

Horndal at Heathrow?  
Incremental Innovation through Procedural Change at a Congested Airport

B. S. Tether and J. S. Metcalfe

Communication to  
European Meeting on Applied Evolutionary Economics

7 - 9 June 1999, Grenoble, France

Organised by the Institute for Energy Politics and Economics  
*Organisé par l'Institut d'Economie et de Politique de l'Energie /*  
IEPE, BP 47, 38040 Grenoble Cedex 9, France

And the INRA-Unit of Sociology and Economics of Research and Development  
*Et l'unité Sociologie et Economie de la Recherche Développement de l'INRA*  
INRA/SERD, BP 47, 38040 Grenoble Cedex 9, France

## **Horndal at Heathrow?**

### **Incremental Innovation through Procedural Change at a Congested Airport**

B. S. Tether and J. S. Metcalfe

ESRC Centre for Research on Innovation and Competition (CRIC),  
The University of Manchester and UMIST,  
Tom Lupton Suite, Precinct Centre, Oxford Road, Manchester, M13 9QH.  
Tel. +44 (0)161 273 7376, e-mail: [bruce.tether@man.ac.uk](mailto:bruce.tether@man.ac.uk)

#### **Abstract**

In this paper, we examine the development of runway capacity at Heathrow over recent years, and show that capacity has increased substantially. Investigating this increase, we emphasise the role played by procedural change, rather than the introduction of new hardware technologies. Beyond this, we show that this process of innovation is based on a close and evolving understanding of how the system works, and how the key actors within the system work together and compromise to provide increased activity.

This paper develops ideas first presented in our paper 'Productivity Growth and Procedural Knowledge - Analysing the Development of London's Airports from a "Systems of Innovation" Perspective', presented at the CRIC, Workshop on 'Systems and Services', on 17<sup>th</sup> – 18<sup>th</sup> March, 1998, in Manchester. We are grateful to the workshop participants for their comments on the work. We are also grateful to the BAA's research department at Gatwick for their help in providing data and discussing 'innovations' for this project. The usual disclaimer applies.

*To many of those involved in getting the maximum out of our airports and runways, the methods used in calculating a runway's capacity remain largely a 'black art'*

Keith Williams, NATS Director of Operations, December 1993 <sup>CAP 627</sup>

## 1. Introduction

Heathrow airport is busy. Indeed, it is Europe's busiest and most congested airport (CAA, 1995). Last year (1998) over 60 million passengers used Heathrow, on 441,000 flights<sup>1</sup>, but what is remarkable is not that Heathrow is busy, but that it keeps getting busier, despite the lack of obvious investments, especially new runways and terminal buildings. In this paper, we explore the development of Heathrow's capacity to handle an ever-increasing number of flights, despite retaining the same basic (runway) infrastructure. Our primary interest is therefore in the development of runway capacity at Heathrow, and how this has been increased through procedural and organisational changes - so called 'soft innovations' - rather than through investments in infrastructure or hardware innovations. This said, our purpose is not to deny the role of hardware innovations, such as new technical equipment or new pieces of infrastructure, such as taxiways. Instead, we wish to highlight the important complementary role of 'soft innovations' which are complementary to 'hardware' innovations but which are based on procedural knowledge and the co-operation between agents..

We begin the paper with a brief historical account of the development of air traffic activity at Heathrow, before introducing the notion of capacity and showing how this has changed over time. Later, we discuss how the development of the airport's capacity depends on co-operation between different groups of actors at the airport. Because the development of the airport's capacity is based on institutional practices and requires co-operation and co-ordination, we describe the process of innovation behind the development of capacity at Heathrow as systemic.

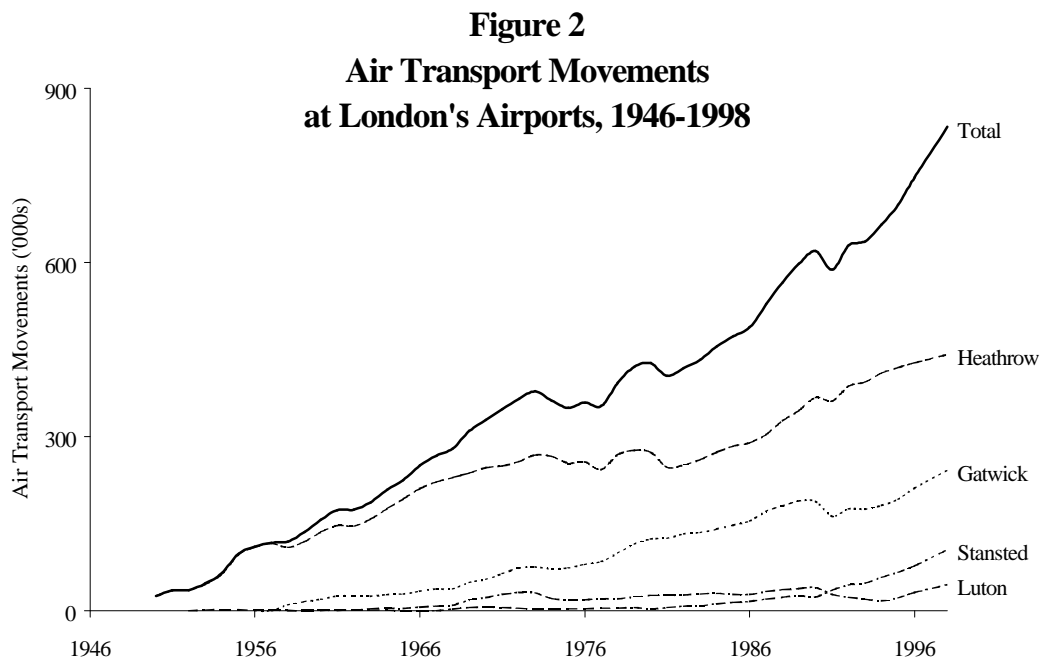
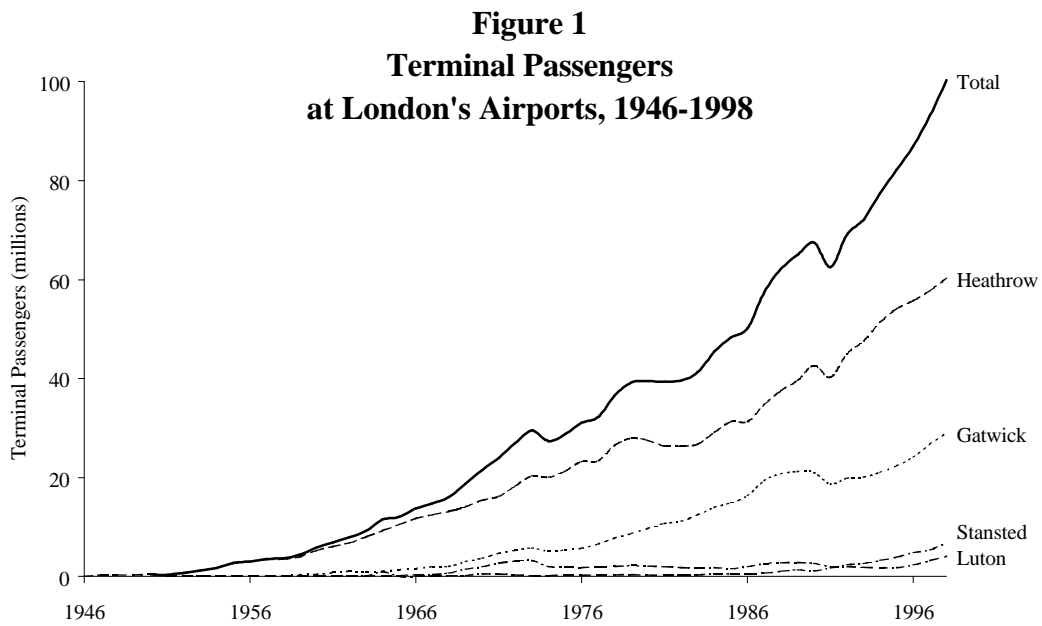
## 2. Heathrow airport – A Brief Historical Review

Heathrow airport opened in 1946, and in that year handled 63,000 passengers and 2,046 commercial flights (or 'air transport movements' (ATMs), in the industry jargon).<sup>2</sup> By 1998 the airport was handling more than two and a half times that number of passengers in a single 'average' day, and two-hundred times as many flights in a year. This growth in traffic,

---

<sup>1</sup> Data from BAA plc, the owner of Heathrow airport. See: <http://www.baa.co.uk>

together with that of London's three other major commercial airports,<sup>3</sup> is shown in Figures 1 (for passengers) and 2 (for ATMs).



The average annual rate of growth in the number of passengers using Heathrow and London's other airports has been high, as Table 1 shows, but the rate of growth has been slowing since

<sup>2</sup> For a short account of the early history of London's airports see Brooks (1957)

<sup>3</sup> The London City Airport is the London's fifth airport, but traffic levels are small relative to the other four.

the 1960s. Because of the use of larger aircraft and higher average ‘seat <sup>4</sup> growth in the number of ATMs has not kept pace with the growth in the number of passengers, yet the number of ATMs per year also continues to increase, at roughly 2% per annum in recent years.

**Table 1**  
**Traffic Growth at Heathrow and London’s Airports, 1960 - 1998**

	Heathrow Airport			London’s Airports		
	Pax. (a)	ATMs (b)	a - b	Pax. (a)	ATMs (b)	a - b
1960 - 1970	11.3%	6.2%	5.1%	14.0%	7.7%	6.3%
1970 - 1980	5.9%	1.0%*	4.9%	6.3%	2.6%	3.7%
1980 - 1990	4.5%	3.3%*	1.2%	5.5%	3.8%	1.7%
1990 - 1995	4.9%	2.8%*	2.1%	4.0%	2.4%	1.6%
1995 - 1998	3.7%	1.8%*	1.9%	6.9%	6.0%	0.9%

Figures are for the mean annual growth over the periods stated. Pax. – terminal passengers; ATMs – air transport movements. \* ATMs for Heathrow after 1970 include only passenger flights.

For our purposes the growth in traffic at Heathrow is not particularly interesting in and of itself, but is interesting because Heathrow’s capacity to handle traffic is constrained by the (supply side) infrastructure of the airport. This leads us into a discussion of the definition of capacity at an airport, but before turning to that we briefly consider the demand side.

### 3. The Demand Side

#### 3.1 – A Case of Exogenous Growth

In this analysis, we will treat demand, and the growth of demand, as exogenous. For Heathrow, we are justified in so doing for a number of reasons. Firstly, econometric studies have shown that the development of demand for air transport services in the UK has historically been closely related to the growth in GDP, and that, from a statistical perspective at least, growth in GDP largely ‘explains’ growth in demand for air transport services.<sup>5</sup> Given that the London airports have dominated air traffic activity in the UK, and that Heathrow is by far the largest airport in the London system, it follows that growth in demand for air transport activities at Heathrow is closely related to growth in GDP. Figure 3 shows a range of forecasts for the growth of air passenger activity at the London airports until 2015. These forecasts predict the number of air passengers will roughly double in the twenty years between 1995 and 2015.

<sup>4</sup> Seat loading is the jargon for passengers on a flight as a proportion of the total number of seats available.

<sup>5</sup> Although air fares, exchange rates and trade variables have also been found to be significant.

Secondly, Heathrow is the airport of preference for most airlines serving, and wishing to serve, London. Consequently, most traffic, given the choice, would prefer to use Heathrow than Gatwick, Stansted or Luton airports. As Heathrow's capacity has not developed - and has not been allowed to develop - in line with demand, the airport has become saturated, and many airlines wishing to serve London have been unable to gain adequate access to Heathrow so they have had to use Gatwick, Stansted, Luton or the City Airport instead, or not to serve London at all.

**Figure 3**  
**Forecasts to 2015 of Demand in Terms of Terminal Passengers at London's Airports**



Evidence of Heathrow's saturation is found in the demand for slots exceeding the available supply, and this situation has existed for many years, particularly at peak times.<sup>6</sup> For example, in 1994: 'The initial demand from airlines over the summer season as a whole was for some 292,000 slots, about 10% more than the total number of available slots in the daytime hours. In the busiest hours demand exceeded the available slots by more than 30%.' (CAA, 1995, para 10).<sup>7</sup>


<sup>6</sup> Some slots do remain unused in certain off-peak periods – particularly late at night. These are unused largely because of airlines are not able to gain adequate access to the airport to maintain a service, or because they are not commercially viable.

<sup>7</sup> Note that because demand is difficult to quantify, there is no definitive proof that Heathrow is saturated, but this remains the belief of the authorities. 'We believe that, unconstrained, demand to use Heathrow would significantly exceed capacity ... although the extent is difficult to quantify' (NATS et al., 1994). 'The potential demand for slots and particular uses cannot be measured precisely given the complex of influences involved and the sometimes unpredictable nature of individual management decisions' (CAA, 1995 – CAP 644, para 43).

Because of Heathrow's saturation, particularly at busy times of the day, much of the traffic which uses the other London airports can be considered overflow traffic from Heathrow.<sup>8</sup> Figures 1 and 2 show that a significant proportion of London's air traffic has been using these other airports, and particularly Gatwick, since the mid-1960s. On these bases, we treat total demand as growing, but exogenous, and in the case of Heathrow significantly in excess of the available supply. We consider the growing demand is the primary stimulus for innovation on the supply side of the airport.

### 3.2 Demand – Some Further Notes on Traffic Mix and the Size of Aircraft Used

We have just argued that demand for air traffic activities at Heathrow significantly exceeds supply, yet one interesting aspect of demand at Heathrow is that the average number of passengers per flight tends to be quite low, relative to that which is theoretically possible. Figure 4 shows that even in August, the busiest month, the average aircraft carried about 150 passengers, yet there are aircraft capable of carrying several hundred passengers. So why is the average so low, and why are larger aircraft not more widely used?



There are a number of reasons for the traffic mix at Heathrow. The first reason is that access to the airport is determined by the possession of airport slots. A pair of slots essentially provides the airline which holds them with the right to land and take-off using the airport's runways, and to use the ground facilities of the airport whilst the aircraft is on the ground.

<sup>8</sup> London's airports have a certain hierarchy, related to their size, such that airlines normally prefer to serve Heathrow, and failing that Gatwick, and failing that Stansted or Luton.

Most slots are retained by the airlines over time by a system of ‘grandfather rights’. This institution means that the airline which held a given slot in one year has the right to retain it for the next year. Moreover, before a legal judgement earlier this year, slots could not be bought or sold openly (although they could be swapped). These institutions mean that many slots are used by airlines which, in a global sense, are not making the most efficient possible use of them.

Secondly, airlines tend to use smaller aircraft on short haul feeder services that supply passengers to their long haul services for which they use large aircraft. In some cases, these feeder routes are too thin, in terms of market demand, to justify the use of larger aircraft, but smaller aircraft are also used when the frequency of a service is high. It is well known in the airline industry that frequency of service is one of the most important aspects of the service considered by passengers, particularly high paying business class passengers. Consequently, when airlines compete on a route they typically try to maintain as many flights a day (or week) as their rivals, even if this means using smaller aircraft (CAA, 1995). For example, in 1997 British Midland introduced a rival service to that of British Airways between Manchester and London Heathrow, matching British Airways in both the number of flights per day and the timing of those flights. This resulted in a loss of passengers to British Airways, and both airlines now use smaller aircraft than British Airways employed when it had a monopoly on the route. Recently, the competition authorities in Europe have been encouraging the introduction of a second or even a third airline on routes which were previously monopolies or duopolies, and this has encouraged the use of smaller aircraft.

It is also widely known in the airline industry that, as demand for a route grows over time, airlines will, if they are able, first increase the number of flights then the size of the aircraft serving the route, rather than using larger aircraft then adding more flights. This behaviour reflects passengers’ preference for more frequent rather than lowest cost services, especially on short haul routes.

Consequently, whilst the total demand for air traffic activities at Heathrow can be considered an exogenous stimulus to the creation of more capacity on the supply side, the actual structure of demand is not wholly unrelated to the supply side capacity constraints. Figure 4 shows the steady growth in passengers per ATM during the 1970s, the period during which large wide-bodied aircraft and especially the Boeing 747 were introduced. This growth in

passengers per ATM slowed during the 1980s, but has since increased again in the 1990s as Heathrow has become increasingly constrained in terms of its available capacity.<sup>9</sup>

#### 4. Introduction to the Supply Side

##### 4.1 Supply Side Strategies in Relation to Capacity Creation

Faced with growing demand, the airport authorities have, theoretically, a number of choices as to how to react in terms of the provision of facilities on the supply side. Firstly, they can do nothing, essentially resisting the increased demand and let the airport stagnate. Whilst this is possible, the value of each additional flight at Heathrow to both the airline which operates it and to the airport is such that the airport has every incentive to increase the capacity of the airport and to run its existing facilities as efficiently as possible, within the constraints posed by the regulatory authorities and by the practices of the airlines.<sup>10</sup> For this reason, we disregard the ‘do nothing’ option in this analysis.

The second, and perhaps most obvious response, is to build more infrastructure – a new runway, terminal building or parking area – as required. This approach has been followed in France, particularly at Charles de Gaulle airport to the north of Paris, because the airport is located in a sparsely populated area and has a large area within which to expand, and the French approach to planning transport infrastructure has been much more strategic than that in the UK.<sup>11</sup> Heathrow, unlike Charles de Gaulle airport in Paris, is located in a highly urbanised area of west London, and is locationally constrained by the small area of the airport. Heathrow occupies a site of 1,200 hectares, which is considerably smaller than that of several other major European airports, including Charles de Gaulle airport which occupies a site of 3,200 hectares.

<sup>9</sup> In fact, the principal constraints at Heathrow are currently terminal and stand capacity. Stand capacity is where aircraft park whilst at the airport. It is these constraints that the proposed Terminal 5 at Heathrow is intended to overcome.

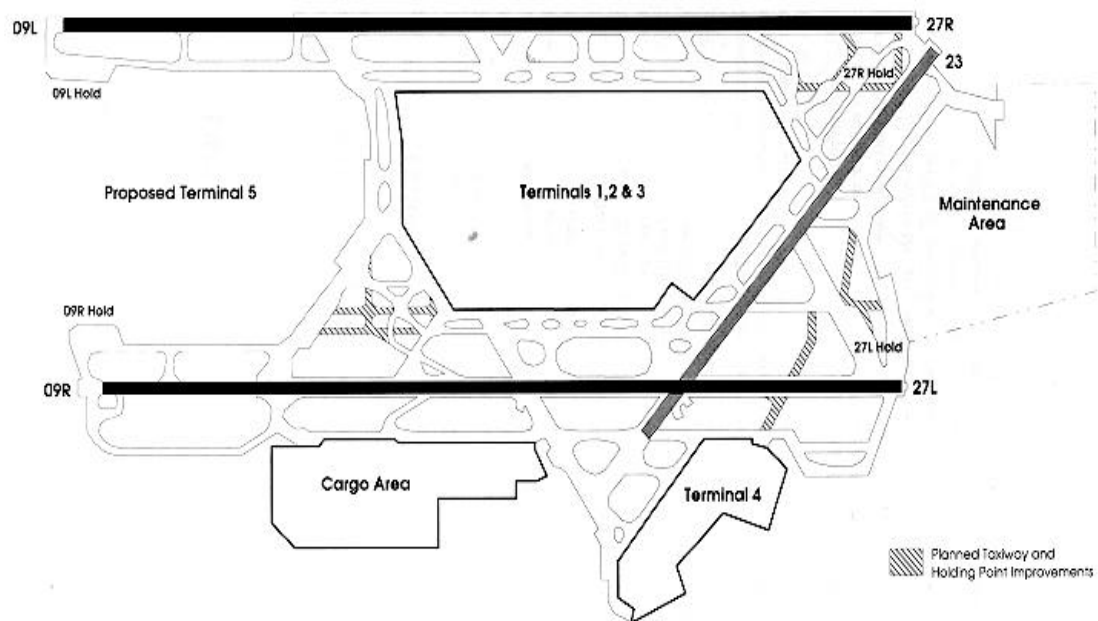
<sup>10</sup> The monetary value of a pair of runway slots is not fully known as a market for slots has only recently become legal, but the value is thought to run into several million pounds, particularly for slots during busy parts of the day. In 1987 a high level study (CAA 1988 – CAP 534) noted: ‘Each additional five slots, i.e., one in each peak hour, was thus associated with an annual average earnings potential of £180 million per annum at Heathrow and £70 millions at Gatwick.’ ‘Total benefits calculated this way were £2.2 million per annum at Gatwick for five additional peak slots, i.e., one in each peak hour (equivalent to £5 million when capitalised). At Heathrow, the benefit would be 50% more because of increased interline revenue.’

<sup>11</sup> French planning law has been favourable to the development of large infrastructural projects considered to be in the national interest. British planning law is very different, as the case of the proposed Terminal 5 at Heathrow makes clear. In fact, air transport in the UK has had repeated battles over planning permission over the years. There were major inquiries into the construction of a third London airport, initially at Maplin, then later at Stansted, and there have been inquiries into the development of major facilities, such as Terminal 4 at Heathrow, which, incidentally, the airport operator initially opposed. (for the history and debate about UK planning in relation to airports and especially London’s third airport – see Higham, 1995)

Because of its relatively small site, the reluctance of the British planning system to allow the air transport industry to develop strategically and unhindered, and the lack of determination on the part of British governments of all political hues to support the industry, Heathrow's basic infrastructure has remained largely unchanged over the years. Figure 5 shows a plan of the airport, with the location of Terminals 1, 2 and 3 between the main runways, Terminal 4 and the cargo area to the south, the maintenance area to the east, and the location of the proposed Terminal 5 to the west. Most importantly from our perspective, the airport has two main runways (and a rarely used cross-runway) – this is the same basic runway infrastructure that it has had since the 1960s. Indeed, previously the airport had six runways, but three of these were built over during the expansion of the central terminal area (Graves, 1998). Consequently, the emphasis at Heathrow has been on making the most of existing facilities, adding incrementally to these facilities, and 'innovating' largely on the basis of procedural changes which improve the efficiency with which the existing facilities are used. The aim of this paper is to highlight and discuss the important role of this sort of 'soft innovation'.

**Figure 5 – Plan of Heathrow Airport**

**Heathrow Airport Layout**



#### **4.2 The Supply Side – Defining Airport and Runway Capacity**

The first point to note in discussing airport capacity is that airports do not have one, but multiple constraints that impinge upon their capacity. These constraints relate to, amongst

other things, the capacity of the runways to handle aircraft, the capacity of the stands to accommodate parked aircraft, the capacity of the terminals to handle passengers and the capacity of the transport infrastructure which links the airport to the outside world. Any of these may be, at a given time, the binding constraint on the capacity of the airport.

Our interest is in runway capacity, and the constraint this imposes on the level of traffic at the airport. We note, however, that runway capacity has not always been the primary, or the binding constraint at Heathrow. Indeed, presently the BAA, the airport's owner, argue that it is terminal and stand capacity are the primary constraints at Heathrow; Terminal 5, is their proposed solution to these problems. Nevertheless, runway capacity is certainly a central element in determining the overall capacity of an airport, particularly in terms of air transport movements.

The maximum number of aircraft that a runway can handle, for a given set of conditions, such as the weather and traffic mix, is known as the runway's 'service rate'. The 'service rate' does not equate with capacity, because it takes no account of delays. Instead, runway capacity is 'the number of aircraft movements that may be scheduled to use a runway such that their average delay during the airport's busy period does not exceed a specified value' (CAA, 1993). Until 1990, the agreed delay criteria was an average of 5 minutes, but in 1990 this was increased to 10 minutes.

Before assessing further how runway capacity is determined, we consider the 'environmental constraints' which have a bearing on runway capacity. These environmental constraints are, to a greater or lesser extent, exogenous to the determination of runway capacity, but are none the less important in influencing the way in which the runways are operated (in a sub-optimal way).

### **4.3 Runway Capacity – Environmental Constraints**

Although a private company, Heathrow is not permitted to operate in the most efficient possible way, because of legal restrictions and by agreement with the airlines and the communities around the airport.

#### 4.3.1 Safety Related Constraints

Aircraft are subject to separation rules, which ensure aircraft are kept a safe distance apart. For aircraft landing and taking-off, separations are enforced for three main reasons:<sup>12</sup>

- Firstly to ensure that the leading aircraft has successfully completed a manoeuvre before the following aircraft is committed to undertake that manoeuvre.
- Secondly, aircraft are separated for wake vortex reasons. Large aircraft especially have powerful engines that create turbulence in the aircraft's wake. Generally the larger the aircraft the greater the turbulence it can endure, so when a small aircraft follows a large aircraft it has to follow at a greater distance than when a large (or 'Heavy') aircraft follows another large (or 'Heavy') aircraft. Aircraft are categorised by wake vortex separation criteria for which separation minima are specified on departure and on arrival (CAA, 1993).
- Thirdly, aircraft operate at different speeds, with turbo-prop aircraft being 'slow' compared with 'fast' jets, and it is necessary to enforce longer separations to allow slow aircraft to manoeuvre in an environment in which fast aircraft dominate.

For our purposes, the wake vortex separations are the most significant separations, as we shall show shortly the cumulative significance of these can be minimised by clever procedures. However, it is notable that wake vortex separations 'are based more on experience than scientific research' (CAA, 1993). They are, however, now under-review, and a more scientific approach is being taken, which implies a shift from knowledge based on experience to knowledge based on scientific principles in the determination of these separations.<sup>13</sup>

---

<sup>12</sup> We recognise that this is a simplification of the reality.

<sup>13</sup> Also of interest is how the categories were established. The International Civil Aviation Organisation (ICAO) divides aircraft into three categories based on maximum take off weight in kilograms. These are Light (7,000 or less), Medium (7,001 to 135,999) and Heavy (136,000 or more). In the UK, however, the Medium Category is divided into Small (7,001 to 40,000) and Medium (40,001 to 135,999). In an industry in which insurance plays a key role, and in which insurance often depends on compliance with international (ICAO) standards, it would be interesting to know how the UK came to introduced the distinction between Small and Medium traffic in the early 1980s. Moreover, around the same time, aircraft such as the Boeing 707 and Douglas DC-8, although in the Heavy category for their take-off weights, were reclassified as Medium category aircraft. Beyond this, there is some suggestion that if the Medium category, which dominates traffic at Heathrow but which includes aircraft of quite widely differing sizes, were divided up, some savings on runway time might be made. It would be interesting to know whether any attempts have been made to seek a further reclassification, and if so whether these attempts have met with resistance.

### 4.3.2 Noise Related Constraints

As Figure 5 shows, Heathrow has two main runways, and these can be operated in different modes. In *segregated mode*, one of the runways is used for landing aircraft, whilst the other is used for taking-off. In *mixed mode* both runways are used for aircraft landing and taking-off, with departing aircraft interleaved between arrivals. It is known that more aircraft are able to use the runways within a given time period in mixed mode than in segregated mode, and that mixed mode is therefore more efficient. A study in 1993 estimated that Heathrow could handle 92 aircraft movements (arrivals and departures) in an hour if mixed mode were used, compared with 82 movements per hour in segregated mode (NATS et al., 1994). ‘The major disadvantage of segregated mode is that – because of aircraft vortex wake - streams of arriving and departing aircraft have to be comparatively widely spaced and therefore potential runway capacity is wasted. Moreover, its lack of flexibility means that capacity may be unused on one runway whilst excess arrivals or departures on the other runway suffer delays.’ (NATS et al., 1994, para 2.2)

Heathrow operates in segregated mode because this imposes less noise pollution on the surrounding communities.<sup>14</sup> Also for noise reasons, aircraft departing Heathrow must follow a restricted number of paths when leaving the airport, called Noise Preferential Routes (or NPRs) under 4,000 feet. Because of the need to separate traffic, NPRs also restrict the possible number of departures in a given period, but again clever procedures can make the best of this situation, Later in this paper, we will discuss how the airport is able to squeeze out extra capacity within these restrictions.

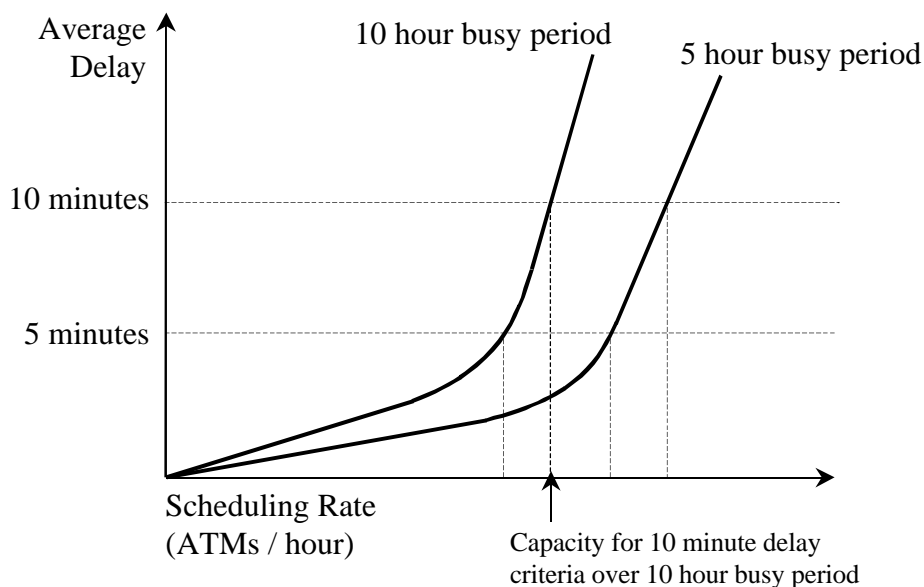
### 4.3.3 Constraints related to User Requirements

Apart from constraints imposed by safety considerations, and by the need to placate the neighbouring communities with respect to the noise nuisance, airport operations are also compromised by the need to satisfy the users. On the airside, that means the airlines.

---

<sup>14</sup> Further details of the operational procedures undertaken to reduce the noise nuisance are (NATS et al., 1994):  
 1. An ‘alteration procedure’ applies to the use of the runways during westerly operations. This means that one runway is used for landing between 7 am and 3 pm, after which the other runway is used until 11pm. This is done to equalise the noise distribution and give periods of respite to the communities under the respective approach flight paths. During westerly operations there are no restrictions to using both runways for departures. This means that some departures can be interleaved between arrivals on the ‘landing runway’.  
 2. With easterly operations, the close proximity of the village of Cranford to the end of the runway severely restricts the use of 09L for departures. The ‘Cranford Agreement’ means that aircraft can only take-off from 09R, but both runways can be used for landing.  
 3. ‘Preferred Runway Selection’ is the procedure whereby, even in light easterly winds, westerly operations are sustained in preference to easterly operations, provided that tailwind and crosswind limits are not exceeded. This is to reduce the number of times that aircraft on ‘take-off power’

**Figure 6**  
**The Relationship between Runway Capacity, Delays and Length of Busy Period**



For airlines, delays are a significant problem, not just because they annoy passengers, but also because they waste fuel and make idle valuable assets. However, if no delays are allowed at a busy airport, this would severely reduce the operating efficiency of the airport, for there is a well understood relationship between the average delay and the number of aircraft movements the runways are able to handle (see Figure 6):

As the volume of traffic scheduled to use the runway increases, delays remain low until a critical point is reached. As a guide, this critical point occurs when aircraft scheduling reaches about 80% of the runway's 'service rate'. At this stage, delays begin to increase rapidly for relatively small increases in movements. This is because the number of aircraft in the arrival or departure queue is building gradually to the point where long queues develop, and with more aircraft waiting for longer and longer periods, the average delay soars. The point at which this dramatic rise in delays starts – and the rate at which delays then continue to increase – depends on the length of the busy period.<sup>15</sup> The longer the busy period, the earlier and steeper is the increase in average delays and the greater the 'peak delay' (NATS, 1993)

---

overfly the built up areas to the east of Heathrow. 4. Departures from Heathrow must follow prescribed noise preferential routes (NPRs) below 4,000 feet.

<sup>15</sup> The 'busy period' is defined below.

The effect of increasing activity on average delays was noted by a study in 1988: ‘[A]n increase at Heathrow by one movement over the declared capacity for arrivals or departures would produce a rise in the average delay from 5 to 7 - 8 minutes’ (CAA, 1988, para 4.1)

A compromise has therefore to be reached between the scheduling rate – the number of movements the runways can handle in an hour – and the average delay the airlines will tolerate. Between the mid-1970s and 1990, the agreement between the airlines, the airport and air traffic control authorities, was that the average delay criteria should be 5 minutes, but in 1990 this standard became 10 minutes.

#### **4.4 The Supply Side – Defining Runway Capacity**

After consideration of the geometry of the airport, and the safety and environmental constraints within which air traffic can operate, runway capacity is primarily determined by the length of the delays regarded as acceptable by the aircraft operators. Thus:

The capacity of a runway is defined as that level of average demand, over a busy period of defined length which on being serviced gives rise to average delays of a specified amount. The capacity, consequently, depends not only upon the ‘hardware’ of the runway, the queue discipline and the air traffic control procedures and separation standards, but also on the length of the busy period and on the level of delay regarded as tolerable by the users of the system, that is, predominantly, commercial airline operators. (CAA, 1983)

We will discuss the importance of ‘queue discipline’ and ‘air traffic control procedures’ later in this paper, as it is innovation in these that have given rise to a considerable amount of increased capacity through ‘soft innovation’, but also of note from the above quote is the significance of capacity being defined over a ‘busy period’. The busy period is the period over which aircraft are arriving and/or departing in a continuous stream, so there is always at least one aircraft demanding to use the runway(s).

A notable feature of this definition is that it shows that runway capacity is largely based on experience, and experience which only arises when the airport reaches a level of activity at which demand exceeds the ability of the airport to handle that demand. It was only in the 1970s that the level of air traffic activity at Heathrow and Gatwick became sufficiently high that capacity studies began. Thus capacity is not fully known in the abstract, but is instead ‘discovered’ through a learning process over time, which is closely related to the practical experience of operating the airport. We discuss this further below (in Section 4.4.1). The

techniques for assessing runway capacity have also developed over time, for example with the introduction of increasingly sophisticated computer simulation.

In the UK, the determination of runway capacity is a distributed process, in which air traffic control, the airport and the airlines all play a role.

Assessment has a number of stages. First there is the monitoring of actual operations to record exactly what is happening, and the analysis of that data. The level of traffic, delays and time intervals for aircraft operations are monitored by observations at Heathrow and Gatwick by DORA personnel.<sup>16</sup> The next step is the use of fast-time computer simulation – using that data – to calculate the levels of delays that would occur for a given schedule. The result is a count of the movements that used the runway in each hour together with delay estimates which translate into ‘capacity profiles’.

Following an assessment of the operational implications by NATS ATC [Air Traffic Control] managers, capacity ‘profiles’ are presented and agreed with the airport’s scheduling committees [which are composed of airline representatives] who then use them to plan their schedules. In considering the various ‘profiles’ available, the scheduling committee and airport operator have the opportunity to consult with NATS and airlines and take decisions on the patterns of their schedules, on the traffic mix and on the level of acceptable delay in order to meet the optimum capacity profile which will meet the needs of all parties. (CAA, 1993)

Our interest in runway capacity is twofold, although these interests are strongly inter-linked:

1. Incremental Increases in Capacity, despite fears of Saturation
2. The Distributed Nature of Capacity Creation

#### 4.4.1 Incremental Increases in Capacity, in the Face of Fears of Saturation

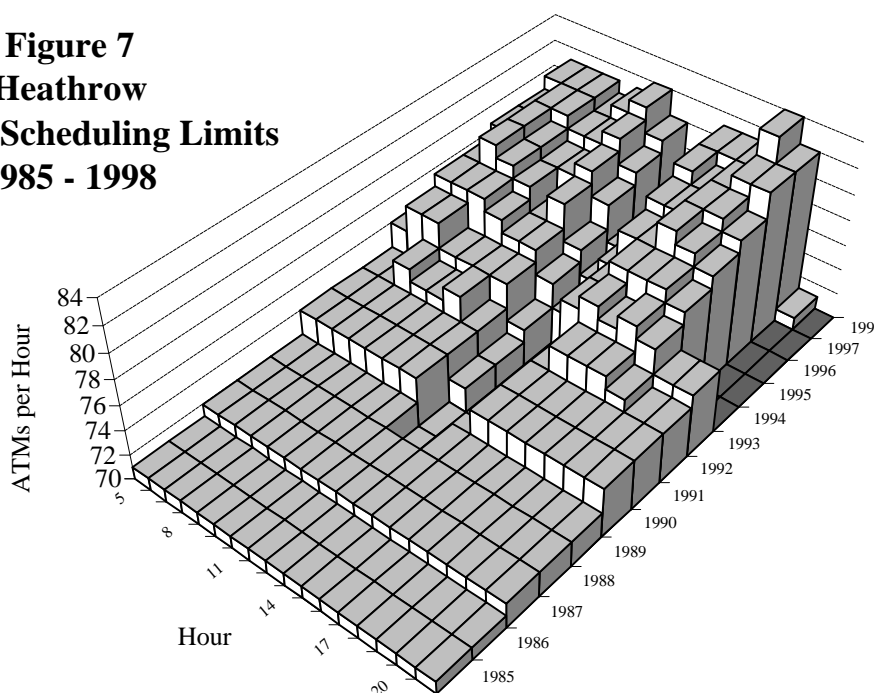
Firstly, we observe that the capacity of the runways at Heathrow (and Gatwick) has increased steadily over time (although throughout the period of our study there has been surplus demand – particularly at Heathrow). Figure 7 and Table 2 show this development in the airport’s capacity over time. The figure shows that in 1981 the airport was able to handle 67 air transport movements in each hour, whilst by 1998 this figure, whilst now varying across

---

<sup>16</sup> DORA = Department of Operational Research and Analysis, in the Chief Scientist’s Division, of National Air Traffic Services (NATS), of the Civil Aviation Authority.

the day, reached 84 movements per hour. The table shows that across the fourteen hours of the day between 0600 and 1959 GMT, the number of available slots increased by 20% over the 18 year period. Although there has been some deterioration in service – particularly with the acceptance of a longer delay criteria – this growth in capacity is quite remarkable, given the context of continually expressed fears that the capacity of London’s airports, and Heathrow in particular, was becoming severely limited.

**Figure 7**  
**Heathrow**  
**Summer Scheduling Limits**  
**1985 - 1998**



**Table 2**

**Development of Capacity at Heathrow – 14 hour day Summer Slots Available, 1981 - 1998**

Year	1981	1982	1983	1984	1985	1986	1987	1988	1989
Slots: 14 hour total	938	966	980	994	994	994	1,008	1,008	1,000
Index 1981 = 100	100	103	104	106	106	106	107	107	107
Year	1990	1991	1992	1993	1994	1995	1996	1997	1998
Slots: 14 hour total	1,020	1,024	1,037	1,061	1,084	1,099	1,110	1,119	1,127
Index, 1981 = 100	109	109	111	113	116	117	118	119	120

This concern is reflected in a reports by the CAA and a Parliamentary Inquiry into Airport Capacity in the South East of England. In 1986, for example, a high level committee was established to investigate runway utilisation at Heathrow and Gatwick. This committee was chaired by the Chairman of the CAA, and included airline industry ‘heavy-weights’, such as Micheal Bishop, chairman of British Midland [and (possibly) Lord King, chairman of British Airways - to be checked].

The Committee saw its task as being to demonstrate conclusively (or to put in hand studies which would ultimately demonstrate conclusively) whether or not it was possible to increase the capacities of Heathrow and Gatwick beyond the levels currently foreseen. (CAA, 1988, para 3.1)

After a year of research, the Commission reported various means by which capacity could be increased, but by only one arrival and one departure movement per hour, and that only in part of the day.<sup>17</sup> Beyond this, and although the Committee recognised that the capacity of Heathrow had increased over the years, the report states firmly that there was little room for further improvement, and that consequently the government should investigate, as a matter of urgency, the issue of providing further air traffic capacity in the south east of England:

Declared hourly capacities have shown a gradual increase at both Heathrow and Gatwick over the last decade or so. This has been the result of changes to ATC procedures, including changes in the vortex wake classification scheme and a general tightening of achieved separations, and, to a lesser extent, improvements in the ground layout at both airports which BAA has made. (para 4.4)

After a detailed and critical examination of the technical issues, the Committee has determined that, within its terms of reference, the scope for increasing runway capacities of Heathrow and Gatwick beyond that now envisaged is necessarily small (para 8.1).

The pressure of demand on runway capacity at Heathrow and Gatwick combined with current forecasts of passenger demand in the South East (of England) strongly suggest that unless Stansted airport is allowed to operate to its full potential as a single runway airport, capacity in the London area will be insufficient to meet likely traffic growth beyond the early to mid 1990s. Even with the full utilisation of Stansted's single runway there is a strong possibility that runway capacity will be insufficient by the turn of the century (para 8.5).

After this, a number of further studies were undertaken, including a major study on 'Runway Capacity to Serve the South East (RUCATSE) (DoT, 1993), a runway capacity enhancement study at Heathrow (NATS et al., 1994), and a Parliamentary Select Committee Inquiry into UK Airport Capacity (HMSO, 1996). Two important points emerge from these studies:

---

<sup>17</sup> This report began the process of declaring a variable capacity – taking into account the traffic mix and how that differed throughout the day – e.g., in the afternoon the proportion of heavy traffic is less, allowing an

1. Firstly, the capacity of the runways at Heathrow and Gatwick are not known in an abstract theoretical sense, instead they are understood in relation to the practical experience of operating the runways. An example helps make the point. In response to fears that runway capacity in the London area would soon become inadequate, the airport's owner (the BAA), together with the air traffic control service (NATS) and the civil aviation authority (CAA) were asked to assess the potential capacity of a second 'close parallel' runway at Gatwick. A 'close parallel' is a runway which is too close to the first runway to be operated with complete independence from the first runway, and it is 'known' that the capacity of a close parallel is less than a fully independent runway. However, what was not known was what that capacity would be. The BAA, NATs and the CAA conducted trials over a three year period to determine the capacity of a 'close parallel' runway at Gatwick, showing that the capacity of a runway is not simply known by experts in advance.
2. The second important point arising from these studies is that whilst it is generally stressed that scope for further capacity formation at Heathrow and Gatwick (with their existing infrastructures and operating restrictions) is extremely limited, and the airports are close to saturation, the studies have also normally found some ways through which a small amount of increased capacity can be provided through incremental change, which depend rather more on changes to operating procedures than to the adoption of new technologies. For example, the Heathrow Airport Runway Capacity Enhancement Study (NATS et al., 1994) found:

Our examination of segregated mode shows that measures could be introduced to increase both arrival and departure capacity from 39 to 41 movements per hour (82 movements per hour in total), subject to changes to NPRs [Noise Preferential Routes] being approved for new SIDs [Standard Instrument Departure routes] (para 2.10).

We concluded that there is unlikely to be any further strategic improvement in peak hour runway capacity at Heathrow in segregated mode beyond 82 movements per hour (para 2.22).

This combination of incremental improvements to capacity using existing infrastructures and operating procedures with an unknown ultimate capacity makes planning difficult – or

---

increase of 0.5 movements per hour during this period. Also the provision of extra holds – 0.3 additional

at least makes it easier for politicians to defer the difficult decision as to when and where new runway capacity should be built. Of course because new capacity is not provided through the construction of new runways, further incremental change using the existing facilities is encouraged.

#### 4.4.2 The Distributed Process of Capacity Creation

Secondly, we are interested in the process of capacity creation, and in particular the distributed nature of this process, which is illustrated by the fact that although the air traffic control (ATC) authority bears the ultimate responsibility for deciding the capacity levels that can safely be accommodated (because it bears a legal responsibility for safely operating the procedures affecting runway capacity, 'the declaration process involves considerable iterative discussion between airport ATC, the airport operator and the airport scheduling committee.' (CAA, 1993) Thus the operation of airports, particularly on the air-side, depends on the interaction and co-operation of the airport operator, the airlines and the safety authorities, particularly air traffic control. Over time, institutions develop by which these different agents understand their role and what is expected of them. Sometimes these expectations are incorporated into legal requirements, but others are norms which are not enshrined in law.

The case of slots is an important example, for airport slots are not really owned but they are controlled by the airlines which have the right to keep their slots from one year to the next under the principle of 'grandfather rights'. A slots is an important institution, which provides the airline which holds it with precedence in the use of the runways within a limited time period around the slot-time. Aircraft which miss their slots times lose this privilege, and must wait for another slot to become available. The existence of slots, and associated rights, in some ways limits the scope for efficiency savings for the airport operators and air traffic control.

We shall show later that if aircraft are processed in strict 'first come first served' order the efficiency of the runways will be less than if the airport operator / air traffic control have some scope for re-ordering the traffic into a more efficient sequence. Principles, such as 'first come first served', can be hard to break, and individual airlines must be persuaded that they will not lose out by such changes, possible through the introduction of new rules that limit the extent of the re-ordering (such as limiting the number of times an individual aircraft can be overtaken in the queue to land or depart).

---

movements per hour, reduced separations from 3 nm to 2.5 nm in good conditions.

We emphasise that while incremental innovation through procedural change often depends on the creation of new rules or institutions, new rules are difficult to negotiate because of the vested (and different) interests of the parties involved. Air transport is a global industry, and airlines can be reluctant to change their ‘standard operating practices’ (or SOPs) to suit the needs of individual airports<sup>18</sup> – particularly as there is no direct link between good performance and rewards in terms of lower airport charges or the award of more slots as these are made available. The performance of the airlines is observed and compared, and discussions are held with those that are less efficient in their use of the runways. Interestingly, however, there is no attempt to single out individual pilots, and appeals are made through the Chief Pilots on the basis of professionalism and good practice, rather than on the basis of penalties and rewards, because of the rule that everyone must be treated equally.<sup>19</sup>

The existence of SOPs is often reinforced by industry wide standards which are agreed at the International Civil Aviation Organisation (ICAO) and other international bodies, but compliance with ICAO standard operating procedures is important for insurance purposes. Airlines are often reluctant to deviate from ICAO standards because their insurance and licensing can depend on compliance with these standards.

Heathrow, and a limited number of other airports around the world, have sought to improve their efficiency which has meant trying to move away from some of the ICAO standard operating procedures, and this has necessitated the negotiation of a new set of standards, called ‘High Intensity Runway Operations’ (HIRO). These provide a number of special operating procedures. For example, an aircraft on final approach may be given clearance to land whilst the preceding aircraft is still on the runway, so long as the first aircraft is at least 2,500 meters from the threshold when the landing aircraft crosses the threshold (Duke, 1996). This contrasts with standard operating procedures whereby the first aircraft must have cleared the runway before the second is cleared to land. As most airports do not have the same need for efficiency related measures as Heathrow, there is a reluctance on the part of many airport operators and many unaffected airlines to renegotiate the existing standards.

---

<sup>18</sup> For example US aircraft don’t power up before turning onto the runway because of what might be behind – i.e., a light aircraft. This, however, slows the manoeuvre and extends runway occupancy, reducing the efficiency of the runways.

<sup>19</sup> For example, it was discovered that one US airline took longer in using the runways on departure than the average. Discussion revealed that this was because the airlines standard practice was to use a satellite downlink to set the flight computers just prior to departure. This led to a delay in the airlines use of the runways.

These changes in practice must also be negotiated because they may have cost implications, and an example helps make this point. Runways operate most efficiently when the occupancy time of each aircraft is minimised, and the gap between aircraft operating the runway is minimised. Clearly, one way to minimise runway occupancy is to request that aircraft brake severely on landing, so that they slow down rapidly and can leave the runway quickly. Apart from not wanting to terrify the passengers, such a request may be refused by the airlines because of the costs it imposes on them – through increased break-wear, the costs of new aircraft breaks being considerable.<sup>20</sup>

The key point of this section is that changes in operating procedures upon which efficiency savings largely depend are themselves the result of negotiations between the different agents involved. Many lines of change are resisted, because vested interests do not make them acceptable to one or more of the interested parties. Thus not only is it important to know what might potentially be done to improve the efficiency of runway operations, but it is also important to know whether a change is likely to be acceptable, and if not whether there are ways of circumventing the objection. In the following section, we will discuss various ‘innovations’ which have gradually been introduced in the last twenty or so years which are largely responsible for improving the efficiency with which Heathrow’s runways are operated, but in general these ‘innovations’ do not depend on new ideas, instead they depend on changing procedures and institutions which tend to be very sticky to change. For example, rapid exit taxiways (or RETS) were ‘invented’ through basic research at the University of Californian in the late 1950 (Hornonjeff et al., 1958), but were only adopted slowly, and little was done about measuring the effectiveness and optimal positioning of RETs until the late 1970s and 1980s (Warskow, 1970).

Clearly, there is a relationship between the highly institutionalised way in which airports and runways in particular are operated, and the lack of definite knowledge about an individual runway’s ultimate capacity – for capacity depends on these institutions and the extent to which they may be changed.

## **5. ‘Innovations’ related to Runway Utilisation**

In this section, we will finally turn to some specific ‘innovations’ which have been introduced and which have contributed to the incremental increases in runway capacity at

---

<sup>20</sup> Moreover, modern aircraft tend to be highly automated, and land on ‘auto-break’ which is designed to ensure comfort and to protect the breaks from excessive wear.

Heathrow over the last twenty or so years, but we begin by showing evidence of improvements in runway occupancy times and inter-aircraft separation times which are the keys to improving runway capacity.

### ***5.1 Reducing Separations and Runway Occupancy – the Keys to Efficiency Savings***

The key to maximising capacity – in terms of air traffic movements within a given time (usually an hour) - is maximising throughput of aircraft in a given period of time, and the keys to this are twofold: minimising separations between aircraft using the runways, and minimising the time taken by each aircraft on the runways.

#### ***5.1.1 Reducing Separations between Aircraft***

The time spacings between aircraft when arrivals follow each other directly is partially determined by the wake vortex criteria, such that when a Heavy aircraft leads a Light aircraft the light aircraft must stay further back than when a Heavy follows a Heavy, for example, or when a Heavy follows a Light. Table 3 provides separation matrices, derived through visual observations, for each combination of leading and following aircraft. These matrices are important for sequencing, an ‘innovation’ which we shall discuss shortly, but our immediate interest is in how the separations have reduced, albeit slightly, over the years. For example, between 1980 and 1988/89 a saving of between 7 and 8% was found in the inter-arrival time between two Medium aircraft, but significant savings of over 10% were common. The only increases were with Light aircraft following Mediums in 1988 and Light aircraft following other Light aircraft in 1989, but these may be ‘rogue’ results, because these particular means are based on very few observations.

At present we do not have more recent data for separations in the 1990s, and will be seeking this in our future research, but for now the key point we wish to stress is the apparent improvement in spacings between aircraft over time. Whilst it is true that the improvements are relatively small, this is a business in which small changes can yield significant benefits. Moreover, these small improvements are consistent with our story of incremental innovation, based largely on procedural change. As the CAA points out, the accuracy of spacings is very important to the efficient use of the runways, and although air traffic control bears ultimate responsibility for maintaining separations,<sup>21</sup> the achieved spacings also depend on the co-

---

<sup>21</sup> In the US, under certain conditions responsibility for separations can pass to the pilots, and this normally results in smaller average separations. Such a change would require a change in the law in the UK. It is also thought that because Heathrow has so many different airlines, many of which have low frequency flights into

operation and skills of the pilots, and on technology, as well as on the skills of the air traffic controllers:

The accuracy with which the required spacings can be achieved in practice also affects capacity substantially. Ideally, the actual spacings achieved between aircraft using the runway should match the required safe separations as closely as possible. This depends on a combination of controller skill, pilot skill and the use of ground-based and aircraft technology to achieve the required spacings' (CAA 1993 – our emphasis).

**Table 3**

**Mean Separation Times between Westerly Arrivals at Heathrow – 1980, 1988 and 1989**

		FOLLOWING			
L		Heavy	Medium	Light	
E	Heavy	104	142	178	- Times are in Seconds
A	Medium	99	94	128	
D	Light	91	92	68	

Separation Matrix - 1988		FOLLOWING			Improvement cf. 1980				
L		Heavy	Medium	Light	L				
E	Heavy	100.7	132.1	None	E	Heavy	-3%	-7%	n.a.
A	Medium	85.8	87.5	146.6	A	Medium	-13%	-7%	+15%
D	Light	82.4	81.4	None	D	Light	-9%	-12%	n.a.

Separation Matrix - 1989		FOLLOWING			Improvement cf. 1980				
L		Heavy	Medium	Light	L				
E	Heavy	98.8	122.7	170.0	E	Heavy	-5%	-14%	-4%
A	Medium	83.8	86.9	122.5	A	Medium	-15%	-8%	-4%
D	Light	74.8	81.4	71.0	D	Light	-18%	-12%	+4%

In relation to this, there is evidence to suggest that controllers learn to separate aircraft more efficiently over time, although this is an unconscious learning-by-doing process, made more difficult by the fact that at the time controllers set up the second arrival behind the first they do not know exactly how long the first arrival will occupy the runway as this is affected by different aircraft operating characteristics and pilot performance. (CAA 1993).<sup>22</sup>

the airport, the lack of familiarity amongst the pilots with the airport might increase the average separations rather than decrease them if responsibility for separations were given to the pilots.

<sup>22</sup> Another example of learning-by-doing amongst air traffic controllers is found in their improved handling of the LINK helicopter which linked Heathrow and Gatwick between 1980 and 1982, after which it was withdrawn

### 5.1.2 Reducing Runway Occupancy Times

The separations just discussed also depend on runway occupancy times – if runway occupancy times can be reduced, separation times can be reduced, and again, small improvements matter: ‘saving an average of 5 seconds on every aircraft’s runway occupancy would add another 1 - 1 ½ movements per hour.’ (CAA 1993).

Table 4 below shows the runway occupancy times for runway 27R in 1980, 1984 and sometime around 1994.<sup>23</sup> This Table appears to provide evidence for a gradual reduction in runway occupancy times, with aircraft making greater use of the early exits rather than progressing down the runway (note the decreased use of the exits in Block 9 and 11, and the increased use exits in Block 13). Moreover, occupancy times for aircraft exiting in most of the blocks appear to be declining.

**Table 4**  
**Runway Occupancy on Heathrow 27R – developments between 1980 and c.1994**

	9	11	13	14	15	16
Proportion per Exit, circa 1994	17%	22%	37%	18%	5%	1%
Proportion per Exit, 1984	22%	37%	14%	18%	6%	3%
Proportion per Exit, 1980	23%	34%	26%	14%	3%	n.a.
Seconds on Runway, circa 1994	70 - 73	50 - 60	46 - 52	43 - 49	32 - 41	33 - 39
Seconds on Runway (mean), 1984	77	58	49	48	41	36
Seconds on Runway (mean), 1980	78	61	50	50	45	n.a.
By Aircraft Type – circa 1994 data:						
B747, Concorde	85%*	15%				
A300/A310, B767, MD11, L101, IL86	36%*	38%*	26%			
B737, MD80, A320, B727, TU54	5%*	26%*	47%*	19%	3%	
B757		9%*	43%*	40%	8%	
BA46, FK28, DH7			55%*	32%	9%	4%
FK27, FK50, ATP				28%*	36%	28%

\* - scope for improvement in runway occupancy

Table drawn from Duke, 1996, except 1980 data (CAA Paper 81014) and 1984 data (CAA Paper 85010)

We should note that direct comparisons are difficult, due to changes in the traffic mix – i.e., the types of aircraft using the runways – larger aircraft tend to spend longer on the runways than smaller aircraft – as the lower part of the Table shows. Although we are still gathering

---

due to the completion of the M25 motorway. Following observations in 1981, it was reported that ‘The effect of the LINK helicopter is to reduce runway capacity (at Gatwick) by 0.7 movements (per hour), an improvement of 30% on the 1979 observations. This seems to be a consequence of the greater experience of controllers in handling the helicopter in a manner which minimises its effect on the other traffic.’ (CAA, 1982 - Paper 82019).

information on this, we would expect that changes in the traffic mix would, *ceteris paribus*, increase runway occupancy times, as more larger aircraft are using the runways. Against this, modern aircraft have better brakes allowing them to slow down sooner.

Finally, the bottom part of the table is interesting because for each category of aircraft a large proportion are identified as having ‘scope for improvement’ in their use of the exits and consequently runway occupancy times. This illustrates the interest of the airport operator in gathering such information in order to improve the efficiency of the system, particularly because pilot skills and technique rather than technology is seen as the key to runway occupancy times (CAA, 1993). Since 1993, the airport has run a ‘runway occupancy awareness campaign’ to make pilots and air traffic controllers aware of the significance of minimising runway occupancy.

## 5.2 *Specific ‘Innovations’ in Runway Occupancy*

In this sub-section, we discuss some specific ‘innovations’ which have been introduced and which have contributed to the incremental increases in runway capacity at Heathrow over the last twenty or so years. We stress that at present we do not have a full picture of the relative importance of these innovations. We hope to develop a more complete understanding of their relative importance in the course of the research.

### 5.2.1 *Rapid Exit (and Access) Taxiways and Procedures*

Traditionally, exit taxiways are perpendicular to the runway, but it has long been recognised that one way of reducing runway occupancy times, and thereby increasing runway capacity, is to a design runway exit at taxiways at 30° to the runway which allow aircraft to exit at a much higher speed than with a conventional exit (Hornonjeff et al., 1958; Warskow, 1970). An expert at the BAA claims the optimal positioning of Rapid Exit Taxiways (RETs) is worth 2 to 3 additional movements per hour.

But the introduction and use of RETs appears to be a classic example of how procedural change is more important than the hardware component of innovation, for RETs can be built but pilots (and airlines) have to be persuaded to use them effectively. Gradually, RETs have become standardised, and indeed the ICAO publishes a book of standard designs for RETs.<sup>24</sup>

---

<sup>23</sup> The actual date was not provided in the reference.

<sup>24</sup> Note that these standards often have to be adapted to local circumstances, for at Heathrow there is very little room between the runways and the outer-taxiways.

Standardisation has encouraged their widespread use. The introduction of indicator lighting systems is also encouraging the use of the earliest possible rapid exit.

Efforts are also being made to introduce Rapid Access procedures, to minimise runway occupancy on departure. These appear to be part of the High Intensity Runway Operation rules in force at Heathrow, whereby aircraft line-up on the runway as soon as the preceding aircraft has commenced its take-off roll. Moreover, pilots are encouraged to have completed all cockpit checks prior to line-up, and to minimise the number of checks requiring completion on the runway. The aim is to ensure pilots commence the take-off roll immediately take-off clearance is issued.

### 5.2.2 Arrival Sequencing

It has long been recognised that whilst a runway can be regarded as a single server queuing system, it need not have a first-come first serve queue discipline, and by skilful choice of the next movement to be handled, the runway controller can increase the number of movements handled in a given period of time (i.e., the rate of service) (CAA, 1983).<sup>25</sup>

The controller has the opportunity to re-order, or sequence, the aircraft, rather than treat each aircraft on a random ‘first come first served basis’, and this can provide some significant benefits in terms of the efficiency with which the runways are used. The key to sequencing is the fact that separation times between different types of aircraft are not the same. (Note that if there were only one type of aircraft using the airport there would be no benefit from sequencing). The controllers aim is therefore to minimise the time between (in this case) arrivals. Sequencing is not simple, partly because it is based on rules, but also because the strategy used can vary with the traffic mix and with the separation matrix.

However, to illustrate the benefits of sequencing, we developed as simple simulation model which sequenced arriving aircraft in a simple way. This model considered just three types of aircraft (Heavy, Medium and Small), and assumed all aircraft were in the same speed category. We assumed a traffic mix of 27% Heavy, 67% Medium, and 6% Small, roughly similar to that at Heathrow in 1993. We also assumed the inter-arrival separation matrix shown below in Table 5. According to this matrix, the separations were greatest when a Small Aircraft followed a Heavy aircraft – due to wake vortex problems - (140 seconds) and were smallest when any type of aircraft followed a Small aircraft (80 seconds). The strategy

of the model was therefore simple – as far as possible, the controller (i.e., computer) should avoid Small aircraft following Heavy aircraft.

**Table 5**

***Arrivals Sequencing Simulation – a. Assumed Separation Matrix***

		Following Aircraft		
		Heavy	Medium	Small
Leading Aircraft	Heavy	100	120	150
	Medium	85	85	100
	Small	80	80	80

***b. Summary of Results of Simulation Exercise over 50 Runs***

	Arrivals in Random Order (a) (Cumulated Time)	Sequenced Arrival Order (Cumulated Time)	Runway Time Saved (b)	(b) as a percentage of (a)
Mean	9377.6	9027.3	350.3	3.7%
S.D.	47.0	20.4		
Maximum	9495	9070	425.0	4.5%
Minimum	9290	8970	320.0	3.4%

The first step was to allow the controller (i.e., computer) to ‘look ahead’ a certain number of aircraft. We chose five. Thus instead of having to land the next arrival in strict ‘first come – first served’ order, the controller is able to assess the next five flights due to land. Initially, the controller selects the first arrival amongst the largest category by size of aircraft available. This lands and is replaced in the pool of five possibilities by the next arrival in the random sequence. The controller then lands any other aircraft of the same size category in the pool of five, each time replacing the aircraft which lands with the next to arrive in the random sequence. When the supply of aircraft of the size being picked is exhausted, the controller moves to the next category of aircraft smaller than that from which it has been selecting (i.e., Medium following Heavy, or Small following Medium). If, however, it has been selecting Small aircraft, there is no smaller category, and it instead has to select the first arrival amongst the largest category of aircraft available by size.

This simple procedure reduces the number of times a Small aircraft follows a Heavy aircraft in the landing sequence, and it is remarkable that such a simple procedure can yield quite significant time savings. Using the inter-arrival separation matrix in Table 5a, and allowing

<sup>25</sup> Consequently, as demand increases queues tend to become longer, which in turn provide the controller with a greater choice as to the next movement. Because of this, the service rate is not independent of the level of demand.

the controller to look ahead just five aircraft, we assessed the total landing time of 100 aircraft (in a traffic mix 27% Heavy, 67% Medium and 6% Small) arriving in a random order and compared this with the total time taken by these aircraft if they were sequenced by the procedure described above. Table 5b provides a summary of the results after 50 simulation runs.

This shows that over the 100 arrivals, an average time saving of 3.7% was achieved by sequencing the aircraft compared with the time the aircraft would have taken had they landed in a random, ‘first come-first served’, order. Moreover, apart from this significant time saving, the other notable feature of the simulation runs was that the standard deviation between the total times taken was substantially reduced, and the difference between the largest amount of time taken and the smallest amount of time taken was reduced by half (from 205 seconds to 100 seconds). Thus, not only was time saved (thus allowing greater capacity to be handled within the time available), but the pattern of activity also became more regular than with the simple random ordering of arrivals.

We know that air traffic controllers have been using sequencing techniques for some time, and the evidence we have collected so far suggests that the extent of sequencing is increasing. The actual sequencing strategies may vary, and we do not yet know how significant the strategies used are in terms of reducing cumulated landing times, however, it is likely that they provide significant time savings.

An important element of sequencing is that it depends on the principle of ‘first come first served’ being abandoned. This means that safeguard rules must be introduced to prevent individual aircraft loosing out badly in the re-sequencing of traffic.<sup>26</sup>

### 5.2.3 Departure Sequencing

Departure sequencing is done for the same reason as arrival sequencing – time can be saved by reordering departures – but it is an even more complex problem. The controller’s objective is to re-order the aircraft at the holding point so as to select the appropriate departure sequence which minimises departure intervals – ideally 1 minute between departures wherever possible. One minute intervals are achieved when the following aircraft

---

<sup>26</sup> Indeed, this was found with our sequencing model, for potentially an aircraft had to wait some considerable time before being allowed to land. We therefore introduced a new rule, such that no aircraft could be passed in the queue by more than 10 other aircraft. This solved the problem, but reduced the efficiency of the re-sequencing, such that the mean saving in runway time was reduced from 3.7% to 2.1% - that is more than 40% of the efficiency gain. Most probably a more sophisticated approach would not lead to such a severe reduction in the time saving.

departs on a different route from the lead aircraft. The minimum time between departures using the same route is two minutes, but can be longer if wake vortex separations apply (Graves, 1998). The departure rate is increased by switching departures between those using the left, right, straight-ahead tracks.

But there is more to it than that. Apart from their wake vortex categories, aircraft are grouped according to size and speed and when an aircraft from a faster group follows a slower one extra time must be added. Thus there is a lot of mental arithmetic involved, and ‘Achieving the optimum departure sequence is by no means easy. The scale of difficulty is illustrated by the fact that at Heathrow the departure controller is expected to memorise rules for over 500 route combinations and then apply the appropriate speed allowances before deciding whether wake vortex separations should take precedence. With say ten holding aircraft, he or she must then determine which of 3½ million possible sequences matches flow control requirements to give the least delay and achieve maximum use of the runway.’ (CAA, 1993). He or she must also take account of the slot times of the aircraft waiting at the hold, as aircraft with slot times have to be fitted in within a limited time-period, because scheduling is based on these slot times (Graves, 1998).

The efficiency of departure sequencing has been further improved by the redesign of the departure routes (known as Standard Instrument Departures – SIDs), so that these do not interfere with each other.

## 6. Discussion and Conclusions

The title of this paper - ‘Horndal at Heathrow?’ - refers to the Horndal iron works in Sweden, which Lundberg (1961) discovered increased its productivity<sup>27</sup> by an average of 2% per annum despite receiving no new investment for 15 years. Lundberg thus identified ‘the Horndal effect’ – an increase in output based on learning, which itself is ‘the product of experience’ (Arrow, 1962), rather than by investments in new capital equipment.

Other examples of productivity due to learning by experience are provided by Arrow (1962) and Rosenberg (1982), and include Wright’s (1936) study which found the number of labour-hours required to produce an airframe was a decreasing function of the total number of airframes of the same type previously produced.

---

<sup>27</sup> Output per manhour.

Arrow emphasises the importance of ‘learning by doing’, which, according to Rosenberg (1982, p. 121), ‘is a form of learning that takes place at the manufacturing stage after the product has been designed. ... Learning at this stage ... consists of developing increased skills in production. This has the effect of reducing real labour costs per unit of output.’ Rosenberg goes on to identify a related form of learning by doing, that has not been highlighted in the literature, which concerns small improvements to the production process which arise out of a direct involvement in the production process. For this type of learning by doing:

The learning involved requires participation in the production process. Such participation is obviously not sufficient, because a perception of possible improvements depends not only upon the opportunity to make certain observations but also upon prior training and experience. The point is that productive activities always involve specialised kinds of knowledge, much of which may be unique to a specific industrial process. There is typically a range of possible improvements that require intimate familiarity with the minutiae of the productive sequence.

(Rosenberg, 1982, p. 122)

Rosenberg (1982) then goes on to identify another sort of learning – learning by using – which is based on an increasing familiarity amongst users with a particular product or technology. Thus ‘learning by using refers to a very different locus of learning than does learning by using’ (p. 122), and although the experiences of users with a technology can feedback to inform producers about possible improvements to the next generation of that technology, learning by using essentially takes place on the technology as given.

In this paper, we have been interested in the growth of air traffic activity at one of the world’s busiest airports – Heathrow, and have observed that air traffic activity has grown substantially despite the fact that the airport retains the same basic infrastructure today – in terms of the number of runways - as it has had for more than 30 years. We must stress that we are not saying Heathrow has not had any investment over this period – indeed, investments have been substantial. What we are claiming is that much of the improvement in the airports capacity is due to procedural changes in which new technological hardware plays a relatively minor role.

Our study is novel because it examines a service activity, whereas most studies of learning in industry have focused on manufacturers. This is not a trivial difference, because an important

feature of services is that services are consumed as they are produced. This is sometimes referred to as ‘co-terminality’. The ‘learning-by-doing’ identified by Arrow (1962) and Rosenberg (1982) is an activity which provides productivity improvements to producers, whilst ‘learning-by-using’, as identified by Rosenberg (1982), is an activity which provides productivity improvements to users. To the extent that such learning is conscious and subject to encouragement and choice,<sup>28</sup> within the context of a manufactured product these activities are essentially separate and the producer is free to make changes to the process of production of the technology, whilst the user is free to make changes to the pattern of use of the technology. This is different in services, because the ‘product’ does not exist as an entity independent of the producer, or indeed the user, and the service is consumed as it is produced. Thus the producer’s freedom of action is likely to be more constrained by the operational requirements of the user. In the case of airports, this situation is complicated further by the importance of the regulatory authorities whose primary interest is safety.

We have been seeking to open up the ‘black art’ by which runway capacity has developed over time. We have stressed the importance of an evolving understanding of how the system works, together with an iterative process by which a relatively (but not optimally) efficient compromise situation emerges between the provider (the airport operator), the user (the airlines) and the regulatory agencies (air traffic control and the civil aviation authority) so that the system may be made to work more efficiently. Improvements are not automatically or easily imposed. For example, Airlines have Standard Operating Procedures which they are reluctant to change for individual airports, not least because of the insurance cost implications. So the (conscious) learning that is significant is not just about what are the possible changes that might be made (the possibility space), but extends to understanding the needs of the other actors in the system, so the number of real options is less than that offered

---

<sup>28</sup> There is a difficulty with ‘learning-by-doing’ and ‘learning-by-using’ in the sense that it is not clear whether these types of learning are confined to sub-conscious processes (which are automatic and costless – Malerba, 1992), or whether they extend to conscious and deliberate changes, based on choice, to the organisation of the production process or the pattern of use of the technology, where the opportunities for organisational change arise out of learning by experience. With reference to the experience of the Horndal iron works, Arrow (1962) presumes that because there was no new investment there was also no ‘significant change’ in the methods of production. This raises the question as to what constitutes a significant change, but our point is that whilst some of the productivity improvements at Horndal may have been due to unconscious and automatic ‘learning-by-doing’, another component may have been due to conscious decisions by management and workforce about how to better organise the production process. If we ascribe productivity increases to ‘learning-by-doing’ in the narrow sense of unconscious automatic processes, then to what extent is really a measure of the importance of this, and to what extent is it a measure of our ignorance about the importance of deliberate organisational and procedural change.

by the possibility space. In this sense, innovation in services through procedural change has a close relation to the design function for manufactured goods.

Another interesting feature of this inter-dependence is that because procedural change at an airport requires the close collaboration of the airport operator, air traffic control (which is independent of the airport), the airlines and the safety regulator, and because the issue of capacity is so important politically, a lot of information has been documented revealing the known options and conscious choices taken in relation to capacity creation. Thus we can ascribe a considerable part of the increased capacity to choices made on the basis of this conscious understanding (although further improvements will also arise through unconscious learning by doing/using is also going on – amongst pilots and air traffic controllers). Compare that with the situation at a factory, which increased its output significantly, but neither customers nor regulators had any direct interest in this process, and the management were free impose any changes they saw fit without the need for extensive documentation. We might well ascribe the productivity improvement of this factory to an unconscious and automatic ‘learning-by-doing’ process, because we have no other explanations.

## References

- Arrow, K. J. (1962) 'The Economic Implications of Learning by Doing', *Review of Economic*
- Brooks, P. W. (1957) 'A Short History of London's Airports', *The Journal of Transport History*, Vol. 3, pp. 12-22.
- CAA (1981) 'Runway Capacity and Aircraft Delays at Gatwick Airport', CAA Paper 81002, Civil Aviation Authority, London.
- CAA (1983) 'The Methodology of Runway Capacity Assessment – A Summary Paper', CAA Paper 83019, Civil Aviation Authority, London.
- CAA (1985) 'Arrival Runway Capacity and Arrival Delays – Heathrow Airport, Summer 1984', CAA Paper 85010, Civil Aviation Authority, London.
- CAA (1988) 'Report of the Committee on Runway Utilisation at Heathrow and Gatwick', CAP 534, Civil Aviation Authority, London.
- CAA (1989) 'Heathrow Demand and Delays: Monitoring of Arrivals, Summer 1988', CAA Paper 89008, Civil Aviation Authority, London.
- CAA (1990) 'Heathrow Demand and Delays: Monitoring of Arrivals, Summer 1989', CAA Paper 90004, Civil Aviation Authority, London.
- CAA (1993) 'A Guide to Runway Capacity – For ATC, Airport and Aircraft Operators', CAP 627, Civil Aviation Authority, London.
- CAA (1995) – 'Slot Allocation: A Proposal for Europe's Airports', CAP 644, Civil Aviation Authority, London.
- DoT (1993) 'Runway Capacity to Serve the South East (RUCATSE) – Report of the Working Party', Department of Transport, London, June.
- Drew, A. E. and Logan, H. K. (1981) 'Arrival Runway Capacity and Arrival Delays at London Heathrow Airport, Summer 1980', CAA Paper 81014, Civil Aviation Authority, London.
- Drew, A. E. and Logan, H. K. (1982) 'The Effect of Certain Facilities on the Runway Capacity and Aircraft Delays at Gatwick Airport, 1981', CAA Paper 82019, Civil Aviation Authority, London.
- Duke, G. (1996) *Air Traffic Control*, 6<sup>th</sup> edition, Ian Allan, Shepperton, Surrey.
- Graves, D. (1998), *UK Air Traffic Control: A Layman's Guide*, 3<sup>rd</sup> Edition, Airlife, Shrewsbury, England.

- Higham, R. (1995) 'A Matter of the Utmost Urgency: The Search for a Third London Airport, 1918-1992', in Leary, W. M. (ed.) *From Airships to Airbus – the History of Civil and Commercial Aviation – Volume 1, Infrastructure and Environment*, Smithsonian Institution Press, Washington and London.
- HMSO (1996) *UK Airport Capacity – Second Report in the Session 1995-96 of the Transport Committee of the House of Commons*, London.
- Horonjjeff, R (and others) (1958), *Exit Taxiway Location and Design*, Institute of Transportation and Traffic Engineering, University of California.
- Malerba, F. (1992) 'Learning by Firms and Incremental Technical Change', *Economic Journal*, 102 (July), pp. 845-859.
- NATS, IATA and BAA (1994) 'Report on the Heathrow Airport Runway Capacity Enhancement Study', National Air Traffic Services, Civil Aviation Authority, Cheltenham, England.
- Rosenberg, N. (1982) *Inside the Black Box – Technology and Economics*, Cambridge University Press, Cambridge (see particularly Chapter 6 on 'Learning by Using').
- Warskow, M. A. (1970) 'Means of Increasing Airport Capacity', in *World Airports the Way Ahead – Proceedings of the Conference held at the Queen Elizabeth Hall, London, 23-25 September, 1969*, Institute of Civil Engineers, London.