airline profitability - conceptual framework and assessment of stakeholder evidence. To the fullest extent permitted by law, Grant Thornton UK Advisory & Tax LLP does not accept a duty of care whether in contract or in tort (including in negligence) or under statute or otherwise nor assume responsibility to anyone other than the CAA for our work or this report or for any opinions or conclusions that we have formed. We do not accept any responsibility for any loss or damages or costs incurred which arise out of the use of this report by any third party.

We do not warrant or represent that the report is appropriate for your purposes. The report was not created for, and should not be treated as suitable for, any purpose other than that set out in our terms of engagement with the CAA. If you do rely upon the report for any purpose, you will do so entirely at your own risk and you will not bring or threaten to bring any actions, proceedings or claims against Grant Thornton UK Advisory & Tax LLP where the action, proceeding or claim in any way relates to or concerns or is connected with the use of or reliance on the report.

All data has been provided by third parties. We have not verified the accuracy or completeness of any such data. There may therefore be errors in such data which could impact on the content of this report. No warranty or representation as to the accuracy or completeness of any such data or of the content of the report relating to such data is given nor can any responsibility be accepted for any loss arising therefrom.

# Contents

## Table of contents

# 1 Executive summary

## Purpose of the report

1.1 This report is intended to support the CAA's assessment of future regulatory options for capacity expansion at Heathrow Airport ("Heathrow"). It relates to two questions central to estimating the consumer benefits of expansion: whether scarcity rents exist at Heathrow and the extent to which these are captured by airlines in the form of economic (or excess) profits; and how airlines are likely to respond to changes in airport capacity, airline competition and aeronautical charges.

1.2 The report does not undertake new empirical analysis. Instead, it develops a conceptual framework grounded in economic theory and uses that framework to assess the stakeholder material provided to us by the CAA for review. This comprises two papers by International Airlines Group ("IAG"), submitted to the CAA and the Department for Transport in March and May 2026, responding to Frontier Economics' estimates of a congestion premium at Heathrow, which underpin part of the cost-benefit case for expansion; a report by Oliver Wyman, commissioned by IAG, on the level of airport charges airlines could afford to fund expansion (March 2026); and responses from Heathrow Airport Limited ("HAL") and its adviser Frontier Economics ("Frontier") to the IAG and Oliver Wyman material, provided in correspondence of April 2026.

## Overview of the conceptual framework

1.3 The framework treats air travel not as a single market but as a set of overlapping sub-markets,[1] defined across routes, passenger types, booking horizons and time of travel, each analysed according to the conditions prevailing within it. The terms sub-market and segment are used interchangeably in this report; both refer to an element of demand within which competitive conditions are broadly uniform, whether defined at the level of a route or more finely within it.

1.4 The framework links what airlines are trying to achieve – profit maximisation, the levers they use to do so (including pricing, service capacity and network configuration), and the conditions in each sub-market, to outcomes for fares, quantity and profitability. The framework is deliberately stylised: its purpose is to identify the conditions under which scarcity rents would be expected to arise, so that observed outcomes can be interpreted in a structured way, rather than to establish whether such rents are in fact being earned at Heathrow.

1.5 Scarcity rents are higher fares than those that would be found under a perfectly competitive equilibrium, due in whole or in part to airport capacity constraints. Capacity constraints can give rise to airline market power by limiting the likelihood and threat of entry, thereby preventing the conditions required for perfect competition to be achieved. There are also features of the aviation market, other than airport capacity constraints, that can contribute to, or reinforce, scarcity rents. Scarcity rents are expected to be strongest where airlines have a

[1] Batley, R., Wheat, P. and Mackie, P. (2019) *Independent Peer Review of Recent Research on the Existence of Scarcity Rents at Heathrow: Final Report*. Leeds: Institute for Transport Studies, University of Leeds. Available at: https://www.caa.co.uk/publication/download/17572 (Accessed: 1 July 2026).

greater degree of market power, where substitutes are weakest and where demand least elastic (noting that these factors are typically correlated).

1.6 Scarcity rents need not appear as headline accounting profit. For example, if an airline were to reinvest economic profits captured into lower margin/loss making routes to maintain market share, this would erode reported accounting profit. Assessing whether returns exceed the competitive level requires adjustment for the opportunity cost of capital and the value of slots, and should be judged over the demand cycle rather than on single-year results.

1.7 The framework considers how airlines would be expected to respond to capacity expansion and to the corresponding increase in charges that may be required to fund it. As expansion relaxes the capacity constraint, airline scarcity rents would be expected to erode, as expansion allows the entry, competition and supply response that a capacity constraint otherwise blocks. Some rents may persist where airlines retain market power from other sources, such as product differentiation or network effects. As the capacity constraint eases, more of any charge increase would be expected to reach passengers. Ultimately, when no capacity constraint or other sources of market power remain, and the market is assumed to reach a perfectly competitive equilibrium, all scarcity rents would disappear (and be passed on to consumers through lower ticket prices) whilst increases in charges would be passed on to consumers in full.

## Assessment of stakeholder evidence

1.8 Four submissions were reviewed, and they disagree about whether airlines at Heathrow benefit from scarcity rents/congestion premium. This report uses scarcity rents and congestion premium interchangeably. IAG, the owner of British Airways, argues in two papers that no such premium is being earned. The first paper states that if airlines were earning a large premium, it would be visible in at least one of: their profits, their costs, how full their aircraft are, and their fares once airport charges are stripped out – and it is visible in none. The second says the high prices paid to buy and sell Heathrow take-off and landing slots cannot simply be turned into a per-passenger fare premium.

1.9 A report by Oliver Wyman, commissioned by IAG, estimates that airlines could absorb an increase in airport charges of only about £4.50 to £5.50 per passenger, enough to fund up to £30 billion of expansion, before their margins become unviable. It argues that the higher charges needed to pay for expanding Heathrow would exceed this, and so airlines could not afford them.

1.10 HAL and its adviser Frontier disagree with the IAG and Oliver Wyman analyses, arguing that their results cannot be reproduced and do not reflect how airlines actually behave, and that the premium is well established by earlier work from the CAA, the Airports Commission and the Department for Transport.[2]

1.11 Overall, our assessment, drawing on both economic theory and our review of the stakeholder evidence, is that scarcity rents are likely present at Heathrow, though their scale and the stage of the supply chain in which they are captured are materially uncertain. From a theoretical perspective, Heathrow operates at capacity (not disputed by IAG), which makes scarcity rents possible. Whether they arise, and at what scale, depends on demand and competition in each sub-market. From an empirical perspective, we see validity in many of IAG's analyses and criticisms of the magnitude of Frontier's estimate of scarcity rents. We view these as reasons to potentially narrow where in the supply chain any rent may sit, and to treat Frontier's

[2] References for the CAA, Airports Commission and DfT analyses are set out in our assessment of the HAL and Frontier response (Section 5, Evidence 1)

quantification as likely towards the upper end of the plausible range. We therefore view IAG's evidence as motivating further analysis, rather than as sufficient to rule out the existence of scarcity rents.

1.12 On the second question this report addresses – how airlines would respond to increased aeronautical charges – the framework describes two analytical extremes: full absorption into scarcity rents while a constraint strictly binds and the rent in a sub-market exceeds the charge, and full pass-through under perfect competition. Because any charge increase considered here would accompany capacity expansion, the share reaching passengers would be expected to grow as markets become more competitive, at a pace depending on the conditions in each sub-market, and alongside the downward pressure on fares from the supply response capacity expansion allows. The Oliver Wyman report adopts the first of these extremes. Its model allows no part of a charge increase to be recovered from passengers through fares, prices are held flat in real terms (OWR Appendix A.1), so every pound of additional charge reduces airline profit directly. On that assumption, airline profits at Heathrow reach an unsustainable level once charges rise by more than £4.50 to £5.50 per passenger, which is what limits the affordable investment to £30 billion. The assumption drives the result: if airlines can pass even part of the charge to passengers, as the framework predicts and as the report's own analysis suggests has happened in the past (Heathrow's charges are already roughly twice those of other London airports, and Oliver Wyman finds the fares on long-haul routes have in effect funded them), then the affordable charge and the affordable investment are both larger than the figures presented. We therefore consider the £4.50 to £5.50 and £30 billion figures likely to understate what is affordable, and do not treat them as established constraints.

1.13 The report also models how airlines would move capacity away from Heathrow if charges rose too far. It assumes whole groups of carriers relocate once the group's average profit margin falls below a threshold. Our conceptual framework suggests instead that these decisions are made route by route, and that Heathrow's network and hub economics would slow any relocation for many carriers, so the modelled response is likely to overstate how readily capacity would leave.

1.14 No party presents evidence on the degree of pass-through itself. Fare responses to historic changes in Heathrow's charges could be analysed to inform where between the extremes the true split lies.

1.15 The analyses by IAG, Oliver Wyman, HAL and Frontier all rely on assumptions we are unable to verify. Our finding that the magnitude of any rents is uncertain therefore applies in both directions: we neither accept that no premium exists nor endorse a single figure for its size.

# 2 Introduction

2.1 Heathrow Airport is one of the busiest airports in the world and operates at nearly full capacity.[3] This scarcity has long been a central feature of the competitive environment at Heathrow and has played a prominent role in regulatory discussions regarding airport expansion, aeronautical charges, and the distribution of economic rents between airlines, the airport, and consumers. In this context, an important question for policymakers is the extent to which airlines operating at Heathrow may derive scarcity rents or earn economic profits as a result of these constraints, and how changes in capacity or competition might affect airline behaviour, pricing, and service provision.

2.2 These issues are particularly relevant in light of ongoing regulatory considerations by the CAA relating to future capacity expansion at Heathrow, and the potential implications for consumer outcomes. A number of stakeholders have provided submissions to the CAA setting out differing perspectives on whether scarcity rents exist, the extent to which such rents translate into airline profitability, and how airlines are likely to respond to changes in airport capacity, competition, and aeronautical charges. These submissions reflect differing interpretations of both economic theory and available evidence, highlighting the need for a structured and theory-informed assessment.

2.3 Against this background, the purpose of this report is to develop a conceptual framework for analysing airline profitability and behaviour in the presence of capacity constraints, and to use this framework to assess the arguments and evidence presented by key stakeholders. The report focuses on a theoretical and conceptual assessment rather than undertaking new empirical analysis. In particular, it draws on established economic theory and relevant literature to set out the mechanisms through which airport capacity constraints, competition, and cost conditions influence airline decisions on pricing, network configuration, and service capacity. This framework is then used to critically review and assess stakeholder submissions written by IAG, Oliver Wyman, HAL and Frontier, with a view to identifying areas of alignment, divergence, and uncertainty.

2.4 The analysis is intended to support the CAA's consideration of two key questions:

- the extent to which airlines operating at Heathrow earn economic profits or benefit from scarcity rents; and,
- how changes in capacity, competition, and aeronautical charges are likely to influence airline behaviour and, ultimately, consumer outcomes.

2.5 In doing so, the report distinguishes between economic concepts of profitability (including economic rents) and observed accounting outcomes, and considers the extent to which these may differ in practice.

2.6 The remainder of the report is structured as follows:

- **Section 3** sets out the theoretical foundations underpinning airline pricing and capacity decisions.
- **Section 4** develops a conceptual framework linking revenue generation, cost structures, and key determinants of airline profitability.

[3] Government update on airport expansion - GOV.UK

- **Section 5** presents an assessment of stakeholder arguments against the conceptual framework.
- **Section 6** provides a synthesis and comparative assessment of the different stakeholder positions.
- **Section 7** sets out an annex on supporting theoretical foundations.

# 3 Theoretical foundations

## Overview of relevant economic theory

3.1 This section outlines the key economic theory that underpins the conceptual framework used to analyse airline behaviour. It does not seek to determine a single "correct" approach, but instead presents a range of theoretical perspectives that help explain firms' behaviour and market outcomes.

3.2 To do so, it adopts a simplified representation of the world, intended to highlight the main drivers of airline decision-making and the mechanisms shaping the air travel sector, and, ultimately, airline profits. While abstract, this approach provides a clear foundation for understanding how pricing, costs, and competitive interactions influence outcomes.

3.3 Importantly, the market for air travel is not a single, unified market, but a set of overlapping sub-markets[4] (for example, across routes, passenger types, or time of travel), each with distinct characteristics. These differences imply that no single theoretical market structure fully captures airline competition.

3.4 For these purposes, a market is the set of products that buyers regard as close substitutes, defined over both a product and a geographic dimension; markets can in principle be delineated more formally using standard competition tools, but this report does not attempt to formally define them. Two related levels in the supply chain are relevant: an upstream market for airport operation services, supplied by airport operators to airlines, and a downstream market or markets for air transport services, supplied by airlines to passengers (and, more widely, to freight and corporate customers), which is the focus of this report. The geographic and product scope of each varies with the segment in question. Describing air transport services as a set of overlapping sub-markets is a deliberate simplification adopted for this analysis, rather than a formal market definition; it reflects the observation[5] that the conditions facing, say, a leisure traveller booking a Mediterranean route months ahead differ markedly from those facing a business traveller booking a long-haul city-pair at short notice. The sub-market is treated as the relevant unit because it captures these differences better than taking airport-pairs or city-pairs as markets and allows for a more nuanced view of how airlines generate profits. Accordingly, this section considers a range of relevant theoretical perspectives at a high-level, structured as follows.

3.5 It begins with an overview of firm objectives, profit concepts and cost structure, then details of various market structures: the benchmark of perfect competition before progressing through monopoly, monopolistic competition and oligopoly and how these apply to the passenger airline market. It then considers capacity constraints, strategic behaviour, including game theory and entry deterrence, price differentiation, network economics and service quality. Each layer of theory introduces features relevant to air travel markets to provide the theoretical

[4] The paper on the existence of scarcity rents at Heathrow prepared for the CAA by Leeds University (2019) also observed that the market for air travel is not a single market but a series of overlapping sub-markets with different characteristics and elasticities, with some segments involving competition between airports as well as airlines, and noted that the conditions facing a leisure traveller booking months ahead and a business traveller booking at short notice are very different. The sub-market framing adopted here is consistent with that observation. Available at: https://www.caa.co.uk/publication/download/17572.

[5] This observation is also made in the paper on the existence of scarcity rents at Heathrow prepared for the CAA by Leeds University (2019). Available at: https://www.caa.co.uk/publication/download/17572.

foundation required to develop a framework for assessing the conditions under which airlines may sustain economic profit and scarcity rents. This is not intended to be a comprehensive explanation of each of these topics, rather it sets out the key conceptual considerations helpful to framing and assessing stakeholder views.

## Firm objectives and profit concepts

3.6 To set the foundation for this chapter, we first introduce some very general concepts which are later applied to airlines. In general, economic theory of the firm is based upon the assumption that a firm acts to maximise its profit. To determine the optimal level of output (and/or price), theory states that firms should produce output at a level where marginal revenue is equal to marginal cost, given average economic cost does not exceed average revenue.[6] Although in practice firms' incentives can be more complex than simple static profit maximisation, this assumption provides a useful baseline to assess firm behaviour.

3.7 Economic profit (or excess profit) refers to the situation where total revenue exceeds total economic cost, including the opportunity cost of capital. Normal profit occurs where revenue covers all economic costs, representing the minimum return required to keep a firm in the market. These concepts underpin entry and exit decisions. Market outcomes are determined by the interaction of supply and demand, with prices adjusting to equate the quantity supplied and demanded. The nature of this interaction varies by the type of market structure under consideration.

## Cost structure

3.8 Firm costs are a combination of fixed and variable costs. Fixed costs do not change in the short run, regardless of the level of output. Variable costs change with output. Average cost is total cost divided by output, representing the cost per unit produced. Marginal cost is the additional cost incurred by producing one further unit of output, and is solely a function of variable costs since fixed costs do not vary with small changes in output.

3.9 These cost structures are central to understanding firm behaviour across market structures. A firm's market power and its cost structure together determine the price it sets and the quantity it produces. The resulting gap between price and marginal cost is a measure of that market power and bears on the efficiency of market outcomes, while the gap between price and average cost determines whether the firm earns economic profit, normal profit, or makes a loss, and therefore underpins entry and exit decisions in the long run.

## Market structures

3.10 This section describes the assumptions and outcomes for different market structures and considers which types of airline sub-markets each might describe. The discussion below concentrates on how each structure applies to air travel and what it implies for the questions in this report. The standard textbook definitions, assumptions and welfare properties of each structure are set out in full in Annex A.

### Perfect competition

3.11 In practice, even the most competitive airline markets are unlikely to satisfy the strict criteria of perfect competition, set out in Annex A. The price-taking and product homogeneity

[6] In the instance where average economic cost is greater than average revenue, it would be optimal for the firm to shut down.

assumptions rarely hold on the supply side, since flights differ by departure airport, timing and schedule and are therefore not perfectly substitutable. The central role of yield management in airline strategy is itself evidence that carriers retain some ability to influence price, so perfect competition is best treated as a benchmark rather than a description of any real airline market.

### Monopoly

3.12 In contrast, as set out in Annex A, monopoly describes a market with a single seller, no close substitutes and large barriers to entry. A pure monopoly of this kind is rare in air travel, because most passengers retain some alternative: a different destination for a leisure trip, an indirect routing, a competing hub for transfer traffic, or another mode of transport. The structure is most closely approximated where these alternatives are least attractive or viable, for example lifeline routes to remote or island communities, or routes where geographic factors, slot scarcity or regulatory barriers leave a single carrier without an effective competitor.

3.13 Where entry is blocked and market power is sustained, the monopolist is able to price above average and marginal cost, and earn positive economic profit in both the short and long run.

### Monopolistic competition

3.14 Monopolistic competition is a market in which competing firms have free entry into the market and compete on price and product differentiation. This product differentiation provides firms with a degree of pricing power (see Annex A). Where the market is a route between two airports, carriers compete not only through price but also through product features, such as schedule (timing and frequency of flights), network (opportunities for onward connection), service quality or brand perception.

3.15 In pricing and welfare terms, the structure sits between the benchmarks of perfect competition and monopoly. Price would tend to exceed marginal cost in the short run, so output could fall short of the welfare-maximising level, but free entry would compete economic profit away in the long run. Where there is product differentiation, passengers can gain from variety despite supernormal profit, meaning the overall welfare outcome is not as clear cut as observing solely price and quantity.

### Oligopoly

3.16 Oligopoly involves a few sellers that are large relative to the market and aware of their mutual interdependence, so each sets strategy taking account of how rivals will respond (see Annex A). This could apply, for example, to a long-haul route between two major cities operated by two or three airlines, where each sets capacity, frequency and fares taking account of competitors' likely reactions. Oligopoly has less determinate implications for economic profit, prices and quantity produced relative to the previously considered market structures. Outcomes range from the competitive level to sustained excess profit, depending on the nature of strategic interaction (for example, whether firms compete on price or quantity), degree of product differentiation, the conditions governing market entry and the historical context.

### Vertical structures

3.17 A vertical structure describes the relationship between firms operating at different, successive levels of a supply chain, where the output of one level is an input to the next. The market structure can differ at each level, and the level with greater market power tends to capture the larger share of economic profit within the system.

3.18 In air travel, the first level is the airport, which supplies infrastructure (runways, terminals, slots) to airlines. An airport could have substantial market power, such as Heathrow.[7] As such, it is subject to economic regulation by the CAA through a 5-year price control.

3.19 The second level is the airlines, which buy access to that infrastructure to serve passengers. It is important to consider where any economic profit would be captured along the chain. In theory, a monopoly airport upstream can use its charges to extract profit from the airlines downstream, raising their costs. This can leave airlines competing for thinner margins while the airport captures the surplus, which is part of why airport charges, such as those at Heathrow, can be subject to economic regulation.[8]

### Strategic behaviour[9]

3.20 The discussion that follows is confined to the most basic strategic models. Game theory admits more elaborate variants, but these are not pursued here. The aim is to draw out the intuitions most relevant to the airline market rather than to provide a comprehensive treatment. Two features of firm decision-making which can be employed strategically in an oligopolistic market structure are quantity-setting and price-setting. These two types of behaviour can be simply represented using a duopoly market structure (with just two firms in the market) and are captured respectively by the Cournot and Bertrand models.

3.21 In the Cournot model, firms compete by simultaneously choosing how much to produce, each taking the other's output as given. Each firm's optimal output decision, governed by its reaction function,[10] depends on what it expects its rival to produce. The market reaches equilibrium where both reaction functions intersect: neither firm can increase its profit by unilaterally changing its output level. The resulting market price exceeds marginal cost, meaning firms earn positive economic profit.[11] Consumer welfare sits between the extremes of perfect competition and monopoly. In welfare terms, the Cournot outcome lies between the two benchmarks: because price exceeds marginal cost, output is below the competitive level and there is a deadweight loss,[12] but one smaller than under monopoly.

3.22 In the Bertrand model, firms compete by simultaneously setting prices, taking their rival's price as given, and letting the market determine the quantity supplied. Where products are homogeneous, competition drives prices down to marginal cost, eliminating economic profit

[7] CAP1133: Market power determination in relation to Heathrow Airport – statement of reasons | UK Civil Aviation Authority

[8] The supply chain is in practice more complex than a simple airport-to-airline relationship. Airlines buy a range of further inputs whose markets cross between the airport and the carrier, including fuel, ground handling and aircraft maintenance, as well as inputs sourced more widely such as aircraft and crew, and the conditions in each of these markets affect where profit is ultimately captured. The framework in Section 4 concentrates on the airport, airline and passenger levels because these are central to the questions on scarcity rents and charge pass-through, while recognising that other levels of the chain can, in theory, absorb part of any rent. Because passengers typically buy air travel as one element of a wider trip, part of their willingness to pay may also be absorbed by complementary markets such as surface access, airport parking and accommodation, rather than by the airline.

[9] In economics, strategic behaviour refers to decision-making where a firm deliberately accounts for the potential actions, reactions, and motives of other market participants. It occurs when market players are directly interdependent, meaning one's success depends on the choices others make.

[10] A reaction function shows a firm's profit-maximising output level for every possible output choice of its rival, holding all other factors (such as costs and demand) constant.

[11] Each firm's profit is a function of the number of firms competing in the market. As the number of firms increases, the profit per firm falls at a disproportionate rate. Eventually, profit will tend to 0 as more firms enter the market.

[12] Deadweight loss is the loss of economic surplus or welfare, either from consumers or firms, from the market equilibrium not being at the perfectly competitive level.

entirely, even with only two firms in the market. This outcome, equivalent to perfect competition, is known as the Bertrand paradox.[13] In contrast, where products are differentiated, prices remain above marginal cost, as each firm retains some degree of pricing power. The welfare implications follow directly. With homogeneous products, price competition drives the outcome to the competitive benchmark, so welfare is maximised and there is no deadweight loss despite there being only two firms. With differentiated products, price sits above marginal cost, so output is below the welfare-maximising level and some deadweight loss arises, though consumers gain from the greater variety on offer.

3.23 In the airline market, entry deterrence is an important factor to consider. It describes strategic behaviour undertaken by an incumbent firm with the aim of discouraging potential rivals from entering the market and can be analysed using game theory. Unlike the Cournot and Bertrand models above, in which firms move simultaneously, entry deterrence is naturally represented as a sequential game, in which one firm acts before the other and the order of moves matters. In a sequential game, the incumbent moves first, taking actions that alter the payoffs facing a potential entrant such that entry becomes unprofitable.

3.24 A central insight from this theory is that the credibility of the deterrent threat determines its effectiveness. Prior to entry, an incumbent that commits to expanding capacity through, for example, acquiring additional aircraft, securing terminal infrastructure, or locking in long-term slot agreements, signals to potential entrants that it will compete aggressively if they enter. This commitment is considered credible because the associated costs are sunk, meaning the incumbent has no incentive to soften its competitive response post-entry. An incumbent that instead signals its intention to cut prices only in the event of entry offers a less credible threat, since the rational response post-entry may be to accommodate the rival rather than engage in costly price competition. In aviation markets, entry deterrence may therefore manifest through fleet expansion ahead of anticipated competition, slot acquisition at capacity-constrained airports, or the strategic use of frequent flyer programmes to raise switching costs.

3.25 The theory implies that an incumbent may have an incentive to hold or acquire slots in ways that limit the capacity available to potential rivals, and that high frequencies on a route may themselves raise the threshold an entrant must meet to establish a competitive schedule.

3.26 The retention of slots through periods of weak demand may also be consistent with profit maximisation where the option value of the slot exceeds the short-run operating loss. Decisions on whether, and to whom, slots are sold or leased where trading is permitted, and arrangements such as alliances and joint ventures that coordinate slot use across carriers, would also be expected to carry competitive implications. The extent to which any of these possibilities are realised in practice is an empirical question beyond the scope of this conceptual discussion.

3.27 In addition, the existence of scarce capacity would be expected to discourage exit, since relinquishing slots means giving up a valuable asset and use-it-or-lose-it requirements create an incentive to retain them. The theory therefore suggests that incumbent positions would tend to be entrenched, and that the position of a carrier holding a large portfolio would be reinforced, where slot holdings combine with network feed and loyalty programmes to raise the scale at which a rival would need to enter to compete effectively.

3.28 It follows that, competitive pressure exerted on an airline under these circumstances, would be through flights at other airports, rival hubs for transfer traffic, or other modes of transport, rather than through new entry at the constrained airport itself.

[13] It is termed a paradox because the competitive outcome of price equal to marginal cost, normally associated with many firms, arises with as few as two, which is counter intuitive.

## Capacity constraints

3.29 A consideration relevant across all market structures is the role of capacity constraints. A capacity constraint arises where a firm's ability to supply is limited by a fixed factor of production that cannot be expanded, meaning that supply cannot respond fully to demand even where it would be profitable to do so. A capacity constraint can raise prices through two parallel channels. The first is dynamic: by limiting the supply response of entrants and incumbents, the constraint slows or blocks the entry and expansion that would otherwise erode profit over time. The second is static: even setting the threat of entry aside, a binding constraint that holds quantity below the level the market would otherwise supply can push price above marginal cost at a point in time, in the manner of a standard supply-and-demand diagram in which a vertical supply segment meets demand above the competitive price. The size of this static effect depends on the price sensitivity of the market in question. This is a function of demand characteristics and the availability of substitutes to meet excess demand. In the air travel context, demand varies by the nature of travel (business or leisure) and by the time of day, week or year (such as higher leisure demand in summer). Substitutes in this case include the availability of alternative nearby airports offering flights to the same destinations. The price increase resulting from capacity constraints is larger where there are fewer good substitutes and/or demand is less price sensitive, and smaller where close substitutes exist and/or demand is more price sensitive.

3.30 The impact of capacity constraints on the outcomes in various market structures[14] are as follows:

- In a monopoly, capacity constraints reinforce the incumbent's market power by raising the cost and difficulty of entry, further entrenching the conditions under which economic profit can be sustained. A strictly binding constraint[15] limits the volume that can be sold to below its profit-maximising optimum and prices higher than the associated optimum price, leading to lower profits.
- In monopolistic competition, a capacity constraint may slow the entry process that would otherwise erode economic profit in the long run.
- In oligopolistic markets, capacity constraints are particularly significant. In a market where capacity is limited, competing firms commit to a proportion of this capacity. The constraint makes this a credible and binding commitment as firms cannot subsequently increase quantity beyond market capacity to increase market share. Firms then compete on price within their allocated capacities, setting prices to maximise profit. As demand exceeds constrained supply, this leads to an outcome where the equilibrium price will be a mark-up on marginal cost.

3.31 Where capacity is scarce and difficult to expand, the conditions for sustained economic (or excess) profit are materially strengthened across all market structures. Excess profits caused by capacity constraints (and their implications for suppliers' market power) are referred to as scarcity rents. In theory, these rents could be absorbed anywhere along the vertical chain. This could be by the airport (through increased charges to airlines) or the airlines (through increased ticket prices to consumers). It could also be through other parts of the supply chain, such as car parking operators, surface access operators, ground handling operators or airport

[14] Perfectly competitive market structure is not addressed here as a capacity constraint reflected a departure from the theory underpinning this market structure.

[15] A strictly binding capacity constraint reflects a scenario where the capacity at which the constraint is fixed is below the equilibrium quantity at which a firm would supply. A capacity constraint more generally is any constraint below the perfectly competitive optimum.

retailers, depending on the extent to which capacity constraints and market power translate into those markets. As described above, the entity which absorbs any scarcity rent which exists depends on the level of market power at each stage of the supply chain. In principle, given Heathrow faces price regulation on the fees charged to airlines, it may be prevented from capturing scarcity rents. Price regulation extends to commercial revenue. Again, in the case of Heathrow, the airport is likely to exert a degree of market power over the car parking and commercial retail markets at the airport, which is likely to constrain the extent to which firms operating in those markets will be able to earn excess profits. Other key inputs in the air travel supply chain, for example plane fuel, likely come from competitive markets, which could constrain their ability to collect scarcity rents. Therefore, if scarcity rents exist in air travel, they would likely be absorbed, at least to some extent, by airlines.

3.32 In the context of Heathrow, where capacity likely binds for a significant proportion of routes, the slots available cap the number of air traffic movements airlines can operate in the short and medium run. Airlines nevertheless sell seats rather than movements, and the supply of seats retains some flexibility even when movements are fixed: an airline can deploy larger aircraft, configure cabins more densely, and raise load factors. The short-run supply of seats at a constrained airport therefore could be characterised as strongly upward sloping as capacity is approached,[16] ahead of the fixed limit on the number of seats that can be supplied when capacity is reached. As demand presses against the slot constraint, fares and yields are bid upwards towards the level demand will bear, which this residual supply response moderates the effect until the flexibility is exhausted.

3.33 On this reasoning, airline fares and yields would be expected to sit above the level that would prevail were entry and expansion possible, with the effect strongest where substitutes are weakest and demand least price sensitive (or least elastic), and weaker on segments disciplined by alternatives at other London airports, competing hubs or other modes. The less price-sensitive demand is (or less elastic), the smaller the volume an airline gives up by raising its fare, so the larger the mark-up it can profitably sustain. Where passengers have few substitutes, the constraint on supply therefore translates into a larger increase in fares relative to an unconstrained scenario.

3.34 In the airport context, runway and slot capacity is an intermediate input: airlines do not sell slots to passengers but combine them with other inputs (aircraft, crew, fuel) to supply the flights and seats they sell, so a constraint on slots is a constraint on an input to airline output rather than on that output directly. When the capacity constraint is released, the fixed factor that was limiting supply is relaxed (e.g. additional slots become available to incumbents and to potential entrants). Several mechanisms from the preceding theory would then be expected to operate together. The entry process that capacity scarcity had previously blocked or slowed can resume and new firms can enter the market. In oligopolistic markets the credibility of strategic commitments weakens, since an incumbent can no longer rely on scarce slots to deter entry or restrict rivals' capacity; the threat of entry, even if not realised, disciplines incumbent pricing. And the loss of the binding constraint removes the supply-side rigidity that had allowed price to sit above marginal cost so supply can now expand in response to higher fares, so prices are bid back down towards the cost of supplying additional output.

3.35 In theory, the combined effect is to erode the scarcity rent. Any residual market power arising from product differentiation, network effects or sunk-cost advantages would continue to support prices above marginal cost, so rents are eroded but not necessarily eliminated.

[16] Batley, R., Wheat, P. and Mackie, P. (2019) *Independent Peer Review of Recent Research on the Existence of Scarcity Rents at Heathrow: Final Report*. Leeds: Institute for Transport Studies, University of Leeds. Available at: https://www.caa.co.uk/publication/download/17572 (Accessed: 1 July 2026).

## Price discrimination

3.36 A further dimension of firm behaviour is price discrimination. Where a firm has market power, it may be able to charge different prices to different customers or across different market segments, rather than setting a single uniform price. This practice is known as price discrimination, and it is a defining feature of pricing in markets where firms can separate customers into groups with different levels of willingness to pay.

3.37 A profit-maximising firm with market power sets price above marginal cost, and the size of that mark-up depends on the price elasticity of demand it faces (the percentage change in quantity demanded in response to a percentage change in price): the less responsive demand is to price (less elastic), the larger the mark-up the firm can profitably sustain. Where a firm can separate customers into segments with different elasticities, it charges a higher price to those whose demand is less elastic. For example, airlines may charge higher mark-ups on business-heavy routes (e.g. London–New York), where demand is relatively inelastic, than on leisure routes (e.g. London–Malaga). The optimal mark-up above marginal cost in each segment is inversely proportional to that segment's price elasticity of demand (less elasticity means higher mark-up over cost). Within a single route, an important mechanism airlines use to separate these segments is the booking horizon: business travellers, whose demand is typically less elastic, tend to book much closer to departure, so fares that rise as the departure date approaches charge higher prices to precisely those passengers least sensitive to price. Another important mechanism is the cabin class, which combines price discrimination between different passenger types and product differentiation in service quality. Certain passengers are not willing to lower quality of travel so are less responsive to price for business or first-class tickets. Other passengers will be more sensitive to price over quality, particularly on short-haul flights.

3.38 The intensity of price discrimination can vary. More granular discrimination is relevant to welfare as well as to profit: by offering lower fares to more price-sensitive passengers who would not travel at a higher, uniform price, an airline can expand output above the level a single-price firm would supply, so its effect on total welfare may be positive even though it raises the firm's profit. In aviation this is operationalised through revenue (or yield) management: airlines allocate seats on a flight across multiple fare classes and adjust availability over time, trading off load factor against average yield. Because the marginal cost of carrying an additional passenger on an operating flight is low, the incentive is to fill otherwise empty seats with lower fares while protecting capacity for later-booking, less elastic demand. The fare structure and the management of seat availability are therefore central to how any market power is converted into revenue. The mark-up that a firm can sustain in any given segment is also a function of its degree of market power in that segment. As discussed above in the context of monopoly and oligopoly, market power arises where products are differentiated and close substitutes are limited. Product differentiation, switching costs and barriers to entry all weaken the competitive pressure a firm faces in a segment. The weaker that pressure, the less elastic demand tends to be, and the higher the mark-ups that are feasible. Conversely, in segments where products are close substitutes and competitive pressure is strong, demand elasticity will be higher and mark-ups will be constrained accordingly.

## Network economics

3.39 A further consideration relevant to the analysis of firm profitability is network economics. Industries with network characteristics are those in which the value of a firm's offering increases as the scale and connectivity of its network grows, generating demand-side effects

that reinforce the competitive position of larger incumbents. On the supply side, network scale can also confer cost advantages through the more efficient utilisation of shared inputs across a larger operation.

3.40 Together, these effects create a self-reinforcing dynamic whereby firms with more extensive networks are better placed to attract customers, generating revenue that can be reinvested to further expand the network.

3.41 This can give established firms a structural advantage over potential entrants, reinforcing barriers to entry and supporting the sustained earning of economic profit. In the context of air travel, for example, a larger route network generates a much greater number of possible origin-destination combinations through transfers. This increases the total potential pool of demand that the airline can serve. It can also make smaller routes commercially viable by consolidating traffic destined to multiple destinations on a route to a main hub airport (the hub-and-spoke model). More importantly, it can also make the airline more attractive relative to competitors on main trunk routes where the additional demand generated from feeder routes can enable it to profitably operate a higher service frequency or larger aircraft (which in turn may have cost advantages). Operating a larger route network may therefore confer an airline with a degree of market power on both trunk and feeder routes.

3.42 Network effects can arise at the level of the airport as well as the individual carrier: a large hub offering many connections is more attractive to passengers, and therefore to airlines, so the connectivity provided by other carriers at the same airport can raise the value of operating there. This helps explain why a carrier facing a binding constraint at a hub would not simply relocate flights to an unconstrained airport: doing so would forgo the connecting traffic, frequency and network value that the hub provides. Meanwhile an entrant could not replicate the incumbent's position without assembling a comparable network. Competitive pressure is therefore more likely to come from rival hubs, from other airports serving the same local region for point-to-point demand, or from other modes, than from new entry at the constrained airport itself.

### Service quality

3.43 Market conditions can also be expected to have an impact on service quality. Key features of service quality are the frequency and range of routes available to consumers. This is related to the consideration of quantity above but aligned more closely with how airline offerings meet consumer needs. Frequency on any route is bounded by the slots available, and therefore capacity growth would tend to come through aircraft size rather than additional services. The range of destinations served may be affected where slots are directed to their highest-contribution uses, which the network logic set out above suggests could favour higher-value routes and the feeder services supporting the long-haul network over thinner point-to-point routes.

3.44 The implication for more abstract service quality (e.g. access to lounges, food quality, in-flight entertainment) is less determinate: the theory points in two directions, since reduced competitive pressure on protected routes may weaken the incentive to improve quality, while competition for premium and time-sensitive demand, and the role of product differentiation in supporting fares, may work in the opposite direction.

3.45 Taken together, the framework suggests that consumer choice over frequency, timing and the range of destinations may be more constrained than at an unconstrained airport, alongside the more visible effect on fares.

# 4 Conceptual framework: airline behaviour and profitability

4.1 This chapter develops a conceptual framework for how airlines make pricing and service capacity decisions, and what those decisions imply for fares, route availability and profitability.

4.2 This section identifies the features of the economic theory in Chapter 3 that are most relevant to air travel markets, and assembles them into a single framework that links market conditions to airline behaviour and to outcomes.

4.3 The features set out below are intrinsic characteristics of the market: they describe the underlying structure of supply and demand, the scope over which airlines pursue profit, how costs are incurred, and what constraints airlines face. They exist independently of how competition subsequently plays out. We identify them because they are the aspects of the economic theory in Chapter 3 most relevant to shaping airline incentives and behaviour, and form the bridge between the theory and its implications for fares, capacity and profitability in each sub-market.

4.4 The framework is deliberately stylised. It does not seek to give a definitive answer on whether any airline earns excess returns; its purpose is to identify the conditions under which different outcomes are likely to arise, so that observed behaviour and profitability outcomes can be interpreted in a structured way.

## Economic theory applied to air travel markets

4.5 Chapter 3 sets out relevant economic theory and how it may be applied to air travel. Six features of that theory are particularly relevant to air travel and form the building blocks of the framework.

4.6 **No single market structure applies;** the relevant unit is the sub-market. Air travel is a set of overlapping sub-markets, defined across routes, passenger types, booking horizons and time of travel, each with distinct characteristics. Some of these sub-markets approximate competitive conditions, others monopoly, and many sit between. The relevant unit is the sub-market rather than the route, because a single route typically spans several sub-markets: an airline may face very different degrees of competition for different passenger types on the same route, so that demand from time-sensitive passengers with few alternatives is subject to less competitive pressure than price-sensitive demand for which substitutes are readily available. The framework must therefore be capable of being applied sub-market by sub-market, with the applicable structure identified from the conditions in each.

4.7 **Airlines optimise at the level of the network.** The profit-maximising theory of the firm translates, in aviation, into a firm that maximises profit across its network rather than on each route independently. A route may be operated at low or negative standalone margins where it feeds passengers into more profitable services elsewhere, a feature central to hubs like Heathrow, where short-haul services supply connecting traffic to long-haul routes. Cargo and corporate revenues reinforce this, since their value to customers depends on the reach and connectivity of the network rather than on any single route. Network economics also implies a self-reinforcing advantage: a larger network is more attractive to passengers and to corporate

and cargo customers, which supports revenue that can fund further expansion, raising barriers to entry.

4.8 **The cost structure is characterised by high fixed costs**. Aircraft, slots, ground operations and overheads are largely fixed. The marginal cost of an additional passenger on an operating flight[17] is comparatively low, comprising only items such as the extra fuel to carry the weight of the passenger and of their baggage, catering, and any per-passenger charges. This creates strong incentives to fill seats, makes the trade-off between load factor and yield central to pricing, and means costs are sticky when demand falls. The incentive to fill seats is nonetheless tempered by the need to hold back capacity for late-booking, less price-sensitive passengers willing to pay a higher fare for flexibility, so airlines do not simply discount to fill every seat but manage availability to protect higher-yielding demand.

4.9 **Price differentiation is feasible and pervasive.** Where airlines have a degree of market power, economic theory predicts mark-ups that vary inversely with the price elasticity of demand of each segment. Airlines can separate segments in practice, so differentiated pricing is a standard feature of the market rather than an exception.

4.10 **Competition is strategic.** On many routes there are few enough carriers that each recognises its influence on the others, and there are some barriers to entry (e.g. high capital requirements to invest in aircraft, network effects – hub-and-spoke incumbents offer far more connections than new entrants can match). These features match the defining conditions of oligopoly. Oligopolists can strategically deter entry through sunk commitments, such as fleet, slot holdings and loyalty programmes which are investments that are costly to reverse and signal a credible long-term commitment.

4.11 **Capacity constraints are pivotal**. Airport slots and infrastructure are fixed factors that cannot be expanded in the short run. As Chapter 3 set out, capacity constraints strengthen the conditions for sustained economic profit under every market structure, because they limit the supply response that would otherwise discipline prices and erode profit.[18]

## The conceptual framework

4.12 The framework connects the following elements: what airlines are trying to do, the levers they pull to do it, the market conditions that shape what those levers achieve, and the range of outcomes that follows.

4.13 First, the airline objective. Airlines maximise profit across the network as a whole, rather than route by route. Revenue comes from several streams – passengers, cargo and corporate contracts – and the value of each grows with the reach of the network.

4.14 To pursue this, airlines make three sets of decisions, and these interact. They set prices, where the fare structure used to segment demand matters more than the headline fare. They set capacity: which routes to serve, how often, with what size of aircraft, and how seats are allocated across fare classes. And they configure the network, choosing between hub-and-spoke and point-to-point operation, and entering alliances and codeshares.

[17] Over the medium to longer run, where the relevant decision concerns capacity itself, such as whether to operate a service, how frequently, and with what size of aircraft, the marginal cost of carrying additional traffic is materially higher, since it includes the cost of the aircraft, crew and slot use required to provide that capacity.
[18] It is noted that because an incumbent's position rests on the connectivity and feed its network provides, relaxing the capacity constraint does not by itself eliminate its advantage. A new entrant would need to assemble a comparable network to compete fully, so some market power would be expected to persist even after expansion, and the erosion of rents would be partial rather than complete.

4.15 What these decisions achieve in practice depends on the conditions in each sub-market. The key conditions are: the number of rivals and how close their substitutes are; the size, mix and elasticity of demand; the height of barriers to entry; and the extent of airport capacity constraints.

4.16 Together these shape the outcomes in each sub-market. A more competitive sub-market – more rivals, less product differentiation, more elastic demand, lower barriers to entry, and few binding capacity constraints – will tend to produce tighter margins and lower fares. The converse tends to produce higher fares and, where the binding constraint is a fixed factor such as slots, scarcity rents.

In the table below, example air travel submarkets are presented alongside typical market structures, conditions and outcomes.

**Table 1: Economic theory linked to air travel sub-markets**

| Sub-market (illustrative) | Approximating structure | Key conditions | Expected conduct | Indicative outcome | Indicative response to increased airport charges |
|---|---|---|---|---|---|
| Dense short-haul routes served by several carriers with feasible entry | Monopolistic competition / differentiated Bertrand | Many suppliers; differentiated but close substitutes | Price-making through differentiation, constrained by substitutes | Prices above marginal cost; economic profit eroded by entry over time | Fuller pass-through to fares likely: with many suppliers and close substitutes demand is relatively elastic, so little of the charge can be absorbed in margin (inverse elasticity rule) |
| Routes served by a single carrier with no close substitute | Monopoly | One seller; entry blocked by geography, slots or regulation | Pricing above marginal cost | Sustained economic profit possible in short and long run | Pass-through governed by demand elasticity alone: with no rival to switch to, a larger share can be passed to fares, though the profit-maximising mark-up still limits how much |

| Sub-market (illustrative) | Approximating structure | Key conditions | Expected conduct | Indicative outcome | Indicative response to increased airport charges |
|---|---|---|---|---|---|
| Trunk routes between major hubs served by a few large carriers | Oligopoly | Few large sellers; mutual interdependence; entry conditions vary | Strategic capacity, frequency and pricing; possible entry deterrence | Indeterminate; from near-competitive to sustained excess profit | Indeterminate: pass-through depends on how rivals are expected to respond and on the mix of elastic and inelastic demand on the route |
| Any of the above serving a slot-constrained airport at one or both ends | As above, with a strictly binding capacity constraint | Scarce slots; supply cannot respond to demand | Conduct as above with limited supply response | Prices sustained above competitive levels; scarcity rents arise | Zero pass-through, up to the point where the capacity constraint is no longer 'strictly' binding: with supply unable to respond as quantity supplied is already below the airlines profit maximising quantity, all of the incidence falls on scarcity rents already earned by carriers. |
| Distinct passenger segments on a single route | Price differentiation within the prevailing structure | Separable segments with differing elasticities | Differentiated fares via fare structures and revenue management | Higher mark-ups on less elastic segments | Uneven pass-through across segments: more of the increase is recovered from less elastic segments (e.g. late-booking, time-sensitive demand) and less from price-sensitive segments |

Note: In this context, a route may be a city-pair or an airport-pair, depending on the substitutes available to passengers in the specific sub-market. For some passenger markets, such as leisure trips where the destination is not fixed, passengers may treat a wide range of destinations as substitutes, so the relevant market is wider than any single route. This is less material for the market Heathrow serves, which is weighted towards passengers with a fixed destination or a specific hub connection.

### Implications for airlines' responses to increased airport charges

4.17 An increase in aeronautical charges (as may be required to fund capacity expansion) raises the cost to airlines of carrying each passenger, but the party that ultimately bears a charge

need not be the party that pays it. The incidence of a per-passenger charge can fall on airlines and/or passengers, and the share that reaches fares is known as the degree of pass-through. The extent of pass-through is important in understanding the proportion of any eventual reduction in scarcity rents arising from capacity expansion that may be eroded by higher airport costs.

4.18 This subsection sets out the theory on airline price and quantity responses, under different airport charge and capacity scenarios. For ease of exposition, the subsection is divided into three parts: the first two consider, respectively, charges pass-through and quantity responses, absent explicit capacity constraints; the third builds on the previous two by considering how airlines might respond when faced with both capacity expansion and an increase in airport charges.

### How a charge increase is shared between airlines and passengers in theory (absent capacity constraints)

4.19 The extent of pass-through depends crucially on the intensity of competition and would be expected to rise as the number of competing firms increases. We first consider a perfectly competitive market. Where firms are price-takers and free entry drives economic profit to zero, an increase in firms' costs is passed through to consumers in full in the long run. An increase in airport charges raises the cost of supplying a seat like any per-unit cost, and in long-run equilibrium no economic profit remains from which any part of it could be absorbed, so the whole charge falls on fares.

4.20 Under monopolistic competition, pass-through could typically be high, and may in some cases still be complete – for example on routes where carriers compete closely on price for near-identical services: with margins already close to cost, and every carrier facing the same increase, there is little scope to hold price below the new cost without trading at a loss.[19]

4.21 Where an airline has market power, the same forces operate through the profit-maximising mark-up rather than through a market supply curve. A firm that sets price where marginal revenue equals marginal cost re-optimises when the per-passenger cost rises, and pass-through will likely be partial.

4.22 In line with the inverse elasticity rule set out in the price discrimination section above, pass-through is larger the less elastic demand is. Raising fares by the full amount of the charge would reduce sales by more than the firm's own mark-up rule justifies, so pass-through may be partial even where no rival constrains the firm. Generally, a charge may therefore not be expected to pass through in full even in the least competitive sub-markets, unless demand is perfectly inelastic. Where a firm has market power, the shape of demand also matters. Most conventional demand shapes give pass-through somewhere below the full amount of the cost increase, but some demand shapes can push pass-through above the full amount of the increase.[20]

[19] Under Cournot competition with linear demand, pass-through rises with the number of competing firms and approaches the full cost increase as the number of competitors becomes large, even without demand or supply being especially elastic.
Under Bertrand competition, where firms compete on price for a homogeneous or near-identical product, price is driven to marginal cost even with as few as two firms, since any firm pricing above marginal cost would be undercut and lose the whole market, while pricing below marginal cost would mean trading at a loss. A shared increase in marginal cost therefore passes through in full regardless of the number of competing firms.
[20] For a monopolist facing constant marginal cost, pass-through is exactly 50% of any cost increase where demand is linear. More generally, pass-through is below 100% wherever demand is log-concave, a broad class

4.23 The elasticity of demand relative to supply also plays a role. Whichever side of the market adjusts its quantity less in response to price (is more inelastic) ends up bearing more of the increase, since price has to move further to bring supply and demand back into balance. Pass-through is expected to be higher where the cost increase is common to the whole market rather than specific to one firm. A firm facing a cost increase alone has to weigh raising its price against losing customers to rivals whose costs, and so prices, have not moved, so it makes sense to absorb more of a firm-specific increase to protect market share. Where every firm faces the same increase, none of them loses competitiveness by raising price, so prices move further to reflect the shared cost.

4.24 Pass-through is also expected to be higher where the increase falls on marginal cost rather than on fixed cost. From a theoretical perspective, pricing decisions are assumed to be driven by marginal cost, so a marginal cost increase shifts the profit-maximising price directly. Pass-through also tends to be greater over a longer horizon. Short-run frictions, such as the cost of republishing fares or contractual commitments already in place, can delay the adjustment even where the long-run pass-through would be high.

### How airlines might adjust quantity supplied in response to a charge increase in theory (absent capacity constraints)

4.25 A charge increase also alters the return to supplying a seat, so it bears on the capacity an airline provides – the routes, frequencies and aircraft sizes it chooses, and where it places capacity across its network.[21] As with pass-through, the response depends on the sub-market structure.

4.26 Under perfect competition, capacity adjusts through expansion, entry and exit rather than any single carrier's choice. As the charge falls on all carriers alike and as price is equal to marginal cost a charge increase could not be absorbed and capacity contracts until the fare rises to cover the higher cost, and the charge is fully reflected in fares.

4.27 Under monopoly, the carrier re-optimises based upon the increased marginal cost, generally moving to a lower profit-maximising quantity so that some capacity is withdrawn. The reduction is larger the more elastic is demand; because the carrier earns economic profit, it can also absorb part of the charge, so the capacity cut tends to be smaller than under competition.

4.28 Under monopolistic competition, the response resembles the competitive case over the entry-and-exit horizon but works through differentiated products, with capacity withdrawn until the remaining carriers can again cover costs. Differentiation gives each carrier some short-run discretion to hold capacity and absorb part of the charge, so the response is more muted and varies with how close substitutes are across carriers.

4.29 In an oligopolistic sub-market, the capacity response is strategic, and its analysis draws on the strategic behaviour discussed in Section 3. Where carriers compete principally in quantity, a cost increase common to all of them would be expected to lower each carrier's optimal output, reducing capacity while sustaining fares above marginal cost; the extent depends on how rivals are expected to react. Capacity decisions also carry a commitment dimension: holding or acquiring slots, and maintaining frequent flights, can deter entry and support fares, so a carrier may retain capacity it would otherwise trim in order to preserve that position, even where the

that includes linear demand as one specific case, and above 100% only where demand is log-convex, for example close to constant-elasticity demand.

[21] Because the charge is levied per passenger rather than per flight, it also strengthens the incentive to adjust the mix of passengers carried on a given aircraft, for example by favouring higher-yielding business or premium passengers over more price-sensitive ones.

immediate return has fallen. The outcome is therefore indeterminate in the same way as the fare response, ranging from a capacity reduction that leaves fares broadly unchanged to little capacity change with more of the charge passed to fares, depending on the mode of competition, the degree of differentiation and the conditions governing entry.

4.30 Across these structures the capacity response is greater where demand is more elastic and the carriers face closer substitutes, and smaller where economic profit can absorb the charge or strategic considerations discourage giving up capacity. Fares and capacity are two margins of the same decision, with the balance set by conditions in each sub-market.

### How airlines might adjust prices and quantity supplied, under capacity expansion and an increase in charges (assuming an initial capacity constraint)

4.31 Airline responses following capacity expansion and a corresponding increase in charges would vary as a function of the initial market structure and degree of capacity constraint. At one extreme, there is a monopoly where the capacity constraint binds strictly[22] before and after expansion, such that the incumbent airline supply would be held below the profit-maximising quantity it would otherwise choose to supply. At the other, there is perfect competition and no capacity constraint, where entry is costless and all airlines set price at marginal cost. An intermediate set of scenarios (which could result in a variety of outcomes) occurs when capacity is between the monopolist's profit-maximising quantity and perfect competition.

4.32 In the monopoly scenario, the number of seats on offer is fixed by the capacity constraint, and the fare is set by what passengers are willing to pay for that fixed quantity, not by airlines' costs or profit-maximising behaviour. The fare sits above the profit-maximising fare for the monopolist with unconstrained capacity (which is, by definition, above marginal cost). A higher charge raises the cost of a seat, but with capacity still a binding constraint and the fare set by willingness to pay, a rise in marginal cost does not of itself justify a higher fare. Pass-through to fares is therefore zero while the constraint strictly binds, and the charge is absorbed into the scarcity rent, falling on the airline rather than the passenger. Note that the airline's profit-maximising quantity falls as the result of the higher charge. However, so long as capacity remains below the airline's revised profit-maximising quantity, the constraint remains binding and the charge continues to be absorbed by the airline in full.

4.33 Under perfect competition, with free entry and price equal to marginal cost, a charge increase (which would form part of marginal cost) would be passed through to fares in full. The higher fare would in turn reduce the quantity demanded, so passenger numbers would be lower than they would have been without the charge increase. Any airline that chose to absorb the charge and hence price below marginal cost would be loss-making and not able to survive in the long run. Note that, by definition, perfect competition requires no barriers to entry and hence, from a theoretical perspective, can only exist where there is no capacity constraint.

4.34 In the intermediate set of scenarios, capacity exceeds the monopolist's profit maximising quantity but remains below the competitive equilibrium quantity and airlines are assumed to retain a degree of market power (as by definition, perfect competition requires no capacity constraint). Under these scenarios, airlines would behave as set out in the preceding two sections depending on the exact market structure, demand conditions and degree of market power. In general, there would likely be a degree of charges pass-through. All else being equal, pass-through would likely increase as capacity came closer to the competitive

[22] Capacity is defined here as strictly binding when it is below an airline's profit maximising quantity under a given market structure.

equilibrium quantity – though, as noted earlier, pass-through could be high even when a small number of firms are present in the market.

4.35 It is worth noting that, as capacity increases (from representing a binding constraint through to no longer being a constraint), market structure and airline behaviour are likely to change due to the increasing threat and likelihood of entry. In turn, as market structure changes towards perfect competition, scarcity rents would be expected to fall (with corresponding reductions in ticket prices) whilst the proportion of charges pass-through would be expected to increase. In the limit, if/when capacity expands beyond the perfectly competitive quantity and is no longer a constraint, prices would be expected to fall, in the aggregate, by the existing scarcity rents minus the overall increase in charges. In practice, the market structure and degree of capacity constraint, before and after expansion, will vary across sub-markets. Therefore, different sub-markets will experience different degrees of pass-through, depending on the factors above.

4.36 A final point to note is that the net effect of the capacity constraint being relaxed (which benefits consumers) and the increased pass-through (which is to their detriment) requires closer consideration of the relative scale of scarcity rents and incremental charges, as well as the likely degree of pass-through. But unless the aero charge increase and associated pass-through to passengers is so large that it exceeds the level of scarcity rents, consumers can be expected to benefit in net terms from a capacity expansion accompanied by an aero charge increase.

# 5 Assessment of stakeholder evidence

## Introduction

5.1 The CAA is developing its position on the case for, and the regulatory treatment of, Heathrow expansion. A central question for that work is whether airfares at Heathrow contain a scarcity rent: an element of the fare that reflects the airport being full, rather than the cost of carrying the passenger. Whether such a premium exists, who captures it, and whether expansion would reduce it (or prevent it from increasing further) bears on both the case for expansion and the regulatory treatment of the charges that would fund it. The CAA has asked us to review the stakeholder evidence on these questions.

5.2 This review covers four sets of stakeholder material provided to us by the CAA:

- IAG's observability paper ("OBS"): "Empirical Tests of the Starting Scarcity Rent Hypothesis at Heathrow: Submission to CAA and the Department for Transport (DfT) for review and consideration" (March 2026);
- IAG's slot-values paper ("SVP"): "Slot Asset Values Exist Without Excess Airline Returns: Lessons Provided from Outside Aviation" (May 2026);
- the Oliver Wyman report ("OWR"): "LHR Affordable Capex Envelope: Report (Phase I)" (March 2026); and
- the HAL and Frontier response: a letter from HAL to IAG of 9 April 2026 and an accompanying Frontier Economics note responding to the observability paper.

The abbreviated report names (OBS, SVP, OWR) are shorthand used in citations in this report, for example OBS, Table 2.

5.3 IAG commissioned the Oliver Wyman report, and HAL commissioned the Frontier Economics work that accompanies its April letter.

5.4 Our assessment distinguishes three questions: whether economic theory supports the possibility of a scarcity rent at Heathrow; whether the submissions provide empirical evidence on its existence, scale or incidence (that is, where in the supply chain any rent is captured: by airlines, the airport, suppliers or passengers); and whether the conclusions each party draws are proportionate to the strength and limitations of its evidence. These are assessed against the conceptual framework in Section 4. We have not undertaken new empirical analysis, so our assessment is limited to the internal logic, theoretical consistency and evidential weight of the submissions.

5.5 For each set of material, we summarise the arguments made, assess each in turn against the framework, and note material set aside as contextual or addressed elsewhere in this report. Our position on each paper, our conclusions across the submissions taken together, and the implications for the CAA's consideration of the scarcity rent question are drawn together in Section 6: Synthesis and comparative assessment.

## IAG, “Empirical Tests of the Starting Scarcity Rent Hypothesis at Heathrow: Submission to CAA and DfT for review and consideration”, March 2026

5.6 IAG’s March 2026 paper tests whether the scarcity rent at Heathrow that Frontier[23] estimated for HAL is visible in airlines’ actual financial results.[24] IAG argues that a scarcity rent of that size would have to surface somewhere: as extra profit kept by the airlines, as inflated costs that absorb it, as unusually full aircraft if slot scarcity resulted in rationing of available seats, or as higher fares paid by passengers. It tests each of these four channels in turn and concludes that the rent appears in none of them, and that this absence is decisive. In IAG’s words, “the simplest explanation is that it does not exist.” (OBS, para 1.7)

5.7 We assess each of IAG's tests against our conceptual framework. Three propositions follow from that framework and bear directly on IAG's approach; we set them out here because they recur throughout the assessment:

- a capacity constraint strengthens the conditions for scarcity rents but does not itself establish that they are being earned: whether a rent arises in a sub-market, and its size, depends on demand relative to the constrained supply;
- a scarcity rent need not show up as profit kept by the airline. It can be absorbed elsewhere along the vertical chain (Section 3, Capacity constraints); where it is captured by suppliers of inputs to airlines, it would appear as higher input costs rather than airline profit. It can be capitalised into the market value of slots, so that it is reflected in the value of an asset the airline holds rather than in its reported margins, a mechanism we set out below (see Our review of SVP: Evidence 4: Virgin Atlantic financing arithmetic). Or it can be used to support a wider network; and
- a scarcity rent, if it exists, would not be spread evenly by a profit maximising airline. It would be concentrated where substitutes are weakest and demand least elastic: on routes served by few carriers with no close alternative, at peak times, and in less price-sensitive segments such as premium cabins and late-booking demand, rather than earned uniformly across every passenger.

5.8 The third proposition matters here in particular, because several of IAG's tests are run at the level of a whole airline or the whole airport. A rent that is real but concentrated would still contribute to aggregate returns, but diluted to the point where it may be hard to distinguish from normal variation between carriers. We return to this when assessing IAG's overall conclusion.

### IAG’s evidence and our assessment

5.9 The paper presents five empirical tests (on profits, costs, how close aircraft are to full capacity, and fares), two methodological critiques of Frontier’s approach, and a set of policy conclusions. We take each in turn, summarising IAG’s argument and evidence before giving our assessment.

[23] Frontier Economics, Estimating the Congestion Premium at Heathrow (2019) and Heathrow Expansion: Cost Benefit Analysis (2025), both commissioned by HAL.
[24] Frontier estimated a starting scarcity rent captured in fares at Heathrow of £2 billion in 2016, rising to at least £3.5 billion by 2024 (OBS para 1.2).

### Evidence 1: Carrier margins test (OBS, Section 5)

5.10 In the test referred to by IAG as the carrier margins test, IAG examines whether the two largest carriers at Heathrow, British Airways ("BA") and Virgin Atlantic, earn above a normal level of profit, and whether any excess is consistent with the scarcity rent Frontier attributes to them. The test runs in two stages.

5.11 In the first stage, IAG sets a profit benchmark. It starts from each airline's cost of capital – the blended return its investors and lenders require on the money tied up in the business – measured as a weighted average cost of capital ("WACC"), drawn from three sources (IATA, Bloomberg, and a broker-sourced WACC). It then converts cost of capital into an operating margin (operating profit as a percentage of revenue) that the airline would have to earn to cover it, using a model of the capital tied up in each airline's fleet and the revenue each aircraft generates. IAG treats the result, which it puts at roughly 10 to 12% operating margin, as the margin each airline must earn simply to make a normal profit.

5.12 In the second stage, IAG makes two comparisons against that benchmark:

- whether each airline's actual operating margin exceeds the benchmark, as a test of whether it earns more than a normal profit (i.e. whether it earns economic profit); and
- whether any economic profit matches the scarcity rent estimated by Frontier, when allocated to individual carriers. IAG allocates Frontier's estimated scarcity rent across carriers in proportion to each carrier's share of direct (point-to-point) passengers.

5.13 IAG reports that the two carriers considered fall well short of the profit the Frontier scarcity rent would imply.[25] Our assessment of each stage follows.

**The profit benchmark**

5.14 Measuring whether a firm earns economic profit by comparing its return on capital against its cost of capital is a standard approach, and the WACC figures IAG starts from appear to be drawn from reputable sources.

5.15 However, IAG's approach converts a benchmark cost of capital into an equivalent operating margin, and the steps in that conversion are not disclosed. A cost of capital is a return on the money invested in the business; an operating margin is profit as a share of revenue. Converting one into the other requires assumptions, about how much capital is invested, how that capital maps to the fleet, and how much revenue each aircraft earns – assumptions that IAG does not set out (OBS, para 5.7). The sensitivity of the resulting benchmark to these assumptions is also not shown, and conversions of this kind can be highly sensitive to them. In addition, it is not clear to us why IAG builds a margin benchmark rather than comparing return on capital against cost of capital directly (for example, through Return on Capital Employed), which would avoid the conversion altogether. Because IAG's conclusion that the carriers earn no economic profit rests on its constructed benchmark, we are not able to verify the reliability of this analysis. The point applies with most force to BA, where the result depends heavily on where the benchmark is set; the position is clearer for Virgin Atlantic, whose absolute margins have been low for many years on any measure.

**The comparison against Frontier's estimate**

5.16 The second stage allocates Frontier's estimated scarcity rent across carriers in proportion to their passenger numbers. However, a rent that arises because Heathrow is capacity

[25] IAG reports a combined shortfall of £1,333 million against expected EBIT, 71.9% below the premium implied by Frontier's estimate (OBS Table 6, para 5.14), with BA and Virgin Atlantic together accounting for 50.6% of 2024 flights at Heathrow (OBS Table 4).

constrained would be expected to be concentrated where the capacity constraint bites, where demand is least elastic and where competition is weakest, on particular routes, at peak times and in premium cabins (this is the heterogeneity the framework describes; Section 4, Profitability and scarcity rents). When calculated across airlines as a whole, the estimated scarcity rents are diluted, any imprecision in the underlying operating margin estimates is magnified, and the test may fail to detect scarcity rents when they exist. The same assumption also weakens Frontier's own position. Frontier's estimate is an airport-wide figure, expressed as an average premium per passenger, and IAG adopts that per-passenger framing unchanged when allocating the rent across carriers. If any rent is in fact concentrated on particular routes, times and segments, as the framework suggests, then the proportional allocation both overstates the profit uplift IAG should expect to find at BA and Virgin Atlantic, and limits what Frontier's airport-wide average can say about where any rent actually sits and who bears the cost of recovering it. The criticism therefore goes in both directions: it undermines the precision of IAG's test and the interpretability of Frontier's aggregate figure alike. Concentration would not remove a rent from company-level results: whatever a carrier captures remains within its reported returns. The effect is on detectability and on the predicted uplift. A rent concentrated on particular routes and segments is a small share of group-wide results, and may sit within the normal variation in margins between carriers and across years. The uplift IAG tests against is also unreliable. It assumes the whole of Frontier's airport-wide estimate is retained as current airline profit, and that it is spread evenly per passenger. The concentration and incidence points above call both assumptions into question.

5.17 Taken together, the IAG carrier margins test could in principle show that a large scarcity rent, captured evenly and kept as profit by these two airlines, is hard to find in their accounts (although as noted above, extreme care is needed in comparing a ‘balance sheet’ metric like the WACC with a ‘P&L’ metric like operating margin). It cannot show that no scarcity rent exists, because a rent that is concentrated, shared with suppliers, or reflected in the value of slots rather than in profit would not appear in this test even if it were real. The cost channel is tested directly by IAG's unit cost test, which we assess at Evidence 3.

### Evidence 2: Supporting margin tests

5.18 Alongside the carrier margins test, IAG offers three further tests, which it considers to show that a scarcity rent of the size Frontier estimates does not surface in the airlines’ profits:

- the first looks at the trend in BA’s margin. Frontier estimates the scarcity rent to be 75% higher in 2024 than in 2016, yet BA’s operating margin shows no matching upward trend over the same period (OBS, paras 5.16–5.17, Table 7), which IAG reads as inconsistent with a large and growing scarcity rent;
- the second compares the two carriers’ profit per passenger. Virgin Atlantic flies an almost entirely long-haul network, the segment to which Frontier attributes roughly three times the premium per round trip, yet it earns less operating profit per passenger (about £26) than BA (about £45) (OBS, paras 5.18–5.20). IAG treats this as an "independent falsification" of the scarcity rent; and
- the third compares BA against carriers at less constrained hubs (Iberia at Madrid, Aer Lingus at Dublin, and Ryanair), which have matched or exceeded BA’s margins in various years (OBS, paras 6.4–6.7). IAG reads this as showing that BA earns no margin advantage from operating at a constrained airport.

5.19 We consider all three of these tests to be relevant comparisons. However, each is exposed to confounding differences that the comparisons themselves cannot control for, and we therefore consider it difficult to draw firm conclusions from them in either direction.

5.20 Regarding the first test, if the scarcity rent grew by three-quarters as estimated by Frontier, and this was absorbed by carriers, all else equal, one would expect BA's margin to have increased. However, there may have been confounding changes over this period. For example, this period includes the pandemic, which had multiple material impacts on the airline business.

5.21 The second test is based on differences between proportion of long-haul flights between BA and Virgin Atlantic. However, there may be many further confounding differences between the airlines which could impact margins, including scale of airline, network maturity, route mix and customer base, which limits the weight the comparison can bear.

5.22 The third test, the cross-hub comparison, similarly may conflate other differences between Heathrow and Madrid or Dublin. For example, carriers at Madrid and Dublin may face different costs, competition and demand.

### Evidence 3: Unit cost test

5.23 In the unit cost test (OBS, Section 7), IAG asks whether the scarcity rent could be hidden in the airlines' costs rather than their profits: whether carriers at Heathrow simply spend the rent on a higher cost base, so that it never shows up as margin. IAG compares the airlines' unit costs (operating cost per available seat kilometre, the headline metric used by IAG) at Heathrow against the comparator airports selected in Frontier's 2019 study, using the Skailark flight-level cost dataset for calendar year 2024.[26] It finds costs higher at Heathrow and attributes the gap mainly to higher airport charges and less dense, more premium cabin configurations (OBS, Tables 8–9).[27] It then compares each carrier against its head-to-head competitors on a like-for-like basis (cost per equivalent seat kilometre,[28] excluding fuel) and finds BA and Virgin Atlantic to be, respectively, 11.8% and 7.8% more cost-efficient than those competitors (OBS, Tables 10–11). IAG concludes that higher costs at Heathrow are structural, explained by higher charges and differences in cabin configuration, and do not conceal a scarcity rent quietly absorbed into the cost base.

5.24 We agree that, if a scarcity rent were being absorbed into costs, one place it could be present is in inefficiency, and IAG's finding that its carriers are more cost-efficient than their direct competitors is relevant evidence. However, three limitations apply.

- First, comparing the relative efficiency of airlines is inherently difficult, and unit cost comparisons of this kind carry significant limitations even when fully disclosed. A single-year snapshot does not account for the fact that airlines operate long-lived assets and sit at different points in their investment cycles, so a carrier's costs in any one year partly reflect the age and financing of its fleet rather than its operating efficiency. Unit costs also capture only the cost side: they say nothing about the quality of the product delivered for that cost, which differs materially across carriers and cabins. In addition, the adjustments required to place different airlines on a comparable per-seat basis rest on assumptions that are not fully disclosed, and the analysis is IAG's own construction on the underlying data. Taken

[26] Skailark Aviation Intelligence (www.skailark.com); IAG cites the 17 January 2026 release of the Skailark cost model (OBS, Table 12).
[27] A more premium cabin spreads the same operating cost over fewer seats, mechanically raising cost per available seat kilometre.
[28] Equivalent seat kilometres normalise each aircraft's capacity to a standard high-density configuration, so that unit costs measure the cost of operating the aircraft independently of the airline's chosen cabin layout.

together, we consider that unit cost comparisons of this kind are unlikely to support firm conclusions about the true relative efficiency of airlines, and we treat IAG's head-to-head findings as indicative rather than conclusive.

- Second, being more efficient than a direct competitor does not rule out a scarcity rent shared by all carriers at the airport. IAG's head-to-head comparison is conducted at Heathrow, matched by route and aircraft type (OBS, Tables 10–11), so both sides of the comparison reflect Heathrow-deployed operations. If scarcity value were passed into input costs common to everyone operating at Heathrow, for example into wages, every carrier's costs would carry it and it would net out of both sides of the comparison.[29]
- Third, IAG's own decomposition attributes much of the Heathrow cost premium to airport charges – roughly 60% of the short-haul gap and 44% of the long-haul gap (OBS, Tables 8–9). On IAG's attribution, the higher costs at Heathrow largely reflect charges rather than airline inefficiency, which would address one version of the cost channel that the rent is dissipated through wasteful spending, but not the others. Heathrow's aeronautical charges are subject to a price control that limits the airport's returns, which should in principle prevent scarcity rent from being embedded in charges. We have not assessed whether this holds in practice, as this is beyond the scope of this report. The test therefore narrows where a rent could sit within airline costs. It does not rule out rent sitting in input costs shared across all carriers at the airport, which a comparison between carriers at the same airport cannot detect.

### Evidence 4: Seat factor test

5.25 In the seat factor test (OBS, Section 8), IAG asks whether scarcity shows up in how full the aircraft are, on the reasoning that if capacity at Heathrow were being rationed by intense demand, its aircraft should be unusually full. IAG finds that Heathrow's 2024 system seat factor (the network-wide ratio of seats sold to seats flown, 80.9%) is below Gatwick's (84.8%) and ranks seventh of the 10 airports in Frontier's comparator set (OBS, Tables 15–17). It reads this as inconsistent with scarcity-driven demand pressure.

5.26 We agree this evidence is inconsistent with scarcity showing up as unusually full aircraft. However, the theory does not predict that it would. The constraint binds on the number of movements, not on seats sold. Seat numbers retain some flexibility through aircraft size and configuration (Section 3, Capacity constraints), and the share of seats filled is a choice made through revenue management rather than a direct measure of demand pressure. A carrier facing excess demand at a constrained airport raises fares and manages the cabin mix rather than simply filling every seat, and holding seats back for late-booking, higher-yield passengers is itself part of how a premium is extracted. Scarcity would therefore be expected to appear in yields and in the mix of what is sold, rather than in load factor. Scarcity may also be concentrated at particular times of day and seasons of the year, which an annual system-wide load factor dilutes. There is also a compositional point: the comparison is not between like-for-like operations. Gatwick's traffic is weighted towards low-cost and leisure charter carriers, whose business models are built around high aircraft utilisation and load factors. Heathrow's mix is weighted towards full-service and long-haul network flying, where lower load factors partly reflect scheduling for connectivity and premium cabin economics. A Heathrow to Gatwick load factor gap therefore partly reflects differences in the markets served and the

[29] A separate validation exercise in IAG's paper is consistent with this possibility. There, IAG compares each foreign carrier's unit costs at Heathrow against that same carrier's costs across the rest of its own network, finding Heathrow costs approximately 7.8% (long-haul) and 21.8% (short-haul) higher (OBS, Tables 13–14). IAG attributes this to Heathrow being a structurally more expensive operating environment; the same evidence is equally consistent with scarcity value partly absorbed into input costs shared across the airport.

business models operating at each airport, rather than differences in scarcity. A load factor that is merely average is consistent with scarcity being managed through these levers. The test rules out one form scarcity could take, but does not rule out scarcity rents being present.

### Evidence 5: Fares test

5.27 In the fares test (OBS, Section 9), IAG tests the fare channel directly, asking whether passengers actually pay more to fly from Heathrow. To isolate the effect of the airport itself, IAG uses matched comparisons: the same carrier's fares on the same route, flown from both Heathrow and Gatwick, drawn from the OAG (a provider of airline schedule data) and INFARE (a provider of observed airfare data) datasets. After stripping out the part of the fare that reflects Heathrow's higher airport charges, it finds no systematic Heathrow fare premium for overseas carriers (OBS, Figures 3–4). Where a premium does appear on BA's routes at Heathrow, IAG attributes it not to scarcity rents but to hub revenue management. Seats on BA's short-haul services at Heathrow are demanded both by point-to-point passengers and by passengers connecting onto long-haul services. The lowest fare BA will accept from a point-to-point passenger (its "bid price") therefore reflects what the seat would earn as part of a connecting itinerary. Where connecting demand is strong, this raises the fares charged to point-to-point passengers at Heathrow relative to Gatwick, where no connecting traffic competes for the seat. IAG illustrates this with a decomposition of a JFK–LHR–AGP itinerary, showing that the implied value of the Heathrow–Malaga leg within connecting itineraries tracks closely to the standalone fare on that route (OBS, paras 9.8–9.11, Table 19).

5.28 Fares are one of the channels through which a scarcity rent, if present, would be expected to appear. We consider the design of IAG's test reasonable in principle, but several features limit the weight we consider should be placed on this evidence:

- The "Heathrow fare premium" (higher fares for the same route at Heathrow than at a comparator) and the "BA premium" (the premium specific to BA's routes) are not separate quantities to be offset against one another. Both are candidate indicators of the same underlying scarcity rent. The finding that one carrier exhibits a scarcity rent while others do not is therefore a matter to be explained rather than a finding that scarcity rents do not exist.
- The results found by IAG are highly heterogeneous across routes. The fare differential narrows on many routes but not all, and varies by route type by amounts that can be material against the per-passenger premia Frontier estimates. IAG's finding of close to zero on average therefore does not eliminate the possibility of scarcity rents present on some routes.
- The adjustment driving much of the result is the removal of the airport-charge component of the fare. This adjustment is not fully explained, and the result is sensitive to it. More fundamentally, the adjustment reframes the question rather than settling it: if fare differentials were fully explained by airport charges, this would show that passengers pay more at Heathrow, with the difference accruing to the airport through its charges rather than to airlines through fares. Heathrow's charges are subject to a price control that should in principle prevent scarcity rent being embedded in them; but as noted at Evidence 3, we have not assessed whether this holds in practice.
- The matched routes are a selected subset, smaller than the set of routes underlying Frontier's econometric fare analysis.[30] By construction, the test covers only routes a carrier

[30] Frontier Economics, Estimating the Congestion Premium at Heathrow (2019), econometric analysis of ticket prices, and the route set underlying the fare analysis in the 2025 CBA. The reference is to Frontier's econometric fare work rather than the slot-value analysis assessed elsewhere in this report.

operates from both Heathrow and Gatwick, so IAG's airport-wide conclusion rests on the assumption that these routes carry the same mix of rents as routes served only from Heathrow, an assumption the test cannot itself verify. The direction of any difference is predictable: routes served only from Heathrow by definition offer passengers fewer substitutes, a strong reason to expect higher fares and larger scarcity rents on them, all else equal. If so, the method would underestimate scarcity rents.

- We also note that Gatwick is itself slot-constrained, albeit less severely than Heathrow, and may carry a scarcity premium of its own. A matched comparison differences away any rent common to both airports, so the test measures the Heathrow–Gatwick differential rather than the premium relative to an unconstrained benchmark, and would understate the absolute rent at both airports accordingly. This is consistent with the observation we make on Frontier's use of Gatwick as a competitive proxy in our assessment of the HAL and Frontier response.

5.29 In addition, IAG presents revenue management and scarcity as competing explanations for the BA premium. At a constrained hub, however, the two are not cleanly separable. The bid price reflects the value of a seat in its best alternative use, as part of a connecting itinerary. Where capacity is fixed, that alternative value is elevated precisely because seats are scarce: if BA could add flights freely, the opportunity cost of selling a seat to a point-to-point passenger would be lower, and so would the fare floor. Hub revenue management is therefore a mechanism through which scarcity can be priced into fares, not an explanation that excludes it (Section 4, Pricing effects and pass-through; Capacity rationing and allocation decisions). Demonstrating that the premium can be described in revenue-management terms does not, in our view, demonstrate that it is not a scarcity rent. We also note that IAG's revenue management account rests on a single illustrative route rather than a systematic test.

5.30 None of this changes the finding by IAG that variation in airport charges explains a substantial share of the fare differences observed, which is a relevant result. It does mean the test cannot distinguish whether the residual premium on BA’s routes is scarcity or network value, and any future fare analysis would need to control for charges at both airport and route level and test the two explanations against the competitive conditions on each route, rather than treating them as alternatives.

### Evidence 6: Methodological critiques of Frontier

5.31 In its methodological critiques (OBS, Sections 3, 5), IAG criticises the reliability of Frontier estimate of scarcity rents. First, it argues that Frontier’s method, a regression that treats the unexplained part of Heathrow fares as the scarcity rent, is exposed to omitted-variable bias, which Frontier acknowledges: any cost or product difference the regression does not control for is swept into the estimated premium, overestimating it (OBS, para 3.7). Second, it argues that the IATA Direct Data Solutions (“DDS”) fare data Frontier relies on is deliberately randomised in markets with limited head-to-head competition, with individual fares perturbed by up to 20% (a factor of 0.8 to 1.2), which could undermine any analysis built on fares at that level of detail (OBS, para 5.12, Table 5).

5.32 We agree that both points raise valid questions about the reliability of Frontier’s estimated scarcity rent, and the DDS randomisation point in particular is specific, checkable, and warrants a direct response from Frontier or HAL. However, while these issues may indicate that Frontier’s estimate is overstated or based on imperfect data, they do not, in our view, demonstrate that the true scarcity rent is zero. IAG offers no alternative estimate at the route,

slot-time or segment level, which is where we would expect a scarcity rent to show if it exists. We consider this critique by IAG as motivating further analysis of scarcity rents.

### Evidence 7: Policy conclusions

5.33 IAG draws three conclusions on scarcity rent from the tests above. These bear on the DfT's appraisal of Heathrow expansion, which counts removal of any scarcity-rent-driven fare premium as part of the benefit case for additional capacity. On this basis, IAG concludes that "the simplest explanation is that it does not exist," (OBS, para 1.7) and that the scarcity rent should be removed from the appraisal unless it can be reconciled to the airlines' realised profits. It also concludes that its failure to appear across the tests shows the appraisal methodology to be flawed (OBS, paras 1.8, 10.1–10.4).

5.34 We agree with part of what the evidence shows, but not with the strength of the conclusion. IAG's tests support two propositions: that a scarcity rent of the size Frontier estimates is not visible as profit kept by the two largest carriers, and that the tests, run as they are at the level of a whole airline or the whole airport, do not reveal it. However, we do not consider this to demonstrate scarcity rents at Heathrow are zero.

5.35 To move from the conclusion that a scarcity rent cannot be seen in these particular tests to the conclusion that no scarcity rent exists, we consider IAG would need both to establish one of two conditions and to exclude a third possibility. The two conditions are:

- that Heathrow is not in fact capacity-constrained, so that there is nothing for a scarcity rent to arise from. This is not the case: the airport has been at capacity for two decades, which IAG does not dispute; or
- that competition at Heathrow is strong enough to compete any scarcity rent away. At a fully slot-constrained hub, where new entry depends on the availability of slots, this is difficult to sustain.

The third possibility is that scarcity rents exist but appear somewhere other than airline operating profit. IAG would need to exclude this rather than establish it, and its tests do not. We consider this plausible given the extent of the capacity constraint at Heathrow and our theoretical review of the market. Rent may be capitalised into the market value of slots, shared with suppliers of inputs such as labour, or used to support the wider network. None of these would present as elevated operating margins in a test of IAG's design.[31]

5.36 IAG's tests address the first explanation only in passing and do not address the second or third. They establish that a large, evenly captured, airline-retained rent is difficult to identify in the airline accounts; they do not rule out a rent that is concentrated on particular routes and slots, shared with others, or held in the value of slots rather than taken as profit. The reasonable inference from the evidence is therefore that the scale and incidence of any scarcity rent are uncertain, not that the scarcity rent is absent. On that basis, the evidence does not, in our view, justify removing the scarcity rent from the appraisal, though it is a legitimate reason to test its scale carefully with further analysis, and to seek evidence at route, slot and segment level.

[31] The slot-value channel operates differently depending on how the slot is held. Where a slot was purchased, the expected rent was paid to the seller upfront and sits in the buyer's invested capital rather than its margins. Where a slot is held under historic rights, the carrier forgoes the income it could earn by selling or leasing the slot, so a normal reported margin can coexist with a substantial rent attaching to the slot. One way to test this explanation would be to compare airline returns against a capital base that includes the market value of slot holdings.

5.37 We also note a temporal mismatch between IAG's evidence and the appraisal it is directed at. IAG's tests are necessarily historical, whereas the DfT's appraisal covers a period extending some 60 years beyond scheme opening, over which DfT modelling projects the capacity constraint to increase in the absence of expansion. Evidence on the historical scale of any rent, even taken at its strongest, does not by itself establish the path of scarcity over the appraisal horizon.

#### Other points noted but not assessed further

5.38 This subsection records, for completeness, the material in IAG's observability paper that we have not assessed in detail, because it is contextual background rather than evidence on the existence, scale or incidence of a scarcity rent. The paper also contains: the argument that an airline-captured rent narrative requires a high evidential standard given the regulatory asymmetry between HAL and the airlines (OBS, paras 2.2, 3.10); the distinction between a slot constraint and a passenger-capacity constraint (OBS, paras 4.2–4.3, Table 1), which is developed further in the slot-values paper (SVP) and assessed in our review of that paper; the evolution of Heathrow's passenger mix (OBS, Figure 1); the industry insolvency record (OBS, para 5.4); and the recommendations to the CAA and DfT at OBS, Section 10, which are policy asks rather than evidence and are considered in Section 6.

### IAG, “Slot Asset Values Exist Without Excess Airline Returns: Lessons Provided from Outside Aviation”, May 2026

5.39 This paper challenges the slot-value strand of Frontier's evidence: the inference, made in Frontier's 2019 study ("Estimating the Congestion Premium at Heathrow", May 2019) and relied on in the 2025 cost-benefit analysis ("Heathrow Expansion CBA", August 2025), that observed Heathrow slot transaction prices corroborate an annual passenger-paid scarcity rent (£2 billion for 2018, rising to £3.5 billion in 2023). IAG characterises this triangulation as "the linchpin of the entire methodology" (SVP, para 1.2) and argues that Frontier has mistaken the price of a scarce asset for a fare premium.

5.40 The paper is primarily a methodological and conceptual critique. Its new empirical content is limited to two elements:

- a compilation of publicly reported Heathrow slot transactions between 2013 and 2025, which IAG uses to show how thin the traded sample underlying Frontier's calibration is (SVP, Table 1); and
- an arithmetic exercise extrapolating from Virgin Atlantic's 2025 slot-backed financing to an implied value for the airport's slot portfolio, and from that to the annual income such a value could support (SVP, paras 2.26–2.30).

Where the paper tests whether the claimed rent is visible in profits, costs, load factors and fares, it restates the evidence of IAG's observability paper rather than presenting new analysis. That evidence is assessed in the preceding section rather than re-assessed here.

#### IAG’s evidence and our assessment

5.41 The paper comprises five substantive arguments, a restatement of the observability paper’s tests, and a set of asks of the CAA, taken in turn below. Read against our conceptual framework, the claim that no rent exists must rest on the capacity constraint not in fact restricting supply, on effective competition or contestability, or on the scarcity rent accruing

elsewhere (Section 4, Summary of implications). IAG argues in substance that the value accrues to the slot asset, which our framework treats as a capitalised return to the scarce input, not as evidence of no scarcity rent.

### Evidence 1: Category error

5.42 For assessing scarcity rents, Frontier estimates a congestion premium, calculated as the amount by which Heathrow fares exceed long-run marginal cost ("LRMC") because the airport is capacity-constrained, paid by passengers and accruing to whoever holds the slot.[32] Because that premium is not directly observable, Frontier infers it from slot prices – treating the price an airline pays for a slot as the capitalised present value of the future premium it expects to recover, and then dividing by an annuity factor to express it as an annual figure (SVP, paras 1.1–1.3, 3.1).

5.43 In IAG's category-error argument (SVP, Sections 2–3), it states that a slot price capitalises a bundle of benefit streams (the access right itself, redeployment across routes and gauge, tradability, network feed, schedule reliability, strategic optionality and buyer-specific value) and cannot be attributed wholly to a route-level congestion premium (SVP, Figure 1, paras 2.11–2.12, 3.8–3.9). IAG states Frontier attributes the entire price to one stream, spreads it over a ten-year asset life and "renames the result a congestion premium" (SVP, para 1.5).

5.44 We agree with the nuance introduced by IAG, though only in part. The slot transaction price may represent an upper bound on the congestion premium for the respective slot, since the price paid may reflect value streams beyond the passenger-paid economic rent alone. For example, part of the price may reflect the option to redeploy the slot to other routes or larger aircraft as the holder's network evolves, or the strategic value of preserving access to the airport where re-entry is uncertain. Neither is a premium currently being paid by passengers. However, an airline's willingness to pay for a Heathrow slot is fundamentally driven by scarcity: the access right, redeployment optionality and strategic positioning that IAG identifies as separate benefit streams largely derive their value from the capacity constraint itself and would not command the prices observed at an unconstrained airport. These streams are therefore not cleanly separable from the economic rent in the way IAG's decomposition implies.

5.45 We also note that IAG's network-feed argument and Frontier's estimate operate at different levels. IAG's point is that a slot's value cannot be pinned to the single route flown from it, because a short-haul slot may earn its keep by feeding passengers onto long-haul flights. That is true, but it does not put the value outside Frontier's estimate, which covers Heathrow as a whole: if the slot's value comes from feeding Heathrow's long-haul services, the money is still being paid by passengers flying through Heathrow. The argument applies to the extent that a slot's value derives from revenue arising outside Heathrow, which IAG suggests may be the case for feeder value attaching to the London market more broadly. In our view, the gap between the observed slot price and the congestion premium is likely to be limited. Most of the value streams IAG identifies are expectations of future earnings from holding the slot, which the constraint protects; they are not sources of value unrelated to any premium. On that basis, these streams affect when and how any premium would be earned, not whether the price reflects one. The precise magnitude of any gap nonetheless remains uncertain, not least because converting a capital value into an annual flow depends on the appraisal horizon assumptions discussed at Evidence 3. We therefore treat IAG's category-error argument as a

[32] Frontier Economics, Estimating the Congestion Premium at Heathrow (2019), which defines the congestion premium as the amount by which fares at a congested airport exceed long-run marginal cost, and infers it from slot prices on the basis that an airline paying for a slot must expect to recover a premium above LRMC (pp. 36–46). The 2025 CBA relies on the same slot-value triangulation (Heathrow Expansion: Cost Benefit Analysis, 2025)

reason to interpret Frontier's estimate with care, not as evidence against the existence of a congestion premium.

### Evidence 2: Wrong benchmark

5.46 In its benchmark argument (SVP, Section 3), IAG argues that a fare premium above long-run marginal cost ("LRMC"), as Frontier defines it, is not the same as an economic rent. LRMC excludes enterprise-level costs such as head office, brand, technology and distribution, so fares above LRMC may simply be recovering costs that the benchmark leaves out. IAG argues that the appropriate test is instead whether returns exceed the cost of capital across the cycle, a test it states Heathrow operators fail, citing a blended 9.7% operating margin against a required margin it puts at approximately 10.5% (SVP, paras 3.11–3.15).

5.47 We agree in part. We agree with the distinction between a fare-cost margin and an economic rent above the full cost of capital; our framework draws the same distinction, directing that reported returns be adjusted for the opportunity cost of capital and for slot values before any conclusion on excess returns is drawn (Section 4, Summary of implications). Two qualifications apply.

- First, the test IAG proposes imports its own required-margin construction and the Oliver Wyman 9.7% estimate, neither of which we can validate as the underlying methodology is not fully disclosed.
- Second, for the purposes of any cost-benefit analysis of expansion, the relevant question is whether fares exceed those of an unconstrained counterfactual, which is distinct from whether airlines retain a rent: a premium absorbed into costs or capitalised into asset values could still imply passenger-facing price effects that expansion would alter (Section 4, Pricing effects and pass-through). IAG's benchmark argument narrows what slot values show about airline profits; it does not resolve the counterfactual fare question on which the potential passenger benefit of expansion rests.

### Evidence 3: Calibration errors

5.48 In IAG’s calibration argument (SVP, Sections 2–3), it considers there to be five errors which inflate Frontier’s estimate:

- (i) restrictive sample, with the traded sample averaging just 2.1 daily slot pairs a year, about 0.3% of slots in use, with over 99% held under historic rights and never market-tested (SVP Table 1, para 3.18);
- (ii) the exclusion of the 2012 acquisition of British Midland International (“bmi”), the largest in the public record and at a far lower implied per-slot value (SVP, paras 2.10, 3.19);
- (iii) the assumption of a ten-year slot asset life. Frontier converts slot prices into an annual premium by spreading each price over an assumed ten-year life. IAG argues that slots are in practice held indefinitely under the 80/20 "use it or lose it" rule, so the annual premium should instead be derived as a perpetuity; spreading the same slot price over ten years rather than in perpetuity mechanically produces a larger annual figure, which IAG calculates inflates the implied rent by approximately 94% at Frontier's discount rate (SVP, paras 1.9, 3.20–3.21);
- (iv) the absence of secondary trading of slots between carriers at other hubs is a regulatory choice, and not evidence of absent scarcity at Heathrow (SVP, para 3.22); and
- (v) in a thin, fixed-supply market the observed price reflects the marginal buyer (i.e. the highest-valuing bidder) rather than the average slot holder, so the few transaction prices overstate the value of the slot base as a whole (SVP para 3.23).

5.49 We agree with IAG that Frontier's calculation may be an overestimate of the congestion premium. However, based on this, we consider the magnitude of the overestimation to be unclear, and consider it possible that a congestion premium is present and material. We assess each point in turn.

5.50 On (i), the sample's thinness limits the applicability of the Frontier estimate to the broader airport, for most of the slots are held administratively rather than purchased (Section 4, Profitability and scarcity rents).

5.51 On (ii), the bmi treatment cuts both ways. The transaction is the largest in the record by number of slots, but because the purchase price also covered bmi's wider operating business, its implied per-slot price is the lowest in the record. Excluding it removes that low per-slot observation, which likely overestimates the average slot value; including it at face value attributes the whole purchase price to slots, assuming bmi had no standalone business value of its own, which likely underestimates it, a point IAG itself notes (SVP para 2.10).

5.52 On (iii), we agree in part. Slots are held indefinitely under the 80/20 rule, so a ten-year asset life is likely too short as a description of the slot's economic life, although investors may appraise returns over shorter periods. A longer horizon would reduce the Frontier estimate. However, investors would not in practice appraise returns over an infinite horizon, so a perpetuity is likely too long. The appropriate appraisal period is an open question. Our parallel analysis of slot transactions presents scenarios at 10 and 30 years, agreed with the CAA, as illustrative shorter- and longer-term appraisal horizons. That report does not take a view on the appropriate horizon.[33]

5.53 On (iv), we agree with the non-traded-hubs argument: the absence of slot prices elsewhere is a regulatory artefact, and may be unrelated to scarcity.

5.54 On (v), IAG's argument concerns how prices are formed. In a thin market with fixed supply, the price is set by the buyer who values the slot most, rather than reflecting what slots are worth to the average holder. Observed prices may therefore overstate the value of the wider slot base. We agree in part. As set out in our analysis of slot transactions, a transaction price is negotiated within a range. The bottom of that range is the least the seller would accept: the value of the profit it expects from the slot as currently used. This may be low where the route is marginal or the seller is in financial difficulty. The top of the range is the most the buyer would pay: the value of the profit it expects from its intended use. In a thin market the buyer is likely to be the one who values the slot most, but the agreed price need not sit at the top of the range. The direction of bias when extrapolating to the whole portfolio is therefore not straightforwardly upward. Prices agreed with distressed or marginal sellers may sit below what comparable slots are worth to the airlines holding them. And airlines that hold slots under historic rights, and choose not to sell at the prices observed, may value their slots above those prices. In that case transaction values would understate the value of slots that have never been traded. Extrapolation from the traded sample is therefore uncertain in both directions rather than reliably biased upward.

[33] Grant Thornton, Slot Transaction Analysis to Estimate Economic Profit at Heathrow Airport (July 2026), prepared for the CAA and accompanying this report. The two scenarios (a 10-year horizon at a 10.0% rate of return and a 30-year horizon at 7.5%) and their input assumptions were agreed with the CAA. That report notes that, because of discounting, a 30-year horizon produces results similar to a perpetuity. That report also states that Grant Thornton has not sought to determine whether they are appropriate scenarios from an investor perspective.

### Evidence 4: Virgin Atlantic financing arithmetic

5.55 IAG's financing argument (SVP, paras 2.26–2.30) uses financing secured on Virgin Atlantic's slot portfolio, scaled to the whole airport, as a check on the size of the congestion premium. It scales Virgin's 2025 $745 million slot-backed financing, at an assumed 49% loan-to-value ratio and Virgin's roughly 4% slot share, to estimate the value of airlines' airport-wide slot portfolio at £25 to 30 billion. Capitalised at HMRC writing-down allowance rates (6 to 18%), this implies annual income to airlines of about £1.5 to £5.4 billion (midpoint about £3.45 billion), comparable to Oliver Wyman's estimate of the entire Heathrow airline EBIT pool (£3.1 billion). For Frontier's congestion premium estimate to exist on top of normal returns, IAG states, capitalised income of about £6.5 billion would be required (the £3.1 billion EBIT pool plus Frontier's £3.5 billion premium), which the financing does not support.

5.56 We agree IAG's test is a relevant consistency check. However, we note it rests on a chain of key assumptions: a 49% loan-to-value ratio, scaling Virgin's 4% share airport-wide despite IAG's own acknowledgement that its portfolio is skewed to higher-value slots, and converting asset value to income using HMRC writing-down allowance rates, which are tax depreciation rates rather than capitalisation yields. IAG's calculation also produces a wide range (£1.5 to £5.4 billion), and its conclusion hinges on the midpoint being representative; it could change if the true value sat at the top of the range.

5.57 The financing also confirms that Heathrow slots are bankable assets of substantial value, and that value requires explanation under any framework. The value of an asset derives from the present value of the cash flows its holder expects it to enable. For a slot, those cash flows arise from operating services at Heathrow, or from selling or leasing the slot to another operator whose willingness to pay in turn reflects its own expected operating cash flows. The streams IAG identifies (access rights, redeployment optionality and strategic positioning) are rights or options over such operating cash flows rather than independent sources of value. An option is valuable only to the extent that, in some future state or use, exercising it is expected to yield earnings above the competitive alternative. Access to an airport at which only competitive returns were expected, in every state and use, would not command a substantial price, because equivalent returns would by definition be available elsewhere without payment. IAG does not explain how the streams it identifies would hold value on that basis, and its own illustration derives the affordability of a £40 million slot from the operating earnings generated by carrying passengers at Heathrow (SVP, para 2.14). Its own appeal to Ricardian rent points the same way. In that concept, a scarce and differentiated resource commands a price reflecting its scarcity, while whoever operates it earns only a normal return once that price has been paid: the same relationship between slot value and operating margin that the framework describes as scarcity rent. IAG's submission treats the rent as real and capitalised into the price of the scarce asset, with the operator earning normal returns after paying it (SVP, para 3.25).

5.58 We therefore consider substantial slot asset values difficult to reconcile with the complete absence of scarcity rents. They are more naturally read as evidence that rents are expected to exist over the holding period and are capitalised into slot values rather than appearing in operating margins, consistent with our assessment of the observability paper (Evidence 7) and with the treatment of slot values in our parallel slot transactions analysis. Two qualifications apply.

- First, this reading does not quantify any rent: given the assumptions set out above, the £1.5 to £5.4 billion range should not be relied on in either direction.

- Second, asset values are forward-looking. They evidence rents expected over the holding period, which may reflect anticipated future scarcity as well as any current premium, and the balance between the two cannot be determined from asset values alone.

### Evidence 5: Carrier behaviour record

5.59 In its behavioural argument (SVP, paras 2.23–2.25), IAG states that the commercial failures of bmi and Little Red, two decades of short-haul retrenchment by Iberia, Air France-KLM, SAS, Lufthansa and TAP, and outright exits by airlines, are inconsistent with carriers protected by an extractable airport-wide premium.

5.60 We agree in part. The record is consistent with relatively low short-haul profitability at Heathrow. It is not, however, inconsistent with a premium concentrated in peak long-haul markets: evidence of this kind does not distinguish the absence of a rent from its concentration. This is consistent with the framework's expectation that rents, where they exist, would be most pronounced on routes with less competition (Section 4, Profitability and scarcity rents).

### Evidence 6: Visibility test

5.61 In the visibility test (SVP, paras 1.7, 2.17–2.22, 3.16), IAG restates the observability paper's four-channel evidence and adds no new analysis. Our assessment of that evidence, including the aggregation limitations and the non-independence of the Oliver Wyman estimate, is set out in the preceding section and applies equally here.

### Evidence 7: Asks of the CAA

5.62 In the asks of the CAA (SVP, Section 4), IAG recommends rejecting the 2019 slot-based figure, testing charging assumptions against realised carrier behaviour, requiring HAL to demonstrate that the £49 billion envelope is consistent with airline operating economics, and pursuing regulatory reform to introduce contestability and risk-sharing. These are recommendations rather than evidence and presuppose the conclusions of the preceding analysis. They are considered in our overall conclusions at Section 6.

### Other points noted but not assessed further

5.63 For completeness, the paper also contains: supporting theoretical frames (Ricardian rent, real options, network supermodularity, and thin-market comparators in water, fisheries and dairy quota, SVP paras 3.24–3.28), which reinforce the category-error argument and are not separately assessed; institutional description of slot allocation (SVP, paras 2.5–2.8); the DEFRA farmland figures used in the analogy (SVP, para 1.6); and the forward-looking "death spiral" scenario under higher charges (SVP, paras 1.10, 2.31–2.35), which is an assertion about future behaviour rather than evidence and overlaps with the affordability arguments in the Oliver Wyman report, assessed next.

5.64 One point from the paper is worth drawing out: the Ricardian framing (SVP, para 3.25) accepts that scarcity has economic value, and locates that value in the slot as an asset rather than in operator profits. The disagreement between the parties is therefore narrower than the paper's title suggests. Both sides accept that Heathrow slots carry substantial scarcity value. What is disputed is what that value represents. Frontier treats it as capitalising a premium paid by passengers through fares, which expansion could erode and charges could recover. IAG treats it as asset value that can exist without any such premium, although, as set out in our

assessment of the financing argument above, the paper does not explain the source of such value independent of the earnings that operating at Heathrow is expected to enable. This is the distinction the framework draws (Section 4, Profitability and scarcity rents; Pricing effects and pass-through). A high slot value shows that scarce access to Heathrow is worth paying for. It does not show what fares would be without the constraint, or how much of a charge increase airlines would pass on to passengers. Those are the two questions that matter for the CAA's assessment, and slot values alone cannot answer either.

### Oliver Wyman: “LHR Affordable Capex Envelope: Report (Phase I)”, March 2026

5.65 This report, commissioned by IAG and written by Oliver Wyman (OW), models the effect of expansion-related airport charge increases on airline economics at Heathrow and derives an "affordable capex envelope" for the third runway. Its headline findings are that airlines could absorb £4.50 to £5.50 per passenger in incremental charges before margins fall below a minimum viable level, giving an envelope of up to £30 billion against HAL’s £49 billion masterplan.

5.66 The report differs in kind from the two IAG submissions: it is a forward-looking scenario model rather than an empirical test, its results follow from its assumptions, the underlying model is not provided, and no sensitivity analysis is given for the headline figures. It deliberately considers only airline economics and behaviours, excluding the broader economic, connectivity and passenger impacts of expansion (OWR Section 1.1). That scope limit is clearly disclosed and appropriate for the question addressed, however it means the £30 billion figure speaks to affordability for airlines under the report’s assumptions, not to the welfare-optimal or policy-appropriate scale of investment.

5.67 Our assessment applies three framework propositions (Section 4): cost pass-through and incidence (Pricing effects and pass-through), the heterogeneity of airline economics across routes and segments (Profitability and scarcity rents), and capacity reallocation behaviour (Capacity rationing and allocation decisions).

#### Oliver Wyman’s evidence and our assessment

5.68 The report comprises five substantive arguments, a charge-to-capex conversion, and a set of supporting assumptions, taken in turn below. Because the results follow from the assumptions, our assessment focuses on the assumptions that most affect the headline figures.

### Evidence 1: Industry economics

5.69 In its industry-economics argument (OWR, Section 2.1), Oliver Wyman states that airlines are structurally low-margin: the global industry earned an 8.0% operating margin in 2024, with return on invested capital (“ROIC”) (6.7%) below WACC (8.8%), and a 6.8% average margin over 2010 to 2024 against 12.6% across other sectors (OWR, Figure 1). BA's 14.2% margin is similar to the margins of other airlines within the same ownership group, IAG, that have minimal London presence (Iberia at 13.6% and Vueling at 12.3%), which Oliver Wyman presents as evidence against a "Heathrow windfall".

5.70 We agree that these figures are broadly consistent with public data and with IATA’s published industry economics. However, this information is, on its own, of limited value in determining the existence of scarcity rents at Heathrow. First, comparing operating margins across sectors (OWR, Figure 1) says little about whether any sector's returns are excessive: margins are not comparable across sectors with different capital intensity and asset turnover. A capital-

intensive business needs a higher margin on revenue to earn the same return on the money invested in it, so a higher margin does not by itself indicate better performance or excess returns. Second, aggregate industry economics are an imperfect tool for identifying a scarcity rent. The scarcity rent question asks whether particular carriers, routes or slot times earn above competitive returns. Industry averages dilute that signal across the whole sector and may fail to detect it even when present.

### Evidence 2: London market and Heathrow margins

5.71 In the London-market argument (OWR, Section 2.2), OW treats London as one aviation market served by six airports, with Heathrow's charges per passenger about 2.1 times the weighted average of the others. Oliver Wyman estimates that airlines at Heathrow earned a 9.7% operating margin in 2024 (long-haul 10.9%, short-haul 2.3%), based on its own allocation of airline financial results to Heathrow routes (OWR, Figure 3). It notes this is in line with the global pattern of long-haul margins running 2 to 3 points above system averages, and presents the figures as indicating limited excess profit today (OWR, Figures 3–4).

5.72 We agree in part. The figures are relevant to the scale of any airline-retained scarcity rent and are the same figures IAG imports in the observability section. However, the method for allocating airline EBIT to Heathrow routes is not set out: route-level profitability requires allocations of joint costs and network revenues that are inherently judgement-laden, so the precision of the 9.7% margin, and of the 2.3% short-haul figure on which the redistribution story depends, cannot be fully assessed. The comparison between BA and the other IAG-owned airlines, Iberia and Vueling, has the same limitation as IAG's cross-hub comparisons assessed in the observability section: those carriers' Madrid and Barcelona-based operations face their own market conditions, and similar margins across different markets do not establish the absence of a premium in any one of them (Section 4, Profitability and scarcity rents).

### Evidence 3: Revenue constraint and pass-through

5.73 In the revenue-constraint argument (OWR, Section 2.3), Oliver Wyman states that airline revenue grows in line with nominal GDP, settling at a lower share after shocks (0.73% of European GDP post-pandemic, OWR, Figure 5); that revenue is "already broadly maximised" through revenue management; and that attempts to pass incremental costs to passengers would fail or divert volume to other London airports and competing European hubs.

5.74 We disagree in part, and this is one of the report's most consequential assumptions. Airline revenue has tracked GDP at the level of the European industry as a whole. It does not follow that revenue at a single airport is capped in the same way. Even if industry-wide revenue tracks GDP, revenue can shift between airports within that total, and Heathrow's passenger and route mix could allow its revenue to grow faster than the industry average.

5.75 More importantly, "revenue is maximised" does not mean that additional costs cannot be passed on to passengers. Revenue management extracts the most revenue available given current costs, demand and competition. A charge increase applying to every carrier at Heathrow changes those conditions. The framework describes zero pass-through as an analytical extreme that holds only while a constraint strictly binds and the rent in a sub-market exceeds the charge. In practice sub-markets sit between that extreme and the competitive case, and the charge increases the report models would accompany the capacity release that erodes the rents available to absorb them (Section 4, Implications for a capacity expansion scenario). Recovery would be expected to be greatest in the less price-sensitive premium segments that Oliver Wyman itself identifies as funding Heathrow's charges.

5.76 The assumption can also be considered against the report's own account of the market today. Oliver Wyman reports that Heathrow's charges are roughly twice those of other London airports and that carriers operating there remain profitable (OWR Section 2.2, Figure 4). Within the framework, there are two ways this could have happened: the higher charges have been recovered from passengers through fares, or they have been absorbed into scarcity rents at least as large as the charge differential.

5.77 The first is inconsistent with the model's assumption that no part of a charge reaches fares. The second is inconsistent with the report's own finding that airline margins at Heathrow are in line with industry norms, leaving no excess profit of the scale required. The model assumes no part of any future increase reaches fares at any point to 2045, and does not address which explanation of the historic record it relies on.

### Evidence 4: Scenario results

5.78 In the scenario argument (OWR, Sections 3.1–3.2), Oliver Wyman models that with two runways Heathrow remains movement-constrained to 2045 (i.e. it cannot increase its number of take-offs and landings), with about 26 million passengers of excess demand, and with the operating margins of airlines serving Heathrow remaining broadly flat at about 10%. With a third runway at unchanged real charges, Heathrow reaches 756,000 movements and attracts short-haul traffic back, and airlines serving the airport again earn operating margins of about 10%, on a larger passenger base.

5.79 The scenarios follow mechanically from the demand, capacity and cost assumptions, and their qualitative features (continued constraint and shift in mix to more long-haul flights under two runways, short-haul re-entry under three) are plausible. We disagree, however, with two demand-side assumptions that materially shape the results. First, unconstrained demand for London as a whole in 2045 is held fixed across scenarios (about 270 million passengers with or without a third runway, OW Figure 8): expansion is assumed to generate no additional demand from improved capacity, connectivity or any change in fares. That is not neutral in the present discussion, since assuming no fare-driven demand response implies there is no scarcity rent for expansion to erode. The report thus embeds as an input the conclusion contested between the parties (Section 4, Pricing effects and pass-through). Second, and following from this, all incremental traffic at an expanded Heathrow is assumed to be abstracted from other London airports rather than newly generated, which drives both the excess-capacity finding for the London system (OWR, Section 3.2) and the redistribution dynamics behind the affordability ceiling. Neither assumption is tested or sensitivity checked.

### Evidence 5: Affordability ceiling

5.80 In the affordability argument (OWR, Section 3.3), OW states that airlines could absorb £4.50 to £5.50 per passenger, eroding margins by 2 to 3 points to about 7%. The minimum viable margin is set at 6 to 7%, derived from the 25th percentile of operating margins that airlines reported, then adjusted upward for Heathrow's higher share of long-haul traffic. If charges rise above this ceiling, the model assumes airlines begin moving flights away from Heathrow to other London airports and competing European hubs, because they can no longer earn an acceptable margin there. This creates a self-reinforcing dynamic: as carriers leave, the fixed cost of the expansion has to be recovered from a smaller number of remaining airlines, pushing their charges up further and prompting still more to leave. If, instead, charges are held at the affordable level (the £4.50 to £5.50 the model says airlines can absorb), Heathrow is projected to reach about 122 million passengers a year by 2045, using roughly 630,000 of its 756,000 movement slots – about 82% of the expanded airport's capacity.

5.81 We accept the broad proposition, which is consistent with our framework (Section 4, Capacity rationing and allocation decisions): a large enough increase in charges would be expected to drive some airlines to move capacity elsewhere, and once that starts, a self-reinforcing decline on a shrinking base of remaining carriers is a real risk in principle. The framework's account of capacity responses is more specific than this, however. Airlines optimise across the network and make deployment decisions at the level of individual routes and sub-markets, so the capacity response to a charge increase would be expected to be uneven: marginal routes and price-sensitive segments would shift first, while services supported by network feed, premium demand and hub economics would be retained longer, and part of the adjustment would fall on rents and fares rather than volumes. Assessed against that account, the model's approach of applying a single airport-wide margin threshold to broad carrier groups is a simplification that determines its results.

5.82 We also note that the self-reinforcing decline the model describes is not an automatic consequence of the current regulatory framework, assuming that framework remains broadly as it is today. The CAA's price control sets a maximum allowable yield; it does not require charges to be set at that level. An airport facing the demand responses the model describes would have a commercial incentive to price below the maximum rather than drive traffic away, since on the model's own logic charging above the affordable level reduces the airport's volumes and, with them, its revenues. The spiral therefore requires not only that airlines relocate, but that Heathrow continues to raise charges into declining demand, a course on which the airport would itself bear risk and which it would have no evident incentive to follow. How far this safeguard would operate in practice depends on the design of the control, including where volume risk sits between the airport and its airline customers, which is a matter of regulatory design beyond the scope of this report. The model does not engage with this charging response.

5.83 Moreover, we consider that little weight should be placed on the specific figures: the £4.50 to £5.50 charge, or the 6 to 7% margin floor beneath which airlines are assumed to withdraw. This is for three reasons, set out below.

- First, the way the margin floor is derived does not measure what it needs to. OW sets it at the 25th percentile of airlines' historical margins. But that tells us how often airlines have earned low margins, not the margin at which an airline actually decides to pull flights out of an airport. Additionally, OW's calculation assumes that the share of airlines unable to cover their cost of capital stays constant over the long run, which is stated rather than demonstrated.
- Second, the trigger is set at the wrong level of detail. Airlines decide whether to keep or move flights route by route and aircraft by aircraft – when charges rise, it is the marginal individual routes that get cut while profitable ones stay. Yet the model applies a single airport-wide margin threshold to broad groups of carriers, assuming a whole group moves once its average margin dips below the floor. That averaging blends profitable and marginal routes into one figure and so hides the very variation that determines which capacity would actually shift (Section 4, Profitability and scarcity rents).
- Third, the floor carries forward the zero pass-through assumption we questioned earlier: the model assumes airlines absorb the whole charge in their own margins rather than passing any of it on to passengers. If even part of the charge can be passed through, as the framework expects in practice, and increasingly so as expansion relaxes the constraint (Section 4, Implications for a capacity expansion scenario), the affordable charge rises.

5.84 The margin floor itself raises two further issues. First, it is drawn from the wrong starting point. The right anchor would be the margin airlines actually need to stay invested, built up from a

benchmarked cost of capital, rather than a point on the historical distribution of past margins. We note that IAG's own cost-of-capital figures carry the transparency limitations set out in the observability section. Second, OW's own data shows some airlines operating persistently below the 25th percentile without exiting. This suggests the real withdrawal point sits below the assumed floor. If so, airlines could absorb more than £4.50 to £5.50 per passenger, and the affordable charge is understated.

5.85 Taken together, the threshold is likely imprecise in both directions rather than reliably biased one way. But three of the assumptions pull in the same direction, towards a lower affordable charge and a smaller capex envelope. The first is zero pass-through, held fixed to 2045 even though the modelled expansion would be eroding the rents that absorption requires. The second is that capacity moves away easily once the margin floor is breached. The third is that the airport itself makes no charging response. On the second of these, Heathrow's network effects, connecting (feeder) traffic, corporate contracts and premium catchment all make it costly for an airline to switch airport – a point OW itself concedes for hub carriers over at least the medium term (OWR, footnote 20). Finally, no sensitivity analysis is offered around the £4.50 to £5.50 per passenger figure, so the apparent narrowness of that range understates how uncertain it really is.

### Evidence 6: Charge-to-Capex conversion

5.86 In the conversion argument (OWR, Sections 3.4, A.5), OW converts the £4.50 to £5.50 ceiling into an envelope of up to £30 billion by back solving the CEPA regulated asset base (“RAB”, cumulative value of past investment not yet recovered from users ) charge model for the capex that can enter the RAB each year, given the assumed WACC, depreciation profile and traffic forecasts. The report states that this result depends on how the capital spending is phased: a programme that is front-loaded, or larger overall, would need funding from sources other than the airlines.

5.87 The method itself is conventional, and the report is clear that the result depends on the phasing of the spending, the assumed cost of capital (WACC), the depreciation profile and the traffic forecasts. However, the CEPA model and its parameters are not provided, so the conversion cannot be independently checked. Nor would we treat the “up to £30 billion” as a firm upper bound: it rests on the £4.50 to £5.50 charge, which we have argued is itself uncertain and, on several of the report’s own assumptions, may be understated.

### Evidence 7: Assumptions

5.88 The supporting assumptions (OWR, Appendix A) have revenue growing one-for-one with GDP, zero real price growth, cost lines escalating on specified indices (fuel −0.3% a year, net of efficiency gains and rising sustainable aviation fuel use, crew +1.5%, maintenance +1.2%), passenger demand following the Airbus Global Market Forecast (a weighted 1.9% growth per year), broadly flat aircraft gauge (the seat count per aircraft), and carriers shifting capacity when margins fall below the viability threshold, assessed by airline category (European low-cost carriers, European full-service carriers, non-European private carriers, and non-European state-backed carriers).

5.89 The assumptions are conventional in form and clearly documented. The most consequential items, the one-for-one revenue-to-GDP link, zero real price growth, the fixed-demand and abstraction assumptions, the Airbus forecast and the capacity-shift rule, are assessed above. The cost escalators and the sustainable aviation fuel assumptions appear reasonable, but their sensitivity is untested in the report.

5.90 While the assumptions are documented and can be examined, none of them is tested. No sensitivity analysis is presented, and the assumptions that drive the headline figures are not the cost escalators, which appear reasonable. They are the behavioural assumptions assessed above, namely zero pass-through, fixed London-wide demand, and whole groups of carriers shifting capacity together. The report is therefore best read as a working through of one scenario, not an estimate of what airlines could in fact afford.

### Other points noted but not assessed further

5.91 For completeness, the report also contains: descriptive material on the London airport ecosystem and capacity shares (OW, Section 2.2, Figure 2); benchmarking of seats per passenger across European hubs (OWR, Figure 6), which indicates Heathrow's capacity provision is in line with comparator hubs and is contextual to the scarcity question; the historical shift of short-haul capacity away from Heathrow (OWR, Sections 2.2, A.4), which overlaps with the carrier-behaviour evidence in the slot-values paper assessed in that section; and the observation that excess capacity could emerge across the London system after expansion (OWR, Section 3.2), which is a consequence of the fixed-demand assumption assessed under the scenario argument above rather than an independent finding.

## HAL and Frontier's critique of IAG/Oliver Wyman submissions

5.92 This section reviews HAL's response to IAG's March 2026 empirical paper: a letter from HAL to IAG on 9 April 2026 and an accompanying note prepared by Frontier (referred to here as the covering letter and the Frontier note). The documents respond to two of the reports discussed above, IAG's report "Empirical Tests of the Starting Scarcity Rent Hypothesis at Heathrow", which argues there is no congestion premium at Heathrow, with brief engagement with IAG's accompanying Oliver Wyman report "LHR Affordable Capex Envelope: Report (Phase I)" on Heathrow expansion. HAL and Frontier's documents predate IAG's May 2026 slot-values paper and do not respond to it.

5.93 The documents differ in kind from the submission they address. They are primarily rebuttals, with incremental analysis, combining critique of IAG's evidence, a restatement of Frontier's earlier multi-method estimates of the premium, and a limited amount of new behavioural evidence (market evidence of what airlines pay to buy or lease scarce Heathrow slots, used to infer that they expect above-normal returns from operating there). HAL commissioned the Frontier work that IAG's submissions contest. As with the IAG and Oliver Wyman material, the covering letter and the Frontier note were commissioned by the same party and cross-refer to one another, so we do not treat them as fully independent evidence. Our assessment applies the framework propositions (Section 4) and cross-refers to our assessments of the IAG papers and the Oliver Wyman report.

### HAL and Frontier's evidence and our assessment

5.94 The material makes six substantive arguments directed at IAG's paper, taken in turn below. Where IAG argues the congestion premium is not evidenced, the covering letter and Frontier note argue the opposite, drawing on institutional recognition (that the premium has been accepted by official bodies rather than being a claim unique to Frontier), IAG's own fare analysis, slot-market evidence and the supply-demand logic of a capacity constraint.

### Evidence 1: Not a Frontier-only proposition

5.95 In the institutional-recognition argument, the Frontier note and covering letter state that the congestion premium is not a Frontier-only claim: it has been accepted by the CAA (which described an assumed range of £0.9 to £2.5 billion a year as "conservative"[34]), the Airports Commission, and the DfT (whose updated appraisal Frontier states attributed substantial passenger fare benefits to expansion[35]), and examined in the peer review of Frontier's congestion premium analysis carried out by the Institute for Transport Studies at the University of Leeds ("the ITS Leeds peer review"). Statements by easyJet and analysis commissioned by Virgin Atlantic (from WPI Economics) are cited as airline acknowledgements of the same effect, and Frontier restates its earlier econometric, slot-value and top-down elasticity estimates, each of the order of £2 to £3 billion a year, as converging evidence.

5.96 We agree on this point: the congestion premium concept is embedded in CAA, Airports Commission and DfT analysis, and IAG's framing of the scarcity rent as "Frontier's claimed" hypothesis understates that institutional history. Two qualifications apply. First, the fact that official bodies have accepted the concept of a congestion premium in the past does not itself evidence its current scale or incidence. The cited figures also rest on the same classes of method whose limitations are examined in the submissions assessed above: econometric residuals, slot values, and shadow-cost modelling (the estimate within the DfT's aviation model of the fare uplift needed to ration demand down to available capacity). Second, the airline statements are consistent with an expectation of lower fares after expansion.

### Evidence 2: IAG's own fare analysis

5.97 In its fare-analysis argument, Frontier states that IAG's own data evidences a premium. On matched BA routes, after controlling for airport charges, IAG finds Heathrow fares about £36 higher (European) and £12 higher (long-haul) than Gatwick. Applying these differentials to BA's point-to-point passengers, Frontier illustrates a premium of about £600 million a year for BA alone, around 25% of its 2025 operating profit (Frontier note, footnote 17).

5.98 We agree in part. The charge-adjusted BA differentials are an output of IAG's own analysis, and the interpretive dispute over them is the central empirical question, as set out in our assessment of IAG's fares test at Section 5, Evidence 5. We find Frontier's reading of this data theoretically coherent, subject to one assumption: that Gatwick fares serve as a reasonable proxy for prices in an unconstrained market. This is not clear-cut – Gatwick is itself seeking to expand and may be experiencing a degree of slot scarcity of its own, which implies Frontier may be underestimating the congestion premium at Heathrow. The £600 million congestion premium estimate by Frontier should be treated as illustrative, for three reasons.

- First, the differentials are estimated only on matched routes, but the calculation applies them to all BA point-to-point passengers, including those on Heathrow-only routes where no differential has been measured. That step is only valid if the matched routes are representative of the rest of the network, and both parties agree they are not, disagreeing only on the direction of the resulting bias.
- Second, Frontier itself notes uncertainty over whether the differentials are one-way or return figures.
- Third, the calculation makes no allowance for product or service differences between the two operations beyond airport charges (Section 4, Profitability and scarcity rents).

[34] CAA, CAP1871, Economic regulation of Heathrow Airport Limited: policy update and consultation on the early costs of capacity expansion (December 2019).
[35] DfT, Updated Appraisal Report: Airport Capacity in the South East (October 2017; cited by Frontier and HAL)

These considerations pull in different directions. The use of Gatwick as the benchmark would tend to understate the premium, as would applying matched-route differentials to Heathrow-only routes where, on our framework, any premium would be expected to be larger, all else equal. The absence of any allowance for product and service differences beyond airport charges would tend to overstate it. The uncertainty over one-way and return figures could move the estimate materially in either direction. The £600 million figure is therefore uncertain in both directions rather than reliably biased one way, and should not be relied on as an estimate of the premium on BA's routes, in either direction.

### Evidence 3: Hub revenue management as a mechanism of scarcity

5.99 In its revenue-management argument, Frontier states that IAG's bid-price account explains how BA allocates seats between point-to-point and connecting passengers, but not why the resulting fare gap survives. In a market with spare capacity, competitors would add seats and compete such margins away. At Heathrow, entry and expansion are blocked by the slot constraint, so the margin persists.

5.100 We agree with the conceptual point, which is the stronger version of the observation in our assessment of IAG's fares test. At a constrained hub, seats must be rationed between point-to-point and connecting demand, and the fare is the means by which that rationing happens. A seat goes to whichever use yields more to the network, so the fare a point-to-point passenger must pay reflects the seat's value in its alternative use as part of a connecting itinerary. A connecting passenger is not inherently more valuable than a direct one; what matters is that the seat has an alternative use, and the yield of that alternative sets the floor under the local fare. The constraint is what makes this floor high: because BA cannot add capacity to serve both demands, the two compete for a fixed set of seats, whereas with free entry and expansion the floor would fall towards cost. Scarcity does not change the seat; it changes the price at which the seat is rationed and the mix of passengers who end up travelling (Section 4, Pricing effects and pass-through; Capacity rationing and allocation decisions). On this account, revenue management and scarcity are not competing explanations. The persistence argument, that entry would compete the margin away if capacity existed, is standard theory and well made. What it does not do is decompose the fare gap between scarcity and network value. Frontier itself accepts that on routes that exist only because of connecting traffic, some gap could arise even without a capacity constraint. So the argument challenges IAG's explanation, but does not quantify the premium.

### Evidence 4: Heathrow-only routes

5.101 In its selection argument, Frontier and the HAL covering letter state that the matched-route sample (the set of routes BA flies from both Heathrow and Gatwick to the same destination, paired up so IAG can compare like-for-like) likely understates the premium: routes served from both airports give passengers a partial outside option, whereas on Heathrow-only routes the scope for premia would be greater (Frontier note, p.9; covering letter, point 5).

5.102 We agree this is a plausible argument about the direction of bias, consistent with the selection limitation identified in our assessment of IAG's fares test. It remains a hypothesis, however: no analysis of Heathrow-only routes is presented. And the argument cuts both ways. If matched routes are unrepresentative of the wider network, as Frontier argues, then the differential measured on them cannot be extrapolated to the rest of the network with confidence in either direction (Section 4, Profitability and scarcity rents).

### Evidence 5: Slot transactions, leases and financing

5.103 In its slot-market argument, Frontier and the covering letter state that slot transactions, leases and slot-backed financing are market evidence of expected incremental returns: airlines would not pay tens of millions for slots, or lease rates reported at up to £6 million a year per pair (about £8,000 per flight), unless they expected returns above a normal level; the same implication follows from carriers retaining rather than selling slots, and from using them as loan collateral (the reported $745 million Virgin Atlantic financing). IAG's March paper is criticised for not engaging with this evidence (Frontier note, pp.9–11; covering letter, point 6).

5.104 We agree, on the same basis as our assessment of IAG's May 2026 slot-values paper. The transactions, leases and financing show that scarcity-related value exists and can support secured borrowing. Frontier's net present value reasoning establishes a limited proposition: a purchaser paying a positive price for a slot must expect incremental returns of at least that value. As set out in the slot-values section, the streams sometimes presented as alternatives to a premium, such as network feed and redeployment options, are themselves claims on expected earnings from operating at Heathrow, so we consider the gap between the slot price and the capitalised premium likely to be limited.

5.105 Two qualifications carry over from that section. The evidence does not quantify the premium, and it is forward-looking: prices, lease rates and loan collateral values reflect returns expected over the holding period, which may include anticipated future scarcity as well as any current premium, and they do not show where within the network those returns are expected to arise (Section 4, Profitability and scarcity rents). We also note that the criticism that IAG "does not engage" with slot values, while accurate for the March paper, is overtaken by IAG's May 2026 slot-values paper, which addresses this evidence directly and which these documents predate.

### Evidence 6: Constrained growth and price rationing

5.106 In its supply-and-demand argument, Frontier states that Heathrow passenger growth of about 1.1% a year over two decades, against about 3.5% for UK airports as a whole, indicates substantial excess demand that must be priced out of the market when capacity is fixed; its top-down elasticity method puts the resulting premium at about £2 to 3 billion a year (Frontier note, pp.11–12).

5.107 We agree with the underlying theory, consistent with the framework: prolonged excess demand at fixed capacity would be expected to be rationed, including through prices (Section 4, Pricing effects and pass-through). The growth comparison is consistent with a capacity constraint. However, rationing also operates through a range of channels other than price, including aircraft gauge, cabin mix, the displacement of short-haul by long-haul, and congestion effects on the passenger experience and on resilience, such as crowding in terminals and pressure on approach roads and car parking. The growth comparison therefore does not by itself establish the scale of any price effect, which depends on substitutability with other London airports and on demand elasticities.

### Other points noted but not assessed further

5.108 For completeness, the documents also contain: the DfT's Aviation Modelling Suite shadow cost, cited as a further independent route to the premium and already reflected in the appraisal figures noted above; and the slot-transfer record (Frontier note, Figure 2) and reported slot-lease rates, which are descriptive support for the slot-market argument rather

than separate findings. The airline statements by easyJet and Virgin Atlantic’s commissioned analysis are assessed under the institutional-recognition argument.

# 6 Synthesis and comparative assessment

6.1 This section draws together our assessments of the four sets of stakeholder material, states what the evidence in each can and cannot establish, and sets out our overall conclusions on the two questions this report addresses: (1) the extent to which airlines at Heathrow earn scarcity rents; and (2) how airlines would be expected to respond to increased aeronautical charges. It replaces the separate conclusions in the individual reviews above.

6.2 Taken together, the submissions approach the existence of scarcity rents and a congestion premium on ticket prices from several directions. IAG's observability paper challenges whether the claimed premium is visible in airline profitability, costs, load factors or fares. Its slot-values paper challenges whether slot transaction values can be converted into an annual passenger-paid fare premium. The Oliver Wyman report challenges whether airlines could absorb the charges required to fund expansion. The HAL and Frontier response challenges the replicability and behavioural consistency of IAG's evidence, restates the institutional history of the premium concept across CAA, Airports Commission and DfT analysis, and argues that IAG's own fare results evidence the premium they were meant to refute.

6.3 Three findings follow from the evidence assessments above. First, neither average fare differentials nor slot transaction values can be read directly as a measure of a scarcity rent. Fare differentials also reflect airport charges, cost differences and product mix, as IAG's charge-adjusted comparisons show. Slot prices may reflect value streams beyond the passenger-paid rent alone, as IAG's slot-values paper argues, although we consider the gap between the slot price and the premium likely to be limited, since those streams are themselves largely a product of the capacity constraint. Second, the affordability results in the Oliver Wyman report follow from its assumptions on cost pass-through, demand response and exit behaviour. These assumptions are documented but not tested against evidence or sensitivity analysis, so we consider that little weight can be placed on the £4.50 to £5.50 and £30 billion figures as constraints on what is affordable. Third, several distinct effects would be expected to raise observed fares at Heathrow relative to comparators: scarcity, hub market power, premium cabin mix, higher operating costs and higher airport charges.[36] The analyses reviewed tend to attribute the observed differential principally to one of these effects, and the evidence currently available does not allow the contributions of each to be separated. That separation is, in our view, the central unresolved measurement question.

6.4 On the second question, the response of airlines to increased aeronautical charges, the framework describes two analytical extremes. While a constraint strictly binds and the rent in a sub-market exceeds the charge, the increase would be absorbed into the rent and none would reach fares; at the competitive extreme, a charge common to all carriers would pass through in full in the long run. Real sub-markets would be expected to sit between these positions, and, because any charge increase considered here would accompany the release of new capacity,

[36] These effects are not fully independent. As set out in the framework, hub market power is itself reinforced by the capacity constraint: entry that would otherwise contest the hub premium requires slots that are not available, so part of what is observed as hub market power is sustained by scarcity (Section 3, Network economics; Section 4, Capacity rationing and allocation decisions). Any decomposition would need to account for this overlap rather than treating the effects as additive.

the strictly binding case would describe the short term only: as expansion erodes rents, a growing share of any charge would be expected to reach passengers, at a pace varying by sub-market (Section 4, Implications for a capacity expansion scenario). The submissions sit at the framework's two extremes rather than testing where between them Heathrow's sub-markets lie. The Oliver Wyman model applies the strictly constrained extreme airport-wide and holds it fixed to 2045, over the very period in which the expansion it models would be relaxing the constraint, and without evidence that its conditions hold, which the report's own observation that long-haul fares have historically funded Heathrow's charge differential cuts against. HAL and Frontier's position implies substantial recoverability of the premium through charges, which presupposes both the premium's existence and a high degree of incidence on rents rather than fares. Neither side presents evidence on the degree of pass-through itself. The submissions do not resolve the underlying economic question. Economic theory supports the possibility of scarcity rents at a constrained airport such as Heathrow (Section 4, Capacity constraints are pivotal), and the persistence of the constraint and the substantial market value of Heathrow slots, both common ground between the parties, are consistent with scarcity having economic value. At the same time, scarcity rents are not directly observable and need not appear as retained accounting profit: they may vary materially across routes, times of day, passengers and airlines, and may be capitalised into assets or reflected in costs and network choices (Section 4, Profitability and scarcity rents).

6.5 Drawing the claim-level assessments together, our position on each set of material is as follows:

- **IAG's observability paper.** The margins, cost and charge-adjusted fare analyses draw on audited or granular data and collectively indicate that a large, widespread, airline-retained premium is difficult to reconcile with the observed record. However, we do not consider the evidence to establish that no scarcity rent exists. At the level of aggregation used, a rent that arises only in parts of the market may not be detectable.
- **IAG's slot-values paper.** We agree with the paper's central point in part: slot prices cannot be converted directly into an annual passenger-paid premium, and IAG's criticisms of Frontier's calibration indicate that Frontier's figure is likely towards the upper end of the plausible range. However, the streams of value the paper says can explain slot prices without any premium are themselves claims on expected earnings from operating at Heathrow, so we consider substantial slot values difficult to reconcile with the complete absence of scarcity rents. They are more naturally read as expected rents capitalised into asset values rather than appearing in operating margins. What they cannot do is quantify the rent: IAG's own financing arithmetic rests on assumptions we cannot verify from the material disclosed, and asset values reflect expected future rents as well as any current premium.
- **The Oliver Wyman report.** The evidence supports the qualitative propositions that incremental charges of sufficient scale would erode airline margins and induce some capacity reallocation, and that affordability depends materially on capex phasing. The headline figures rest on four assumptions. Zero pass-through is applied airport-wide and held fixed to 2045, over the very period in which the modelled expansion would be relaxing the constraint, and in tension with the report's own account of how Heathrow's historic charge differential has been sustained. London demand is held fixed across scenarios. The viability threshold is drawn from the historical margin distribution rather than from observed exit behaviour or a benchmarked cost of capital. And the charge-to-capex conversion cannot be verified. None of these are sensitivity-tested. We view the envelope as one possible scenario rather than an established ceiling, and, because zero pass-through

produces the smallest affordable charge, one more likely to understate than overstate what is affordable.

- **The HAL and Frontier response.** The criticisms it makes of IAG's evidence are specific, largely checkable, and consistent with limitations we identify independently. However, the new figures HAL and Frontier put forward in the same documents, including the £600 million estimate for BA's routes, rest on assumptions that cannot be checked from the material provided, and the remainder largely restates Frontier's earlier estimates rather than offering new evidence.

6.6 A consistent pattern emerges across the exchange: each party's evidence identifies limitations in the opposing party's quantification, and neither establishes its own. The principal models on each side (IAG's required-margin model, Oliver Wyman's profitability and scenario model, and Frontier's residual-fare regression) are not provided in a form that permits straightforward replication. Each side's submissions share a commissioning party and cross-cite, so we have treated each side's material as a single body of evidence rather than as mutually confirming sources. Most tellingly, the same observations (slot transaction values, the Virgin Atlantic slot-backed financing, and the charge-adjusted BA fare differentials) are advanced by both sides in support of opposite conclusions under different unobserved assumptions. This sensitivity of interpretation to assumptions, rather than any shortage of analysis, is the principal evidential problem.

6.7 On the evidence reviewed, five propositions are not materially in dispute:

- Heathrow is fully slot constrained and has been for two decades.
- Its slots carry substantial capitalised market value.
- Short-haul capacity has shifted away from Heathrow over time. This reflects short-haul operations at Heathrow being relatively less profitable than long-haul operations there, so that under the movement cap, slots have been progressively redeployed to higher-contribution long-haul use.
- Charge-adjusted fare differentials persist on BA's matched routes.
- Airline operating margins at Heathrow are not materially above industry norms, on the estimates available in the reviewed material. Within that material, HAL and Frontier do not contest these estimates or offer alternatives; their position locates any premium in fares and slot values rather than in elevated reported margins.

6.8 We note that the displacement of short-haul by long-haul capacity is itself an observable form of the rationing behaviour the framework predicts under a capacity constraint, although it does not quantify any associated price effect. What remains contested is interpretation: whether capitalised slot value and observed fare differentials reflect a passenger-paid premium recoverable through charges, or cost recovery, product mix and network value; where any premium is concentrated; and who would bear the incidence of higher expansion-related charges.

6.9 The submissions also make recommendations to the CAA and, in the case of IAG's observability paper, to the DfT: that the slot-based figure be rejected and the scarcity rent removed from the appraisal unless reconciled to realised profits, that charging assumptions be tested against realised carrier behaviour, that HAL be required to demonstrate the £49 billion envelope is consistent with airline operating economics, and that regulatory reform introduce contestability and risk-sharing into delivery. These are policy matters rather than evidence, and several presuppose the conclusions of the analysis advanced in support of them. Our assessment speaks to these recommendations only in part. The evidence reviewed does not

justify removing the scarcity rent from the appraisal, nor relying on any single point estimate of it. Testing charging assumptions against observed carrier behaviour is sensible in principle, subject to the limitations of behavioural evidence identified above. The questions of contestability, risk-sharing and the size of the capital envelope are matters of regulatory design that fall outside the scope of this report.

6.10 Consistent with the framework, stakeholder claims should continue to be evaluated against three principles (Section 4, Summary of implications): whether the claim is consistent with the economic theory of capacity constraints and scarcity rents; whether the empirical evidence is granular enough to identify where scarcity effects would be expected to arise; and whether the conclusion drawn is proportionate to the strength and limitations of the evidence. Measured against these principles, no submission reviewed here supports a categorical conclusion in either direction.

6.11 The most balanced conclusion is that the stakeholder evidence establishes material uncertainty around the scale of any scarcity rent, where in the supply chain it is captured, and how far it could be recovered through charges, rather than proving either its existence at the level estimated by Frontier or its absence as concluded by IAG. In our view, the appropriate regulatory position is to treat scarcity rents as uncertain, and as unlikely to be spread evenly across routes, times of day and passenger segments. No single estimate in the material reviewed is reliable on its own. Any figure used should be built from several independent sources, tested for sensitivity to its main assumptions, and revisited as expansion changes the degree of scarcity over time.

# 7 Annex A: Market structure theoretical definitions

This annex sets out the standard textbook definitions, assumptions and welfare properties of the market structures referred to in Section 3. It is provided as background for readers who wish to follow the underlying theory; the application of each structure to air travel, and the implications for the questions in this report, are discussed in the main text.

## Perfect competition

The basic model of perfect competition relies upon four assumptions: (1) there are many buyers and sellers, each small relative to the overall market, all of whom are price takers[37]; (2) entry to and exit from the market are free; (3) the output of sellers is homogeneous, so that buyers regard the products of different sellers as identical; and (4) there is perfect information. The homogeneity and perfect information assumptions are integral to the model, since it is these that make every seller a price taker. Because each seller is small and takes the market price as given, no seller behaves strategically. In other words, none makes its decisions about what to produce, or whether to enter or leave the market, based on how rivals might react.

In this setting, the existence of economic profit would attract entry, increasing supply and putting downward pressure on prices until profit is driven to zero, such that price equals long-run average cost. The perfectly competitive model results in an efficient allocation of resources that maximises total surplus. In welfare terms, perfect competition is the benchmark against which the other structures are judged: because price equals marginal cost, every transaction that is worth more to a consumer than it costs to supply takes place, so welfare is maximised and there is no deadweight loss.[38]

In this market structure, firms have limited strategic levers. They cannot influence price and instead focus on cost minimisation and operational efficiency. Decisions relate primarily to input choice, production technology and scale, rather than to pricing or product differentiation.

## Monopoly

The monopoly market structure sits at the opposite extreme of the spectrum from perfect competition and introduces the concept of a price maker: a firm that can influence the market price through its decision on the quantity it produces or offers to the market. This ability is called market power and is a deviation from the assumptions of perfect competition. The basic monopoly model characterises a market with a single seller, on the following assumptions: (1) the seller is a price maker; (2) the seller does not behave strategically; (3) market entry is blocked; and (4) buyers are price takers. The key differences from perfect competition are that

[37] A price taker treats the market price as given. Because its output is small relative to the market, its own decisions have no perceptible effect on price, so it decides how much to sell at the prevailing price rather than setting a price of its own.

[38] Deadweight loss is the reduction in total welfare (the combined value to buyers and sellers) that arises when output is held below the level at which every transaction worth more to a buyer than it costs to supply would take place. It is the value of the welfare forgone, not simply the number of transactions that do not occur.

there is only one seller, with no close substitutes for its product, and significant barriers to entry.

Where entry is blocked and market power is sustained, the monopolist is able to price above average and marginal cost and to earn positive economic profit in both the short and the long run. Because the monopolist restricts output and sets price above marginal cost, output is held below the level that would prevail under perfect competition, and some consumers who value the product above the cost of supplying it do not purchase. The resulting reduction in total welfare from this lower level of output is the deadweight loss, so total welfare is below the social optimum achieved under perfect competition.

### Monopolistic competition

Monopolistic competition is a market structure that sits between perfect competition and monopoly. It combines the supplier price setting of monopoly with the large number of suppliers and free entry of perfect competition. Buyers may or may not be well informed about the offerings of competing suppliers. Price making arises in these circumstances from product differentiation between firms: while firms' products are sufficiently differentiated that each firm is a price maker, they are also sufficiently close substitutes that the actions of one firm affect the prices and profits of others. Where buyers have perfect information on the alternatives, they may still be willing to pay more for a differentiated product; where buyers are imperfectly informed, that lack of information may itself be the source of price making.

In the short run, over which the number of firms is fixed because entry and exit cannot yet occur, firms in monopolistic competition face the same optimisation as in the basic monopoly scenario. In the long run, entry is assumed to be free and new firms enter if price exceeds long-run average total cost. Entry shifts the demand and marginal revenue curves of incumbent firms inwards, so each existing firm can sell less at any given price as the available demand is shared among more firms. Firms continue to enter as long as there is an output level at which a firm's own demand curve lies above its long-run average cost curve. Long-run equilibrium is reached when no further firm would find it profitable to enter and the price charged equals long-run average total cost, so that economic profit has been eliminated.

In welfare terms, monopolistic competition sits between the two benchmarks. As under monopoly, price exceeds marginal cost, so output is below the welfare-maximising level and there is some deadweight loss. Unlike monopoly, free entry competes away economic profit in the long run, but it does so with each firm operating at less than its minimum-efficient scale, so that capacity is below the level that would minimise average cost. Set against this, consumers gain from the greater product variety that differentiation provides, so the net welfare comparison is less clear-cut than for monopoly.

### Oligopoly

Oligopoly is the structure in which firms behave strategically. The feature that distinguishes it from the other structures is the sellers' recognition of their mutual interdependence. Oligopolies can have products ranging from perfect substitutes to highly differentiated, the only requirement being that products are close enough substitutes that producers recognise their influence on one another. The structure is characterised by the following assumptions: (1) sellers are price makers; (2) sellers behave strategically; (3) entry conditions are not fixed by the model but vary from one oligopoly to another, ranging from blocked to completely free; and (4) buyers are price takers. There are few sellers, each large relative to the market.

Oligopoly has less determinate implications for economic profit, prices and quantity than the other structures. Outcomes range from the competitive level to sustained excess profit, depending on the nature of the strategic interaction (for example, whether firms compete on price or on quantity), the degree of product differentiation, the conditions governing entry and the historical context. The standard reference cases are the Cournot model, in which firms choose quantities and the resulting price lies above marginal cost but below the monopoly level, and the Bertrand model, in which firms choose prices and, with homogeneous products, competition can drive price to marginal cost even with only two firms. Strategic conduct such as entry deterrence, discussed in the main text, is a feature of this structure rather than of the others.