

The ins-and-outs of airport pricing

Airport charges are often viewed in a regulated or semi-regulated framework, where there is a presumption that they should be both non-discriminatory and reflect the “cost” of the service provided. Unfortunately, the economics of airports suggests that cost-based pricing may be neither the most profitable nor the most socially desirable method. Moreover, one-size-fits-all pricing, based on a narrow view of what constitutes non-discrimination, may result in inefficient use of existing infrastructure and potentially wasteful development of unnecessary facilities, raising overall costs.

Most airports function in a regulated or semi-regulated environment with restrictions on the level and structure of charges they can levy for “aeronautical services” like aircraft landing and processing passengers through their terminals. In today’s markets, direct public intervention in pricing tends to occur only in the more extreme instances of market power, specifically utility networks, where conditions akin to natural monopoly are most likely to persist. This makes it unlikely that there will be competition to provide the network (as opposed to competition over the network), for instance in the case of the transmission and distribution of electricity or gas.

The idea of public oversight of airport charges is so frequently taken for granted that the EU even has a directive¹ requiring all member states to establish independent supervisory bodies for that purpose. The directive imposes a number of process requirements with regard to charges on any airport with more than 5m passengers per year (or even fewer if it happens to be a country’s largest airport). These restrictions stop short of formal price regulation, but they do require the airport to consult its users over the level of charges on a regular basis and to provide an “objective justification” for them. All of this without any reference to whether the airport has any degree of market power. All that matters is size (and, by the way, an airport with 5m passengers a year is by no means large).

To anyone not immersed in the sector, such an approach to charging is immediately strange. Can you imagine Tesco being forced to consult shoppers over the cost of baked beans on its shelves, or Apple having to sound out customers before setting the price of its latest iPhone, let alone having to provide an “objective justification” for its decision? In a competitive market you choose your price and take the consequences. Indeed, in a perfectly competitive market you do not even do that: efficient prices emerge from the market itself. And note, the requirements of the Airport Charging Directive (ACD) apply based on scale, not market power. That means even airports like Stansted or Newcastle, which the UK authorities deem to require no form of ex-ante price intervention and so are free to set charges as they please, must still adhere to the directive.

While the ACD does not explicitly require airports to set their charges based on cost, the requirement for consultation and objective justification pushes them in that direction. What better, indeed other, objective justification can an airport produce for its charges than what it costs to offer the service?

At first glance this may not seem such a bad thing. As with utility networks, cost-reflective charging (at least combined with some form of incentive-based scheme or regulation) can have the effect of

¹ Directive 2009/12/EC

promoting efficiency and passing the resulting benefits on to customers in the medium term. If this works for water electricity distribution or fixed telephony networks, why not for airports?

But while the drivers of airport pricing are very different from how a supermarket sets its prices for tinned goods or how a tech company prices its latest gadget, unfortunately the pricing issues they face are also quite unlike those of access to monopoly utility networks.

Like utility networks, airports are characterised by high capital costs that are fixed in the short to medium term, meaning that average costs are usually higher than marginal costs. Utility pricing is typically based either on some variant of average cost pricing (to ensure cost recovery), as in water, electricity and gas, or on long-run marginal cost pricing to promote efficient deployment of technology, as in telephone networks.

But in the case of airports neither strategy is likely to lead to an efficient outcome. Furthermore, because of the particular markets that airports serve, correctly applied price differentiation may be both desirable from the point of view of economic efficiency and absolve the operators of charges of discriminatory conduct.

In practice, we witness examples of airport pricing that diverge from cost-based pricing for sound commercial and economic reasons. But the “ideology” of cost-reflective pricing also leads to some airports acting inefficiently as well as to unhelpful or downright harmful policy prescriptions from some regulatory authorities.

Airports do not face a one-size-fits-all problem

While it is not absolutely true to say that utility networks act as wholesale suppliers to only one retail market, the number they serve is very limited. Certainly, when it comes to residential customers, it is reasonable to view the network as serving only one market. Consequently, one-size-fits-all network, or network access, pricing tends to be the rule, especially in the domestic market.

By contrast, airports act as wholesale suppliers to airlines operating in a wide range of different markets, which drastically expands the scope of both efficient and permissible price differentiation. This is simply because every route flown from a given airport is potentially a distinct economic market, when viewed from the demand side: if a passenger is travelling from Frankfurt to Malaga, a flight from Frankfurt to Chicago is not a substitute. The direction of travel matters too. Arriving and departing passengers are not the same. Frankfurt to Chicago is not a substitute for Chicago to Frankfurt for someone starting in Chicago..

Transfer passengers constitute a different category, with additional wrinkles. Someone travelling from London to Windhoek via Schiphol Amsterdam is not in the same market as the passenger simply flying between London and Amsterdam, even though they may well arrive in Amsterdam on the same plane. The complications start to become clearer: not just one airport serving multiple markets, but even the aircraft coming and going are carrying passengers who make up different markets.

Then, of course, there are the differences in airline business model. In short-haul point-to-point markets the distinction is usually made between full-service and low-cost carriers (LCCs). Full-service airlines, in Europe, are often formerly state-owned flag carriers like BA or Lufthansa with a mix of business class and economy class passengers on a single flight; they often also act as feeders of traffic to their hub airport for connections to further-flung destinations. LCCs were the insurgents 20 years ago but are very much the driving force of the sector today. Although pressure from LCCs is gradually eroding the differences between the two groups of airlines, their demands on the airport remain quite different. To maximise utilisation of their aircraft, LCCs are focused on fast turnarounds using the front and rear doors. This favours contact stands close to the terminal (i.e. no buses) but without air bridges. Fast turnarounds may be less important for full-service carriers because their planes often have to fit in with the schedule of connecting long-haul flights. By contrast, they tend to place a little more weight on customer comfort, providing lounges for business class and frequent flyers as well as air bridges so passengers don't have to queue on the aircraft steps in the rain while waiting to board.

Then there is the distinction between short-haul and long-haul services. Long-haul planes are typically bigger, requiring larger stands and sometimes multiple air bridges (three for the A380 double-decker) and longer turnaround times. LCCs have largely shied away from long haul, although carriers including

Norwegian are starting to develop that segment as well. For these aircraft, not just the terminal but the runway itself needs to be different. The larger birds need more tarmac to lumber into the air.

Airlines, we know, become masters of price differentiation. But how should airports, subject to either direct regulation or vague regulatory oversight, react commercially in this environment?

Why airport price differentiation is not just OK, but positively desirable

Despite the rich diversity they face, it is neither practical nor desirable for airports to build separate infrastructure for each set of customers. So, for the most part, they build common infrastructure that can serve all these markets. That is the practical, cost-effective response. It does not mean that airlines need to be charged the same amount to use these facilities if their needs are different.

The key economic principle at stake here is how to apportion the costs of common infrastructure in ways that promote efficient use and maximise consumer welfare. Cost allocation processes often follow simple rules of thumb: share of floor space used or share of available runway time. But if the costs are there already and do not change with use, allocations based on these straightforward principles are just as arbitrary as a one-size-fits-all charge.

If the airport wants to make maximum use of its infrastructure, it should vary its charges to reflect different demand conditions in separate markets - even when its costs are fixed and common across different customers in different markets.

In principle, even if the airport does have market power, appropriate differentiation can benefit customers as well as the airport. Some good examples are driven by commercial incentives, but others seem to reflect an excessive obsession with “cost reflectiveness” to justify and quantify price differentiation, even when there are sound economic reasons for the practice.

Airports adopt several strategies to vary their charges by carrier. The most common approach remains essentially cost-based and entails the disaggregation of charges.

In the old days airports traditionally levied a single uniform Departing Passenger Charge (DPC), payable by the airline, which was perceived to meet all the costs of providing “in terminal” services. The charge would cover all the infrastructure of the terminal, provision of check-in desks for the airlines or their ground agents, the baggage system (for departures and arrivals, despite the charge being levied only on departures) and security screening of both passengers and bags.

In addition, the airport would charge the airline a “landing charge” for each plane’s use of the runway², airport air traffic control, taxiways and parking at the stand.

Breaking down these charges allows airports much more discretion in how they price different types of service without leaving themselves open to claims of unfair discrimination.

For instance, the DPC can be split into separate charges for passengers (mainly basic infrastructure and security screening), a fee per bag, rental of the check-in desk, self-service check-in kiosks, etc. And rather than billing carriers a fixed sum per check-in desk, the airport can charge based on how long the desk is used, thus encouraging them to be more efficient. Note: all these charges are still levied on the airline. Typically, airports do not charge passengers directly³.

The landing charge can also be disaggregated. Larger aircraft can justifiably be charged more than smaller ones because they need longer runways. And at congested airports, bigger planes may fill up

² The landing charge would also cover taking off. As with the DPC, the charge is made only in one direction, but the strong obvious correlation between arrival and departure means it is efficient to charge for only one of these movements. This historical reason why airports typically charge for passengers leaving and aeroplanes arriving is beyond the scope of this short piece. But suffice to say there is no truth in the old aviation joke that there are many more landings than take-offs: because take-offs are often cancelled, landings hardly ever.

³³ This is another area of complexity beyond the scope of this note. Because airports provide a platform allowing passengers and airlines to interact, it is not silly to think of airports as a “two-sided market”. This is a fashionable context that is sometimes mistakenly used to describe the fact that airports act as both an airport and a shopping centre. The true two-sidedness is in the airport’s connecting role, which could allow it in principle to charge passengers directly as well as airlines. Two-sided market theories establish conditions under which a platform strikes a balance in charges between the two sides. In this case it is clear that airports choose to collect all their aeronautical charges from airlines, even if they relate to costs imposed by passengers. Assuming reasonably competitive airline markets, there is every reason to think these charges are ultimately passed through to fares.

more runway time. That is because they need extra take-off/landing distances and the turbulence their larger engines creates requires greater separation between aircraft.

In addition to runway charges, separate time-related charges for parking at a gate will both encourage the efficient use of scarce infrastructure and reward those airlines that turn their aeroplanes around quickly. Airports can also break down the charge for air bridges, so an airline that does not require one does not pay for it, regardless of whether the bridge is physically present or not.

It should be noted that this sort of disaggregation has as much to do with the fairness of charges as it does with efficiency, at least in the short run. Of course, in the present the number of available check-in desks or air bridges is fixed, regardless of whether an airline uses them. So, it may seem that no efficiency is achieved by this form of disaggregation. Indeed, having a low-cost A320 parked at a contact stand by an unused air bridge may seem wasteful. But there is logic to this approach. By breaking down charges in this way, provided the individual fees are some reasonable reflection of the costs, we can expect airlines to make more thoughtful use of their infrastructure. This would bring two long-term benefits. First, greater operational efficiency means there should be less waste of capacity; hence, expansion of terminals, piers, etc. can be deferred up to a point. Second, when the time comes to expand, the airport is much better placed to judge where the needs for extra infrastructure capacity are the greatest. By contrast, one-size-fits-all pricing tends to encourage an indiscriminate approach to infrastructure, which is likely to be more wasteful in the long run.

Disaggregation is not all about separating costs

Airports practise other forms of price discrimination that are perfectly rational in commercial terms but are distorted to an extent by the regulatory fixation with cost reflectiveness. Three that immediately spring to mind are: justifying LCC discounts by building additional LCC terminals (when they still have spare terminal capacity); discounts for transfer passengers; and new-route incentives. Let me focus in this short piece on new-route incentives.

Most airports offer a discount for airline operators to open new routes. These discounts are often a significant proportion of the DPC and may last for up to three years after the route is opened. They have started to attract scrutiny, including from the Thessaloniki forum of airport regulators.

The concern here is that the cost of serving a given passenger may depend on which of the airport's various services they use, but it cannot depend on whether the airline has been flying the customer's route for only six months rather than five years. These charges can't be cost-reflective, so they must constitute a form of unfair discrimination. Let's leave aside the fact that the discounted passenger is flying on a different route to the full-fare customer and is therefore in a separate market, so no discrimination between carriers seems to be occurring. Regulators have nevertheless become concerned that discounting fares on new routes must mean that prices have to be elevated on existing routes - at least above the level that would obtain without the discounts - so the airport can recover its overall costs. This is viewed as unfair somehow. There have been moves to ask airports to demonstrate that new-route discounts reflect costs or that the recipient of the discount pays it back in the future so that other carriers are insulated.

Brief thought about this issue makes it clear that the latter requirement is impossible to achieve and, moreover, that the whole cost orientation premise is misguided.

First, from a practical point of view the operator of a new route cannot be asked to "pay back" a discount it previously received. To do so it must be required to pay more than the standard tariff after the discount has expired. One could imagine the complication: the tariff is discounted for, say, years 1-3 of a new route, then has a mark-up for the next few years (to reimburse the discount) and then eventually reverts to the "normal" level. It is not easy to see how such an arrangement could be made to work. But in any event, how does it constitute a discount? As the expected reimbursement will have to cover the expected cost of the discount, the result is no actual discount overall.

So, if new-route discounts are not cost-related and there is no option but for established routes to "fund" those discounts, are they warranted from a commercial or welfare point of view? The answer must be an unequivocal "yes", although this does not automatically justify the size of a particular discount.

The reason why discounts *are* justified stems from the way uncertainty and risk are factored into airline route economics. In offering a route, a carrier commits significant capacity and a lot of operating and marketing capital to its choice. It generally makes this commitment well before the route is launched, exposing itself to material financial risk. The airline does so after judging the amount of traffic it can generate and the fares people will be willing to pay. It can estimate these things, but it will not know if it is right until it starts operating.

The uncertainty, and hence risk, involved in opening an untried route is obviously much greater than in expanding capacity on an established route. In simple terms, the uncertainty boils down to a higher price elasticity of demand for the new service. So, how can an airport take this into account in its pricing to achieve a greater good? The answer is some form of demand-led pricing (economists often call this Ramsey pricing) that explicitly seeks to recover less of the airport's fixed costs from new routes (a discount) and charges a somewhat higher passenger fee for established routes. This is standard commercial pricing practice in many unregulated markets. It can be shown to lead to higher demand overall than in the case of one-size-fits-all pricing.

If an airport is looking to recover fixed costs from a finite number of passengers, new-route discounts, even though not cost-reflective, result in more traffic, more routes and lower average charges, because the fixed costs are spread across more traffic. This makes good commercial sense for the airport and benefits passengers too.

However, the problem with demand-led pricing is that neither concrete cost benchmarks nor accounting cost allocations are useful to regulators asking whether the size of any discount is excessive. The first test must remain this: does the airport have significant market power over its passengers, or a key segment of its passengers, and therefore the ability to abuse its position? If not, then regulators have no business interfering, despite the Airport Charging Directive. If the airport does in fact have such power, there is a risk it might discount too heavily, knowing it is relatively free to recover that discount by inflating charges to existing users. In this situation regulators could attempt to quantify the difference in risk between new and existing routes to establish the "efficient" level of discount. But to do that could be fraught with difficulty. The safest approach is likely to be to benchmark new-route incentives against those prevailing in more competitive markets in order to establish the reasonable range of commercial judgments made by airports and airlines.