Quantitative Methods in Aviation Management

Quantitative Methods in Aviation Management

By

Tony Webber

Cambridge Scholars Publishing

Quantitative Methods in Aviation Management

By Tony Webber

This book first published 2023

Cambridge Scholars Publishing

Lady Stephenson Library, Newcastle upon Tyne, NE6 2PA, UK

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library

ISBN (10): 1-5275-9152-2
ISBN (13): 978-1-5275-9152-3

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

FOREWORD

With his vast experience in the industry, Tony approached this book through a lens of practicality by brilliantly blending industry practices, data modelling techniques, and econometric theories. In this book, Tony uses many industry examples to demonstrate how econometric models can solve industry challenges and guide business decisions. With a strong sense of aviation data and a passion for applying econometric theories in aviation, Tony introduces key aviation concepts in Part A of this book. This lays a strong foundation for readers to understand how this key information contributes to econometric modelling.

In Part B of the book, Tony introduces the concept of stochasticity, following that of aviation market volatility. Aviation practitioners and students cannot avoid facing the fact that most data we deal with in aviation is uncertain. For example, passenger demand for a particular route is uncertain, as is the seasonal effects on ticket prices or competition amongst airlines. It's particularly appealing that the concept of distribution sampling is brought to the reader's attention in an aviation context. These statistical techniques and tools are essential for readers when they face imperfect data in analysis, and when trying to make an informed decision under an acceptable level of business risk.

Time series data models and regression data models are introduced in Part C/D of this book. This part provides the readers with a great opportunity to utilise learned knowledge and skills in developing econometric models in an aviation context. The best feature of this book is that each section is filled with examples illustrating how econometrics is applied in decision-making and how econometric models can be used to solve real aviation challenges.

Teaching an aviation economics course at a top university takes work. If the course is too economics-focused, then aviation students will lose interest because there is a gap between economic theories and aviation. When the course is too aviation and less technical, we often find that students need to develop economic senses to approach the aviation business from a 'business lens'. After all, aviation is a big industry, and profitability is essential for every company.

This book is certainly one of its kind that will quickly become a popular textbook on aviation econometrics in university classrooms. Beyond the classroom, this book is an essential reference title for anyone who deals with aviation economic data and is keen to apply econometric modelling techniques in the era of data-driven decision-making. Learning should be challenging, informative and fun. Tony's book achieves that and beyond!

Cheng-Lung Wu
December 20, 2022

PREFACE

In 2020 I was handed the task of teaching a subject at the University of New South Wales in Sydney Australia called Quantitative Analysis in Aviation Economics. When I inquired to a colleague about the text that the course was using the response was that there wasn't one. The students relied on notes that were prepared by previous course coordinators, which were in reasonable shape but needed some work. In response I set about writing the book that you are currently reading.

There are a wide variety of quantitative books that are written for subjects outside of mathematics, such as business and finance, economics, medical science and healthcare, criminology, psychology, law and even real estate. What these subject areas have in common is that researchers can develop a theory about how a variable behaves, but unless they can validate that theory by interrogating a set of data then they are solving only half the problem. The other quite important characteristic of these subject areas is that they are usually supported by the availability of very rich datasets, providing the potential for confident data analysis to occur.

This leads us into quantitative analysis in aviation, in which the same principles apply. Aviation managers may have excellent theories about what is happening within their businesses, but unless they can substantiate these theories with the right data analysis then they are just theories, and you can't take just theories to the executive management of an airline because they will tell you to come back when you have the data. To illustrate, an airline group which flies both a low-cost carrier brand and a full-service airline brand may believe that it is best to fly its low-cost carrier on a route because they believe that most passengers on the route are price sensitive. Unless analysts within the airline can present to executive management data that supports this view, such as data on average income levels for the route, or market survey data that is able to determine how airfares rank in terms of the most important variables that drive decision to fly, then they are presenting a data-less view which does not carry nearly as much weight as a data-supported theory.

Consider another example. Senior management within an airport believe that if they set landing charges relatively low, this will enable the airport to attract more airlines to the airport. The increase in passenger numbers that this brings to the airport will generate more revenue from the sale of non-aeronautical goods and services, such as duty-free goods, food and beverages at food outlets, more vehicles using the airport carpark and more transit passengers using the in-airport hotel. Unless analysts can use data to support their view that lowering landing charges will attract more airlines to the airport, then it is only a view or theory without any supporting evidence.

I started my quantitative journey in the main part when my honours supervisor at the time Professor Barry Hughes began calling me an econometrics elephant, which I gathered was a semi-affectionate way of saying that I was more analytics focussed rather than an elegant economics writer. Under the guiding eyes of my applied econometrics professor, Ron Bewley, back in the early 90s I improved my skillset in developing econometric models to interrogate data, and that stepped-up a notch when I began teaching Masters and PhD students in these subjects during the late 90's. In 2004 I managed to land a job as General Manager of Microeconomics and then Chief Economist of Qantas Airways, which lasted for 7 years. It was during this episode in my working life where I could test to see if my quantitative skills could make a difference in a commercial setting.

I used these skills in many different areas within the airline. One of my first quantitative tasks was to examine the extent of the correlation between the Australian dollar and oil prices. This correlation was used to demonstrate that the exchange rate could provide the airline with up to 40% protection against higher fuel prices, reducing the need to use expensive financial instruments such as swaps and options to hedge the airline's fuel exposure. Another major task involved building an econometric model to forecast the airline's revenue and present the findings of that model to the Qantas Board. A similar type of modelling exercise was used to determine the average airfare that maximises revenue on key routes. Quantitative models were also used to forecast the path of passenger demand across various routes, which in turn was used to determine future fleet requirements.

This book is an introduction to quantitative methods for those choosing to become aviation managers. It is designed mostly for undergraduate students, or postgraduate students studying statistical methods applied to aviation for the first time. There are many topics covered in this book that only touch the surface of the subject matter, and could have books entirely devoted to them, such as time series analysis and bivariate and multivariate regression. The focal point of this book is on the logic of the topic being described and the aviation applications of that topic. The intention is not to reinvent the wheel because the quantitative methods 'wheel' has well and truly been invented. The intention is to focus on my area of expertise, which is to take true and tried methods of data investigation and show how those methods can be applied to topics in aviation management. Hopefully one fine day you will use the contents of this book as the basis for your own quantitative modelling within an airline, an airport or air traffic control, amongst many other important aviation organisations.

PART A:
VARIABLES, GROWTH RATES AND DESCRIPTIVE STATISTICS

Chapter 1

Introduction

1.1 The Importance of Data for Airline Businesses

Airlines collect an enormous amount of data about their businesses, their operations and their financial performances. The engineers of the airline collect information about the performance of their aircraft to understand whether there might be an issue with their safe flying. The sustainability and fuel management analysts within an airline collect information about the quantity of jet fuel that the airline's fleet of aircraft burn in operations so that they might understand aircraft carbon emissions and fuel productivity. The managers within human resources collect information about how happy the airline's labour force is with their job and how the company is being run, to gain insight into the likelihood that it may lose key personnel to other airlines or to other companies outside of aviation. Marketing and advertising executives and personnel collect data to understand how well potential passengers are aware of their marketing campaigns, and the extent to which that awareness translates into greater sales volumes and revenue. Frequent flyer program analysts collect data on the ways that members earn and burn points to better understand the effectiveness of their loyalty program. Operations managers collect information about the reliability of the airline's scheduled services and on-time performance because they understand that this will influence future demand and revenue generated by high-yielding passengers. Managers in the sales area collect data on seats sold and the reaction of sales to discounting strategies to better understand the way that passengers react to lower and higher prices. Treasury, Finance and Strategy executives and management collect data on the revenue and costs of the airline to form a view about the current and future financial performance of the airline, which will in turn drive their budgets, forecasts and future strategies.

While the accurate collection and dissemination of aviation data is pivotable to gaining a better understanding of the outcomes and results of different airline business areas, another incredibly important aspect of the data is the ability to extract from it meaningful insights. To do this, it is necessary that airline management have the required statistical and analytical skills. Without this, and with an enormous sea of data to deal with, the costs associated with the collection of data will not see an adequate return. These statistical and analytical skills will allow an analyst to understand the data that needs to be collected from the set of data that is available. It will allow the analyst to understand what type of relationships may exist between the different variables in the dataset. It will allow the analyst to determine whether additional external data may be necessary to understand the relationships that exist in the internal data. It will also allow aviation analysts a better chance of choosing the right techniques and approach to extracting insights from the data.

There are many different techniques that airlines and airports can use to query the data that they collect. This ranges from simple measures such as estimating the business-as-usual growth rates for key variables. While this may seem a relatively straight-forward task, there are many ways to estimate business-as-usual growth rates, with most possessing significant issues that may generate misleading answers depending on the data that is used. It also includes other simple measures such as descriptive statistics, estimates of central tendency, dispersion and correlation. It may also include more sophisticated hypothesis testing, which include tests associated with a single mean, tests associated with more than one mean, tests of independence and goodness of fit, and analysis of variance. Airlines may also use regression techniques, including both bivariate and multivariate regression techniques. Regression analysis can be used to understand relationships between variables, which may be used to make strategic changes to the business, to compare the airline to other competitor airlines and to forecast key variables.

To see some of these issues at play, let me present an example of the way that I used data when I was Chief Economist at Qantas. Towards the end of my tenure at the airline, I was tasked with providing forecasts of the operating revenue of the Qantas Group, and to present those forecasts at a monthly Board meeting. This took place during the Global Financial Crisis in 2009, which saw a 19.6% reduction in Qantas Group revenue over the financial year 2008/09, with revenue falling from AU$8,068m over the 6 months to December 2008 down to AU$6,484m over the 6 months to June 2009 – refer to Figure 1-1 below. To put this decline in revenue into perspective, when revenue falls by 1% this is a big deal for an airline, so you can just imagine how the senior management of the airline and the Board of Directors felt when they were presented with a 19.6% decline in revenue! They wanted to understand why revenue fell by such an extraordinary quantum, and how likely that reduction was to repeat itself in the 12 months ahead. My bosses at the time, the Chief Financial Officer Colin Storrie, and the Head of Strategy Grant Fenn asked me to prepare a board paper on the likely outcome for the economy and the Qantas Group's revenue over the year ahead.

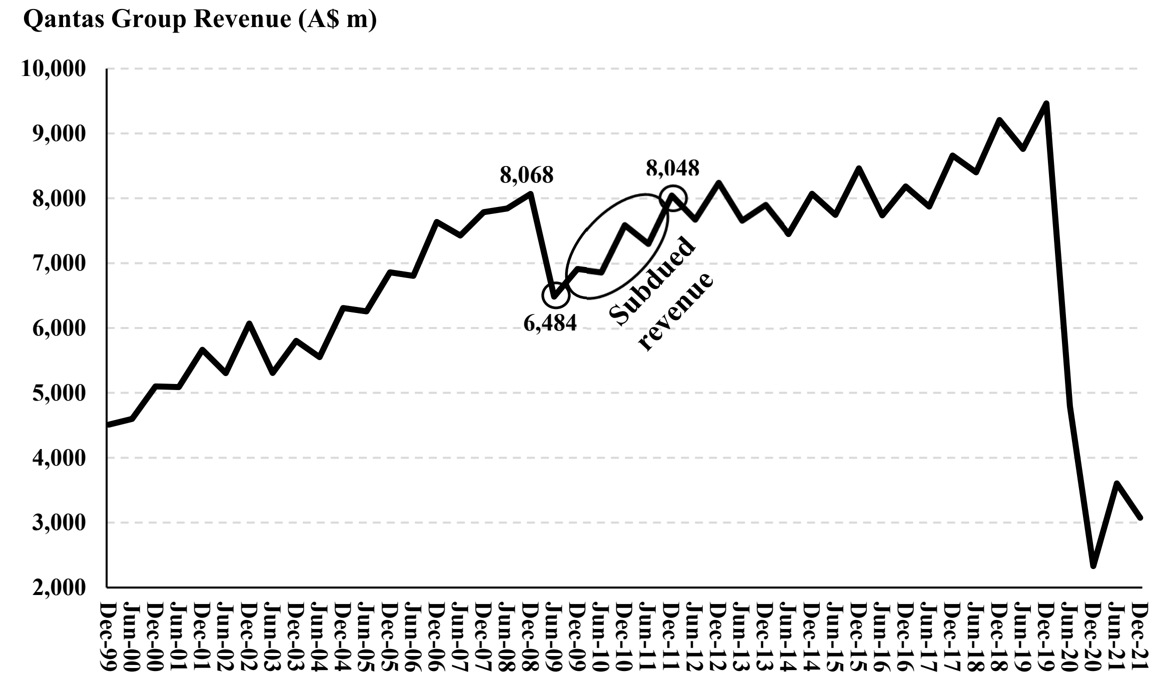

***Source:* Airline Intelligence and Research Database 2022**

Fig. 1-1: Qantas Group Biannual Operating Revenue – Dec-99 to Dec-21

My first step in pulling the revenue board paper and the presentation to the board together was to collect the right historical data on revenue and the right data with which to produce a forecast. If I was to understand the recent movements in the revenue data, I needed to collect the right data. This data was collected in an internal database held by the airline called the Corporate Data Warehouse, which was used as the source data for a program that was developed by the airline called Route Profitability. The Route Profitability program held the financial and operational data of the airline by route and by aircraft, including the data on revenue at the Group level that I required.

Once the revenue data was collected, I needed to eyeball the data. If you are tasked with presenting a forecast of data over the year ahead, then the first thing that you must do is eyeball the data, and the best way of doing this is to simply draw a line graph. At the time of producing the revenue report and forecast for the board, I only had data up to June 2009 available to me, so I only saw the reduction in revenue from A$8,068m in Figure 1-1 down to $6,484m, which is the first circled point in the graph. What I was unaware of was the subdued period thereafter, where revenue didn't recover to pre-global financial crisis levels until December 2011, which is the second circled point in the graph.

When I saw the graph of revenue for the first time, what struck me the most was that Qantas Group revenue had fallen enormously despite the Australian economy, as measured by its real Gross Domestic Product, holding up reasonably well. Figure 1-2 below presents the calendar annual growth rate in Australian real GDP between 2000 and 2019. We can see in 2009 that the Australian economy's growth had slowed to 1.9% (circled) when it's average or trend growth rate is around 3%, but it certainly was nowhere near the negative growth that many economies had experienced over this time-period. This suggested to me that Australian economic growth could not in isolation explain the almost 20% drop in Qantas Group revenue. I had to look for an alternative explanation.

While income is an important driver of airline demand, and this is often represented by an economy's gross domestic product, it is not the only important driver of airline demand. Another important driver is the wealth of passengers and of companies. This is because there are many consumers who are retired and don't finance their consumption through their labour income. These consumers finance their consumption through their wealth, and the returns to that wealth. These wealth effects are not necessarily picked-up in gross domestic product data.

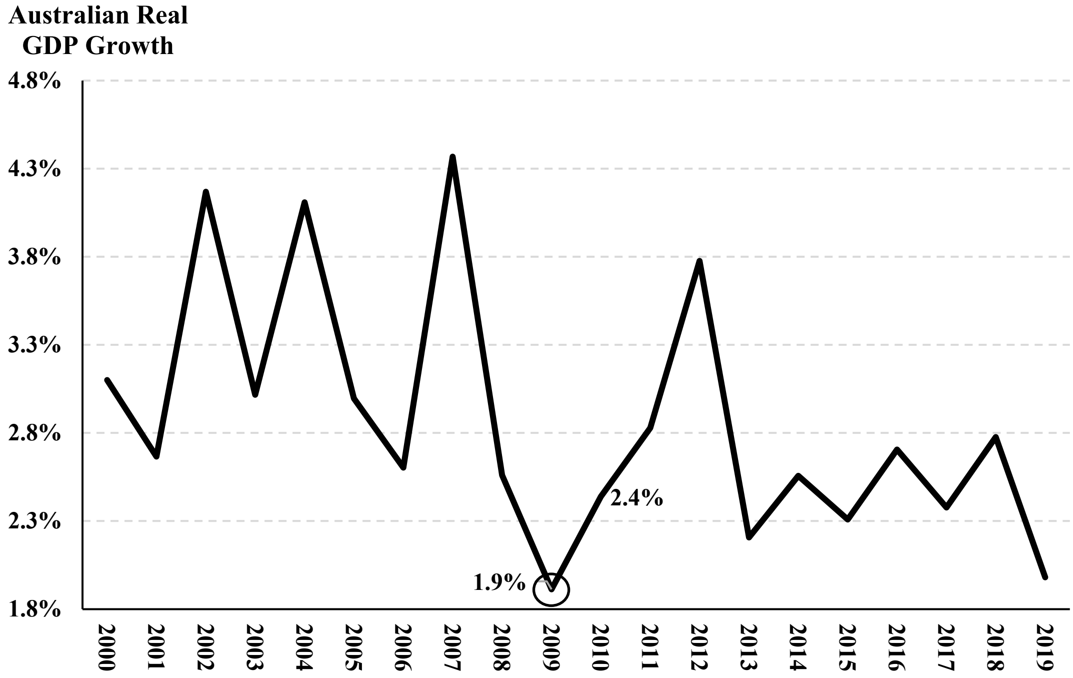

Source: **Australian Bureau of Statistics 2022**

Fig. 1-2: Australian Real GDP Growth – Calendar Annual period 2000 to 2019

One of the best indicators of the movement in wealth in the economy is the major stock market index of that economy. In Australia this is the ASX/S&P 200, which comprises the top 200 companies listed on the Australian stock exchange. This represents consumer wealth because households keep some of their wealth in company shares, and they hold a significant proportion of their assts in superannuation or pension funds, of which a large component is invested in the stock market. Stock markets also present information about financial confidence, for both businesses and consumers. From a statistical perspective, the ASX/S&P 200 was likely to do a better job of explaining what was happening to the Qantas Group's revenue simply because there was far greater volatility in the stock market than in Australian gross domestic product, and this volatility may perhaps explain what was happening to Qantas' yields and revenue.

Figure 1-3 below presents a line graph of the relationship between Qantas' passenger revenue per available seat kilometre (*PRASK*) and the ASX/S&P 200. Qantas' *PRASK* is a measure of the yield or price that it obtains from passengers, which we will discuss in more detail in chapter 2. We can see in the line graph that there is a positive relationship between the two variables. The looseness of the positive relationship tells us that there are many other variables that affect *PRASK* aside from the ASX/S&P 200, including both demand-side and supply-side variables, however some of the major movements in *PRASK* are indeed reasonably well explained by the ASX/S&P 200. When I spoke about this to the Qantas Board and the senior executive, I referred to this as the wealth effect. Qantas' demand was affected by both income and wealth effects, and up until the middle of 2009 the airline had only been affected by the wealth effect. The data was telling me that the income effect was about to arrive, and this could hang around for some time. This is indeed what eventuated.

Consider another data-based change that I initiated while I was Chief Economist at Qantas. When I first started at the airline back in 2004, one of the first things that I was asked to do was to investigate the way that the airline hedged its fuel exposure. The hedging approach of choice of the airline back in 2004 was the use of financial instruments such as swaps and call options. The use of instruments such as these can be problematic for at least two reasons. In the case of swaps, this allowed the airline to effectively lock-in an oil or jet fuel price. The problem with locking in a price is that it works swimmingly well when the market or spot price of jet fuel increases, but if the price goes down the airline loses money because they have locked themselves into paying a higher price for fuel. In the case of using call options, this avoids the problem of being locked into a price because the call option allows the airline to buy fuel at lower jet fuel prices, but the problem with call options is that they were, and still are, exceptionally expensive. When airlines use call options to protect themselves against higher fuel prices, they pay an option premium no matter what happens to the market price of jet fuel. This option premium is like paying an insurance premium against the risk that jet fuel prices increase. The level of the option premium depends on uncertainty about the direction in which jet fuel prices are headed. If this uncertainty is relatively high, or the jet fuel price is volatile, then the option premium that is paid by airlines to, usually, investment banks is high. In fact, during the Global Financial Crisis the cost of option premium typically exceeded US$20 per barrel.

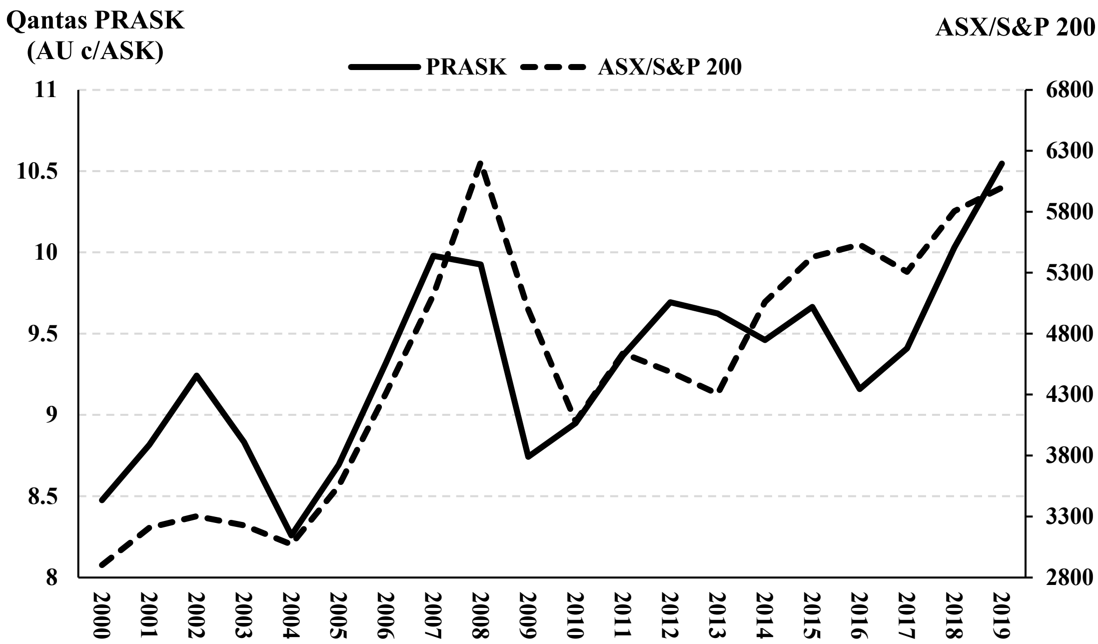

Source: **Airline Intelligence and Research Database 2022**

Fig. 1-3: Qantas PRASK Versus ASX/S&P 200 – Calendar Annual Period 2000 to 2019

One of the first things that I set out to do in 2004 when I started at Qantas was to determine whether there was a less expensive way to protect itself against higher jet fuel prices. I quickly learnt that the Qantas Group's fuel cost was essentially equal to the into-plane price of jet fuel multiplied by the Australian dollar price of 1 US dollar, multiplied by the quantity of jet fuel that the airline consumed. The airline therefore faced a dual exposure to both a commodity price and an asset price, the commodity price being jet fuel and the asset price being the Australian dollar price of 1 US dollar. Given this exposure, I wondered whether there was a connection between the jet fuel price and the Australian dollar exchange rate. If there was negative connection between the Australian dollar and the jet fuel price, then there was a possibility that the natural fluctuations of the Australian dollar price of 1 US dollar with the jet fuel price could act as a natural hedge against higher fuel prices. This means that the exchange rate would provide at least some protection against higher fuel prices for the airline, because each time the jet fuel price increases the Australian dollar price of 1 US dollar would fall.

What I found in the fuel price and exchange rate data was quite striking. There was a clear negative relationship between the Australian dollar price of 1 US dollar and the oil price. It didn't hold for all time periods, but it held for most time periods. My advice to the Qantas Group Treasury Risk Management team was that the Australian dollar offered round 40% protection against higher fuel prices, and that this would allow the airline to reduce the extent to which it needed to buy very expensive call options. Fortunately, the correlation between the jet fuel price and the Australian dollar held for most periods, as can be seen in Figure 1-4 below. It must be noted in Figure 1-4 that the secondary vertical axis, which is the vertical axis on the right-hand side, is in reverse order so that evidence that the lines move together in the same direction in the graph means that they move in opposite directions.

Another important piece of analysis that I conducted for the yield and revenue management area of the airline was the estimation of average airfares that maximise revenue on different routes. The yield and revenue management area used this information to benchmark there realised airfare outcomes to the estimated airfares that maximise revenue. If the two sets of airfares were close in proximity, then they could be reasonably confident that they were setting airfares at right levels. To determine the revenue maximising airfare, it was necessary to build a model of revenue for each route, and for this I mean come up with an approach to explaining route revenue. After deciding on an approach, it was then necessary to collect the data, which would be used to implement this approach. Once I had the data and the approach, I used statistical software to estimate a revenue equation that was shaped like a concave parabola, which allowed me to find the maximum turning point, which was the revenue maximising airfare.

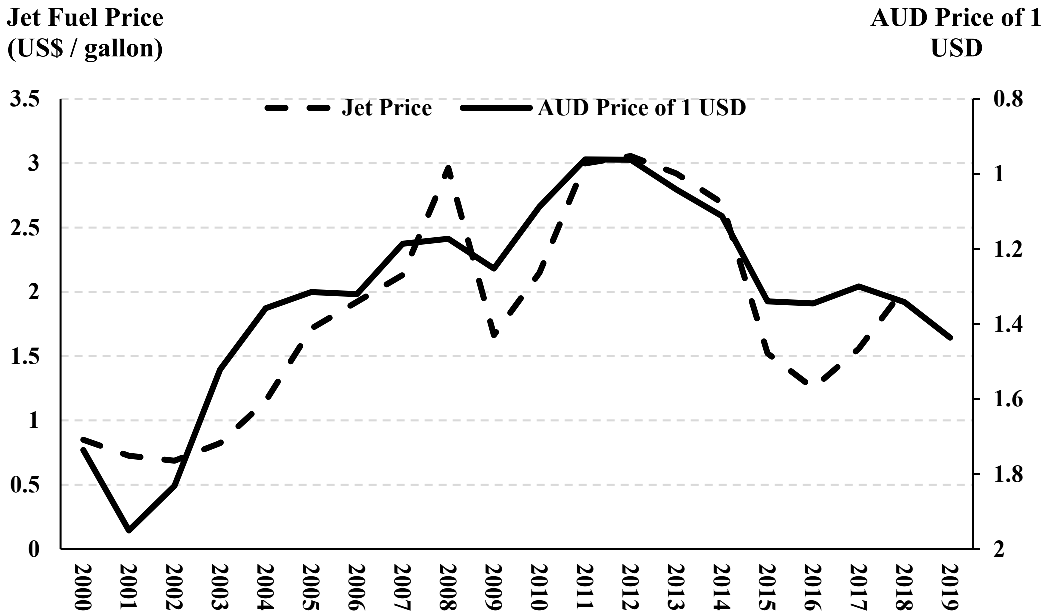

Source: **Energy Information Administration 2022, Reserve Bank of Australia 2022**

Fig. 1-4: Jet Fuel Price Versus Australian Dollar Price of 1 US Dollar

1.2 The Importance of Data for Airport Businesses

Of course, airlines are just one part of the aviation supply chain, although they are a very important part. The other important part are the airports. Like airlines, airports rely on the accurate collection of the right data, and the use of statistical techniques to extract the best information and insights from that data.

Data is incredibly important for airport businesses. Analysts in airport operations will need to collect data about the scheduled capacity of airlines, so that they can control the number of take-off and landings that occur, particularly around the peak times of operations, which is usually in the morning between 7.30am and 9am and in the afternoon between 4pm and 6pm. This data will also be required to determine the extent to which the runways, taxiways and aprons are being used for the purpose of scheduling maintenance for these pieces of infrastructure. The airport will also need to understand the maximum take-off weight of aircraft and the number of passengers that airlines are carrying into and out of the airport, as these variables are typically the basis for setting the aeronautical charges that airports impose on airlines. The number of passengers inside the terminal at any one point in time is also data that is likely to be collected by airports, because they need to understand the extent to which there may be points of congestion around the terminal, such as at passenger screening, the use of the terminal amenities, baggage reclaim, inter-terminal train services and the use of travelators, elevators and escalators.

Airports also collect data on the number of transit passengers expected to pass-through the airport because this has a significant impact on sales at merchants located within the terminal and occupancy at airport hotels, as transit passengers spend a significant amount of time roaming around the airport. Airports will collect data to analyse their energy and water consumption, because this in turn will help the airport to understand whether it is meeting its sustainability objectives. Airports will also collect data about the wind, rain, storm activity and fog at the airport because this affects the capacity of the runways to land aircraft and allow aircraft to take-off. Airports will also collect data on the number of hours that vehicles are using the airport carpark because this is an important determinant of non-aeronautical airport revenue.

1.3 The Remainder of this Book

The first topic that we must address in this book is the language and data of the airline business. The language of the airline business examines the enormous number of acronyms that airline and airport people use for different types of aviation variables, and for different types of aviation institutions. Chapter 2 will examine the acronyms that are used to describe the load on aircraft such as passengers carried (*PAX*), revenue passenger kilometres (*RPKs*) and revenue freight tonne kilometres (*RFTK*). It will analyse the language that is used for the capacity of airlines to carry passengers and freight, such as available seat kilometres (*ASKs*), available freight tonne kilometres (*AFTKs*) and available tonne kilometres (*ATKs*). Performance metrics such as on-time performance (*OTP*), reliability and the passenger seat factor (*PSF*) will also be discussed in detail. Basic financial variables will also be examined, including measures of yield (PRASK and *RASK*), unit cost (*CASK*), and earnings (*EBIT*, *EBITDAR* and *PBT*).

In chapter 3 I will take you through the different methods that are used to estimate growth rates. I consider this to be an incredibly important topic. During my time at Qantas and thereafter I would spend countless hours estimating the trend or business-as-usual growth rates for the airline, and I would see many instances in which analysts used the incorrect method for estimating growth rates. Chapter 3 will analyse four different approaches to estimating the long run average growth rate and the compound annual growth rate. They are the average annual growth rate (*AAGR*), the end-point compound annual growth rate (*EP-CAGR*), the log-slope compound annual growth rate (*LS-CAGR*) and the regression-slope compound annual growth rate (*RS-CAGR*). The chapter will persuade you to use the *RS-CAGR* because this generates the most accurate results, it is relatively simple to understand, to use and apply, and it is excellent if you need to estimate several growth rates at once.

When we analyse aviation data we often do so with the aid of a model. A key part of any model is its parameters and variables. When we build a model or an approach to understanding data, it is important that we have a strong grasp of the role of those variables. We will need to determine whether those variables are dependent or independent, exogenous or endogenous, and deterministic or stochastic. The many different types of variables that we encounter in aviation modelling is the topic of discussion of chapter 4. Alongside types of variables, chapter 4 will also discuss basic descriptive statistics, including frequency, relative frequency and cumulative frequency, and how these relate to measures of central tendency, variability, and covariability. These are key concepts that we will examine in detail in chapter 4.

The basis of much of the statistical analysis that we will perform with aviation variables and models is the probability distribution. The probability distribution measures the probability that a random or stochastic variable takes on different values. Probability distributions can either be empirical or theoretical. They are empirical when they are developed from frequency data that is observed in a sample. They are theoretical when they are built from theoretical formulae, as is the case for the normal, standard normal, log-normal, students-t, chi-squared, F, binomial, Poisson and uniform distributions. Chapter 5 will present all that you need to know about these distributions and discover how they can be used to better understand and draw insights from aviation data.

Much of what is done statistically with aviation data involves the conduct of hypothesis tests. Broadly a hypothesis test uses data from a sample to make an inference about a population. It does this by determining the extent to which information in a sample, say the mean of a sample, deviates from the hypothesised population parameter (mean in this case). The way that the aviation analyst determines whether the deviation of the sample mean from the population mean is sufficient to change beliefs about the population mean, is to determine where the information in the sample sits in a sampling probability distribution. The sampling probability distribution is the relative frequency of a sample statistic. We can use the sampling distribution to understand whether the sample information that we have obtained sits in the tails of the sampling distribution, which means it's deviation from the centre of the sampling distribution is statistically significant. If the sample statistic sits in the tail, which is usually defined as an area under the sampling distribution that includes 1%, 5% or 10% of values, then this suggests that perhaps the population parameter is not what it is believed to be. The sampling distribution and hypothesis testing are important parts of chapters 6, 7 and 8.

Chapters 9 and 10 introduce you to different ways of modelling time series aviation data. Aviation time series data consist of 5 elements – the trend, the cycle, seasonality, structural elements, and irregularity. The trend describes the long run direction, the cycle examines the movement around the trend, the seasonal component examines the peaks and troughs that occur in intra-year data, the structural elements are frequently occurring adverse and beneficial shocks, and the irregular component represents the part of the time series that cannot be explained. If aviation analysts wish to forecast aviation variables, then they must have a strong understanding of time series analysis and the tools that are required to effectively conduct time series analysis, including the autocorrelation and partial autocorrelation functions. Chapters 9 and 10 focus on linear time series models called *ARIMA* (Autoregressive Integrated Moving Average) models and the use of the autocorrelation and partial autocorrelation functions in identifying these processes.

Chapters 11 and 12 introduce readers to bivariate and multivariate regression analysis applied to both time series and cross-sectional data. Bivariate regression involves understanding the movements in a single dependent variable as a function of a single independent variable. Multivariate regression analysis involves understanding the movements in a single dependent variable as a function of multiple independent variables. In chapters 11 and 12 we will focus on least squares estimation of regression equations, and we will also discuss a range of tests that need to be conducted to determine whether these estimated regression equations are statistically healthy or not. Regression analysis is a tool that can be put to exceptionally effective use in aviation, for forecasting different aviation variables, for understanding the extent to which a variable may be causally related to another variable, and for comparing such critical variables as unit cost, margins, and productivity across airlines.

Chapter 13 introduces readers to the use of regression analysis to estimate relationships between variables using panel data, which is a combination of time series and cross-sectional data. Panel data is a more complex set of data to deal with than time series and cross-sectional data, and therefore requires a more sophisticated approach to modelling.

I hope that you develop your analytical and statistical skills sets by reading this book. If you work diligently at understanding the topics that are presented in this book, then the techniques that you learn can be put to exceptionally rewarding use while working for an airline, an airport, air traffic control or an aviation safety organisation to name but a few elements of the aviation supply chain.

CHAPTER 2

AVIATION METRICS

As an analyst for an airline or an airport, it is important that you understand the language of the aviation business and the key metrics and variables that aviation analysts consider. If your senior manager asks you to collect data on the airline's *ASKs* (available seat kilometres) and you don't understand what is meant by *ASKs*, then you will be placed in an awkward position. If you attend meetings and senior management and your peers are using aviation terms that you haven't heard of before, then it's difficult to engage in conversations in those meetings. Having worked for an airline for many years, attending meetings will be an integral part of your working life and it is an important medium for your boss to assess how you are performing. Aviation terms and acronyms are also used in aviation reports to Government officials and the investment community, aviation articles in the media, and internal communications to employees, so a knowledge of the key terms and acronyms is incredibly important across a wide range of activities in which you are likely to engage when you start your aviation career.

We start this chapter by examining those variables which represent the capacity of airlines to carry passengers and freight. They include the seats carried by the airline, the *ASKs*, the available freight tonne kilometres (*AFTKs*), the available passenger tonne kilometres (*APTKs*) and the available tonne kilometres (*ATKs*). We then examine metrics that represent the load that is carried by airlines, including the passengers carried (*PAX*), the cargo carried (*FRT*), the revenue passenger kilometres (*RPKs*), the revenue freight tonne kilometres (*RFTKs*), the revenue passenger tonne kilometres (*RPTKs*) and the revenue tonne kilometres (*RTKs*). By dividing the load statistics by the corresponding capacity statistics this generates an estimate of airline capacity utilisation. This includes passenger capacity utilisation also known as the passenger seat factor (*PSF*), freight capacity utilisation or the freight load factor (*FLF*), and total capacity utilisation or the total load factor (*TLF*).

An airline's unit cost is driven by both the price of resources that the airline employs, such as labour, fuel and aircraft, as well as the productivity of those resources. An improvement in productivity will lower an airline's unit cost, other things being equal, which will allow the airline to set lower prices or improve its profit margin. We will examine how the productivity of labour, fuel and fleet are calculated in this chapter.

This chapter also examines how we compute key operational performance measures of the airline, including on-time performance (*OTP*) and reliability. On-time performance determines the percentage of flights that leave or arrive on-time, which is usually defined as flights that leave or arrive within fifteen minutes of their scheduled departure and arrival times respectively. The reliability of an airline measures the extent to which its scheduled flights actually fly. This is an important operational performance metric for airlines that carry a high proportion of business-purpose traffic, because this type of traffic places a high level of importance on reliability as well as arriving and departing on time.

The last section of this chapter examines financial performance metrics. This starts with an examination of the many measures of passenger, freight, and total yield, all of which describe the price that passengers and freight pay for air transport services. The measures of passenger yield that we examine include passenger revenue per revenue passenger kilometre, the average airfare and passenger revenue per available seat kilometre or *PRASK*. Freight yield and total yield are then examined for the price that freight distributors pay for having freight transported by air and the price that passengers and freight combined pay for air transport services.

Low-cost carriers focus on being the lowest cost airline along particular routes. To achieve this end, it is necessary that low-cost carriers can estimate their unit cost. There are two measures of airline unit cost that we examine in this chapter. The first is operating cost per available seat kilometre (*CASK*) and the second is operating cost per available tonne kilometre (*CATK*). It is often the case that airlines will exclude fuel costs in the calculation of unit cost to take the volatility out of *CASK* and *CATK* estimates attributable to fuel prices. If fuel costs are excluded this generates *NFCASK* (non-fuel cost per available seat kilometre) and *NFCATK* (non-fuel cost per available tonne kilometre).

The final section of this chapter examines key measures of airline earnings, including earnings before interest and tax (*EBIT*), earnings before interest, tax, depreciation, and rentals (*EBITDAR*) and profit before tax (*PBT*). Enjoy learning about the language of the airline business in this chapter.

2.1 Capacity Data

2.1.1 Seats Carried

Seats Carried (*SEATS*) represents the total number of seats carried over a single leg. For example, if a Boeing 737-800 with 174 seats flies from Brunei's Bandar Seri Begawan (BWN) Airport to the main airport in East Timor, which is Presidente Nicolau Lobato International Airport (DIL), then the SEATS for that sector flown will be 174. A return journey represents double the seats carried, which is 348 seats. *SEATS* includes all revenue earning seats on the aircraft, excluding cabin crew seats and seats in the cockpit.[1]

Figure 2-1 below presents the *SEATS* across the entire network of Air Asia X, which is one of the biggest and most successful long-haul, low-cost carriers in world aviation with a hub out of Kuala Lumpur, Malaysia and subsidiaries Thai AirAsia, AirAsia India and Philippines AirAsia.

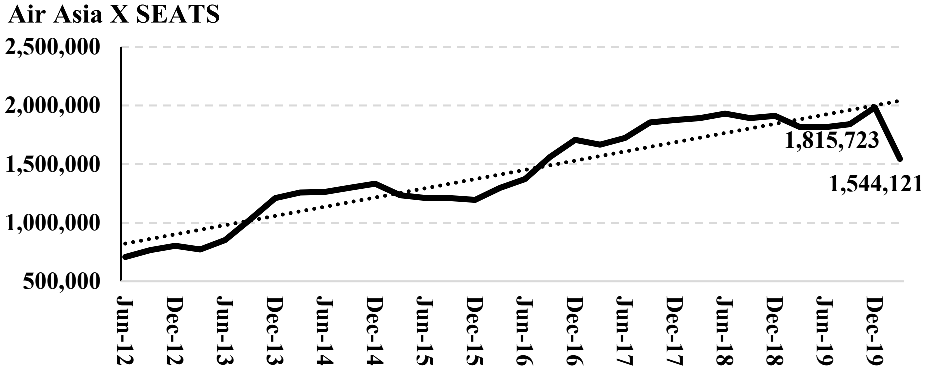

Source: Airline Intelligence and Research Database 2022

Fig. 2-1: Air Asia Seats Carried – June 2012 to March 2020

We can see that Air Asia X carried 1,544,121 seats during the March quarter 2020. The March 2020 *SEATS* of the airline has been adversely affected by the Coronavirus. This can be seen by comparing *SEATS* in March 2020 to the same quarter in the year prior, which is March 2019, where the airline carried 1,815,723 SEATS. The March 2019 *SEATS* are higher than the March 2019 SEATS by 17.6%.

Numerical Example

Air Asia X flight D7 206 departs Kuala Lumpur (KUL) at 23.55 and arrives on the Gold Coast (OOL) at 09:55+1 with a total flying time of 8 hours. The airline operates an A330-300 aircraft on the route with 377 seats. The same aircraft type is used for the return flight D7 207 departing OOL at 22.20 and arriving in KUL at 04:55+1. Both flights are operated daily. What is the number of seats carried over a year? The answer is:

$\text{SEATS}_{\text{KUL-OOL}} = 365 \times 377 \times 2 = 275{,}210$

2.1.2 Available Seat Kilometres (ASK)

This is the most popular metric in aviation that is used to describe airline passenger capacity. It represents the sum over the distance that an aircraft carries each available seat and is computed using the following formula:

$$\text{ASK} = \text{SEATS} \times \text{PASL} \qquad (2.1)$$

The *PASL* term in (2.1) is the passenger average sector length (see also section 2.2.3 below). This is the average distance that a seat is carried. The *ASK* expression at (2.1) tells us that *ASKs* can increase because the same aircraft with the same number of seats travels a longer distance (*PASL* is higher), or an airline travels the same distance on a flight but uses a bigger aircraft with more seats. Aviation management and analysts often refer to the use of a bigger aircraft as up-gauging.

Figure 2-2 below presents the available seat kilometres of Allegiant Airlines, which is an ultra low-cost carrier operating in the U.S. domestic market. Figure 2-2 indicates that 28,791m available seat kilometres were flown by Allegiant Airlines in calendar 2019 which is up from just 412m in 2002.

[1] You will notice throughout this book that I will provide information about the 3-letter IATA code that is used for airports. For example, BWN for Brunei's Bandar Seri Begawan Airport. It is a very good idea for you to become familiar with these codes for they are used extensively in aviation communications. If you wish to know a code for an airport, Google search "name of airport airport code", for example to find the airport code for Bandar Seri Begawan google search "Bandar Seri Begawan airport code".

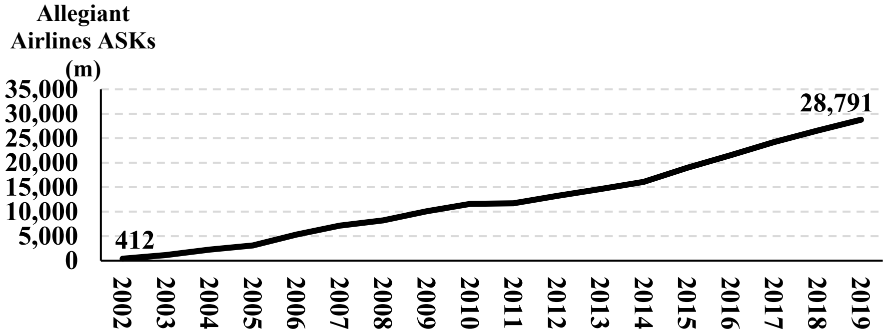

Source: Airline Intelligence and Research Database 2022

Fig. 2-2: Allegiant Airlines available seat kilometres – 2002 to 2019

Numerical Example

Returning to our Air Asia X example from section 2.1.1, let us try to estimate the annual number of *ASKs* flown by flights D7 206 and D7 207. We must first determine the average sector length between OOL and KUL. Using the Great Circle Mapper 2022 website, which is the best website on the internet for finding the great circle or flying distance, this is found to be 6,506 km (we will talk more about this website in section 2.2.3 below). We can now find the yearly *ASKs* for these flights by multiplying *SEATS* by distance as follows:

$ASK_{KUL\text{-}OOL} = 275{,}210 \times 6{,}506 = 1{,}791m$

2.1.3 Available Freight Tonne Kilometres (AFTK)

AFTKs is the most popular metric for the capacity to carry air freight over a certain distance. It represents the sum over the distance that an aircraft carries each available tonne of freight. It is calculated by multiplying the available tonnes of freight that an aircraft can carry by the average distance that the aircraft carries this available freight. This can be summarised as follows:

AFTK = Available Freight x Freight Average Sector Length (FASL) (2.2)

Figure 2-3 below presents the monthly *AFTKs* in the case of mainland China domestic services operated by Air China.

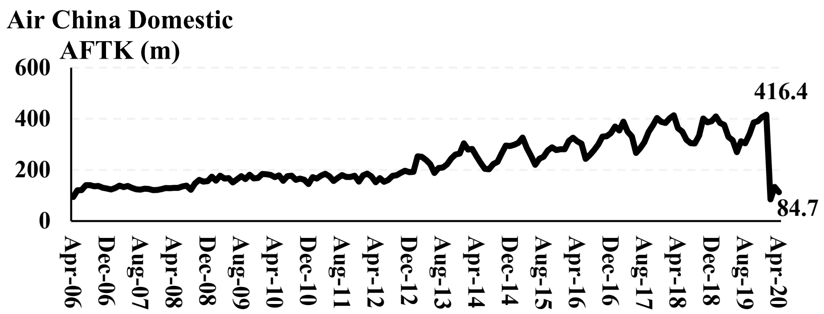

Source: Airline Intelligence and Research Database 2022

Fig. 2-3: Air China Available Freight Tonne Kilometres – April 2006 to April 2020

Figure 2-3 shows that available freight tonne kilometres for Air China domestic services was at a low of 84.7m in February 2020 because of the Coronavirus. This is down from a peak in domestic *AFTKs* of 416.4m in January 2020.

2.1.4 Available Passenger Tonne Kilometres (APTK)

APTKs is a measure of the capability of the aircraft to carry the weight of the aggregate of passengers over a certain distance. It is a statistic that is not routinely published by airlines. *APTK* is calculated according to the following formula:

APTK = APW × ASK (2.3)

where *APW* is the Average Passenger Weight, which is the average weight brought onto the aircraft by the passenger defined in tonnes, including own body weight, checked-in baggage and cabin baggage. The *APTK* metric is simply the passenger capacity of the aircraft defined in seat kilometre terms multiplied by the average weight brought onto the aircraft by the passenger in tonnes.

Figure 2-4 below presents the monthly time series movements in the *APTKs* of Hainan Airlines, which is one of the biggest airlines in Mainland China with headquarters in Haikou, Hainan.