

An empirical analysis of route-based differences in Australian air fares

David Mitchell¹

¹Bureau of Infrastructure, Transport and Regional Economics,
Department of Infrastructure, Transport, Regional Development and Cities, Australia

Email for correspondence: david.mitchell@infrastructure.gov.au

Abstract

This paper presents the results of empirical analysis of differences in domestic airfares across Australian domestic commercial trunk air routes, and provides some estimates of differences in fares across regional routes in Australia. The analysis considered a range of potentially relevant factors, including world oil prices, route distance, number of flights (by route), total route passengers, total number of route operators, route load factor and regional populations.

The results show that all of these factors have a statistically significant impact on average fares, and together they explain close to 95 per cent of the variation observed in collected airfares. The results also imply the presence of scale economies in airline pricing with respect to route distance, market size and load factor, but diseconomies with respect to the number of flights. Importantly, route competition, delineated by the number of different airlines operating on a route, has a statistically significant effect in reducing air fares.

The results also suggest that, taking into account all of the factors above, there remain statistically significant differences in average fare levels across trunk routes, which align into three broad groups: i) high mark-up routes, ii) mid-tier mark-up routes, and iii) low-mark-up routes.

1 Introduction


Prior to domestic aviation deregulation introduced in October 1989, Australia's aviation market was a regulated duopoly with incumbents—Ansett Airlines and Trans Australia Airways (TAA)—shielded from new competitors on trunk routes and fares reviewed and set by the Federal Government's Independent Air Fares Committee (IAFC). Deregulation involved removing restrictions on new market entrants and letting airlines set fares unimpeded. Upon deregulation, fares decreased significantly across most routes, most particularly on very long-distance routes.

Australian government monitoring of domestic air fares commenced in October 1992. Over the 30 years since domestic deregulation, while full economy and business class fares have increased slightly in real terms, real discount airfares have fallen almost 50 per cent below equivalent fares in 1993 (see Figure 1). The entry of low cost carriers in the early-2000s, initially Virgin Australia and subsequently Jetstar Airways and Tiger

Airways, resulted in increased competition on trunk routes and periods of intense competition for market share. This resulted in significant reductions in real airfares, particularly best discount fares. For example, the period between mid-2008 and mid-2011 exhibits significant reduction in both nominal and real best discount fares, before fares stabilised somewhat.

At the route-specific level, trend movements in fares vary across different routes, and in some cases are different from the national trends. This is addressed in subsequent sections. (Note that all of the subsequent analysis of movements in fares presented in the remainder of this paper is in nominal terms, i.e. fares have not adjusted for inflation.)

Figure 1: Trends in nominal and real Australian domestic air fares, 1991 to present



Source BITRE (2018).

1.1 BITRE airfares collection

BITRE's airfares collection (BITRE 2018) contains monthly records of the best available business, restricted economy and best discount fares for each airline across the 70-largest routes. Importantly, the fares recorded in the BITRE collection are the lowest available fare for each fare class on the last Thursday of each month, for a hypothetical trip in three weeks' time (i.e. the third Thursday of the following month). Hence, the BITRE's fares collection cannot provide any insight into the time-profile of fares for a particular flight in the lead up to flight departure (discussed further below).

Routes covered in BITRE's collection vary over the collection period, based on variations in the composition of the top 70 routes. In particular, the set of routes covered in BITRE's collection do not include many of the lower volume routes.

1.2 How airlines set fares

Though details about how Australian airlines set fares is not generally publicly available, it is apparent that nearly all airlines today use ‘sophisticated’ yield management techniques to optimise operating revenues. Yield management techniques typically involve using historical data and real-time booking information to vary the menu of available fares, by fare class, for each flight up until just prior to flight departure. From an airline’s perspective, aircraft seats on any particular flight are a finite and limited-life resource—once the flight departs any unsold seats disappear. Therefore, it is in airlines’ interests to sell as many seats as possible at as high a price as possible to maximise revenue, hence the availability of ‘last minute’ and ‘mystery’ flight deals. The corollary to this is that as passenger bookings on any single flight increase, capacity diminishes and seats are likely to become more valuable and hence prices can rise. Hence, any unanticipated increase in demand on a route can cause fares to increase significantly.

In the medium term airlines can vary the overall level of fares on a particular route by adding more capacity, through increasing the number of flights or operating larger aircraft. For example, competition for increased market share is conducted through increasing available seat capacity which has a flow-on effect to fares—in the short term, fares are subject to seat availability.

Also, passenger preferences vary with regard to price and service quality. Airlines take advantage of this by offering a mix of different fare classes. First/business class tickets, for example, cater for travellers who are relatively price insensitive, time sensitive (last-minute booking) and/or willing to pay a premium for extra services. Other travellers may be both price and time sensitive, for instance, business travellers able to book in advance, but not far enough in advance to secure the lowest fares. Leisure travellers, on the other hand are typically both price and time insensitive (by time of day rather than travel date), and seek the cheapest available fare.

BITRE’s fares collection does not provide the multiple snapshots, over several points in time during each month, which would be needed to provide an indication of how fares change in the lead up to flight departure. Moreover, multiple sampling of air fares at different points in time for a particular flight, does not provide the accompanying booking information necessary to fully understand the factors that influence the variation in fares observed by customers.

1.3 Air fares and airline cost structures

In markets where there is some degree of competition, or even risk of new market entry, changes in prices tend to reflect changes in input costs. The major input costs for airlines are capital, maintenance and parts, labour, fuel and air navigation and airport charges. Capital costs mainly comprise aircraft and airport terminal leasing costs. Labour includes pilots, aircraft cabin personnel, maintenance engineers and administrative staff. Nominal cost shares for Australia’s two major domestic airlines, Qantas Group and Virgin Australia, in 2016–17 imply fuel costs represent between 17 and 21 per cent of total costs, labour is between 23.6 and 27.5 per cent, capital-related costs are around 34–35 per cent and other costs around 17 per cent for Qantas and 26 per cent for Virgin Australia.

Input costs, and other factors, will also vary across different routes, thereby resulting in some systematic differences in fares across routes. For example, fuel costs may be a higher share of variable costs for shorter-distance routes than longer-distance

routes, due to the disproportionately higher fuel use involved in aircraft take-off/landing movements. Similarly, average fuel costs may be a lower share of overall costs on higher-volume routes, where higher total fuel costs can be defrayed over proportionately more passengers. Increased route competition may lead to lower-than-average fares relative to routes where there is much less competition. Finally, variation in the characteristics of aviation demand across different routes may also affect the level of fares. For example, on routes where there are fewer (or no) time-comparable modal alternatives, such as very-long distance routes, passenger demand may be less elastic—i.e. less responsive to changes in fares—airlines may be able to set prices with a small premium. Conversely, on routes where demand is more elastic—i.e. more responsive to changes in fares—fares may be comparatively lower.

1.4 Paper structure

The remainder of this paper is structured as follows. Section 2 outlines trends in average domestic airfares in Australia since 1992 and considers trends in fares across the various routes covered by BITRE's fares collection over that period. Section 3 briefly describes the analytical methodology used to modelling fares across different routes. Section 4 briefly outlines the key raw data and main data sources. Section 5 presents some of the key results. Section 6 uses the empirical results outlined in Section 5 to provide a comparison with fares on lower-volume, regional routes. Finally, Section 7 draws out some of the implications of the results and provides some concluding remarks.

2 Trends in Australian domestic airfares

2.1 Trends in airfares by route

Appendix Figure A.1 shows trends in best discount airfares since 1992, for all BITRE-monitored top-70 domestic aviation routes (by number of passengers). As a result of variation in the composition of the top-70 routes over that period, BITRE has collected fares on over 130 separate routes. Consequently, for the largest 30-odd routes BITRE's collection provides a complete record of the best available monthly airfare, on smaller volume routes (outside the top 30 or so routes), there are often significant periods for which no observations are available. Examples of the latter include Adelaide–Alice Springs, Adelaide–Kalgoorlie, Albury–Melbourne, Brisbane–Hobart, Carnarvon–Perth, Esperance–Perth, Melbourne–Burnie and Sydney–Norfolk Island, among others. For some of these smaller routes, there are likely to be insufficient observations to draw reliable conclusions about the factors influencing air fares on these routes.

Closer consideration of nominal fares by route suggests a few broad differences in fare trends across different routes. Firstly, across many routes, best discount fares have remained relatively stable, with little monthly variation—this appears to be particularly so on major intercapital routes, e.g. Sydney–Melbourne, Sydney–Brisbane, Melbourne–Brisbane, Adelaide–Brisbane, Sydney–Perth, etc. By contrast, on many lower-volume and/or non-intercapital routes, average best discount fares exhibit significantly more month-to-month variation. Examples include Sydney–Tamworth, Sydney–Port Macquarie, Sydney–Coffs Harbour, Sydney–Dubbo, Perth–Geraldton, Melbourne–Devonport, Melbourne–Hobart, Sydney–Albury and Adelaide–Port Lincoln.

Secondly, on many of the major intercapital routes, nominal average best discount fares exhibit no discernible trend, either up or down, over the observation period. However, for a small proportion of routes average best discount fares exhibit a long-term increasing trend (in nominal terms).

In many cases, there appears to be a peak (or spike) in fares before the route disappears from the records, which may reflect airlines attempting to improve profitability by lifting fares before abandoning a route, but it has not yet been verified whether this is indeed the case.

2.2 Australian air fares, crude oil prices and airline hedging

Anecdotal evidence about airline pricing behaviour suggests that not only do airlines use yield management methods to maximise revenue for each flight, but airlines also forward hedge fuel purchases to protect against volatile movements in fuel prices. Hedging limits the fuel cost impact on airlines in periods of increasing fuel spot prices, but equally limits the scope for airlines to take advantage of future reductions in spot fuel prices.

Further, different airlines hedge to different extents and using different contracts/strategies. Details are typically not made publicly available as hedging potentially provides a competitive advantage, by enabling airlines to lower costs *vis-a-vis* their competitors. Qantas' annual report (Qantas 2017, p. 78) notes that fuel consumption may be hedged out to two years ahead within specific parameters, or more with management board approval. Similarly, Virgin Australia's annual report (Virgin Australia 2017, p. 79) notes that, subject to limits determined by its management board, it hedges anticipated fuel consumption to protect against sudden and significant increase in fuel prices whilst ensuring that the airline is not disadvantaged by a significant reduction in fuel prices.

This suggests that it should not surprise that fares exhibit little short-term (e.g. monthly or quarterly) correlation with changes in international crude oil spot prices, but one might expect some degree of longer-term correlation between average fares and long-term prevailing crude oil prices. It also suggests some degree of lagged ('sticky') relationship between fuel costs and crude oil spot prices. Other input costs may also exhibit a degree of 'stickiness', for example, wages tend to be agreed in two- or three-year contracts with annual increments.

Significantly, average best discount (nominal) fares appear to exhibit far less month-to-month variation than world oil prices (not shown here). This in part reflects the fact that, while fuel costs are presumably a significant share of individual flight costs, they comprise only a part of the total cost of air services—for example, as previously mentioned, Qantas and Virgin Australia's fuel costs are around 20 per cent of total operating costs.¹ Airline fuel hedging practices will also act to dampen the impact of changes in world oil prices on airline operating costs, and hence fares.

These factors also suggest that there may be only a weak relationship between air fares and world fuel prices.

¹ Qantas' 2016–17 Annual Report indicates aviation fuel costs were around 21 per cent of total operating expenditures and Virgin Australia's 2016–17 Annual Report implies fuel costs were around 19 per cent of its total operating costs. See Appendix C for further details.

3 Methodology

The analysis was undertaken using a panel (cross section–time series) regression model to estimate systematic difference in average airfares across different routes. The analysis considered movements in the lowest fare on each route. Additional separate models were estimated to test for differences in fares across routes by fare class and airline.

Relevant factors considered in the analysis include:

- world oil prices (and exchange rates)
- route distances
- total flights (by route)
- total passengers (by route)
- total operators (by route)
- load factor
- regional populations.

Differences in fares across different routes manifest as systematic differences in the modelled average fare level across routes and/or systematic differences in how fares change with respect to different factors across routes. We test these for systematic differences in these parameters to gauge whether, and if so, how large the differing impacts are across routes.

The preferred empirical model specification is shown in the equation below. For the models with fares split by fare class and airline, we also include separate fare class and airline specific factors.

The simple static empirical model specification is:

$$y_{it} = \alpha_i + \beta x_{it} + \sum_{j=0}^J \gamma_j w_{t-j} + \delta z_i + \varepsilon_{it}, \quad \varepsilon_{it} \sim \text{iid}(0, \Omega)$$

where

y_{it} – denotes the natural logarithm of average fare per route kilometre for route i at time t .

x_{it} – denotes the natural logarithm of route-specific factors—i.e. total flights, total passengers, total operators and load factor—for route i at time t .

w_t – denotes the natural logarithm of airline inputs costs at time $t - j$ —the current analysis includes only oil prices.

z_i – denotes the logarithm of time-invariant route-specific variables (e.g. route distance)

α_i , β , γ_j and δ – are model parameters.

ε_{it} – denotes the random error, which is assumed to be distributed with mean zero and variance–covariance matrix Ω .

4 Data

As previously noted, the BITRE's fares collection contains monthly records of the best available business, full economy (till February 2015), restricted economy (from March 2003) and best discount fares for each airline across the 70-largest routes. The collection extends back to October 1992. Routes covered in the collection vary over the collection period, in accord with variations in the top 70 routes. Importantly, BITRE's collection does not include many lower-volume regional routes.

Airlines covered in BITRE's collection include:

- Qantas
- Jetstar Airways
- QantasLink
- Virgin Australia Airlines
- Tiger Airways
- Regional Express (Rex) Airlines

The fares data was combined with monthly BITRE data on passengers, flights, number of operators and load factors, by airline. World oil prices were sourced from the World Bank (2018) commodity database, which reports spot prices for UK Brent, Dubai and West Texas Intermediate crude oil. Regional populations were sourced from the Australian Bureau of Statistics' regional population data, and matched to airport according to BITRE-defined regional catchment areas.

The data comprise an unbalanced time series cross-section data set, which may be estimated using panel data estimation methods. All models were estimated on log-transformed data, and hence all parameter estimates may be directly interpreted as elasticities.

5 Empirical results

5.1 Best discount fares by route

We first modelled movements in average best discount airfares (per route kilometre) across all Australian routes—averaging fares by route distance enables a more direct comparison of relative differences in fares across different distance routes—and tested the significance of all parameters and for panel-data specific effects. The results imply there are statistically significant route-specific fixed effects—i.e. statistically significant differences in the average fare level across different routes—and also significant time-invariant fixed effects with respect to route distance. All other terms were also statistically significant.²

Two sets of fixed effects are included in the preferred specification: i) route-specific dummy variables, and ii) seasonal (monthly) dummy variables. The preferred specification includes lagged oil prices.

² We also tested a dynamic specification, including one-month lagged fares, but due to periodic gaps in the air fares data, the inclusion of a lagged dependent variable term significantly reduced the number of complete observations and number of routes analysed in the model, and hence was not further investigated.

We also tested the significance of adding airport-specific variables to the panel data model and replacing the route-specific variables with airport-specific variables. The former does not significantly improve the model fit, while the latter does not fit the data as well as the route-specific variable model. Moreover, because of missing observations in the data, airport-specific constants could not be estimated for all airports.

Table 1 shows summary estimation results, excluding route-specific fixed effects and time-specific dummy variables, for the preferred route-specific best discount fares model specification. The preferred specification includes a general time trend term, lagged oil prices (up to six months). Two sets of results are presented, ordinary least squares (OLS) and (feasible) generalised least squares (GLS) estimates, which better account for differences in variance across routes.

The model results show that both specifications fit the data quite well, the R^2 values imply the model explains over 96 per cent of the observed variation in average fares. The GLS specification is preferred, as the estimates are both consistent and efficient.

All of the variables included in the model are statistically significant and the effects relatively consistent across all model variants. We briefly outline the key effects below.

The model results show that both specifications fit the data quite well, the R^2 values imply the model explains over 90 per cent of the observed variation in average best discount fares.⁵ The GLS specification is preferred, as the estimates are both consistent and efficient.

All of the variables included in the model are statistically significant and the effects relatively consistent across all model variants. We briefly outline the key effects below.

5.1.1 Route distance

The estimated route distance elasticity is approximately -0.57, and implies that the average air fares decline by about 6 per cent for every 10 per cent increase in route distance. In other words, average fares are generally lower for longer distance routes, all else equal. As previously mentioned, this accords with a priori expectations, as non-distance related flight costs can be defrayed over longer distances.

5.1.2 Number of passengers

The passenger volume parameter suggests there are significant scale economies in airline pricing, with average fares declining on more heavily trafficked routes. The impact is not only statistically significant but also substantial—for every one per cent increase in route passenger volume, average fares decline by around 1.75 per cent.

5.1.3 Number of route operators

There is also a statistically significant observed competition effect, with average fares declining as the number airlines operating on a route increases. The implied elasticity is approximately -0.31, which implies that average fares on a route with three operators will be about 15 per cent lower than average fares on a route with just two operators, all else equal. Likewise, average fares on a route with four operators will be about 7 per cent lower than equivalent on a route with three operators.

5.1.4 Number of flights and aircraft size

Conversely, average fares increase as the number of flights and the average aircraft size employed in each route increase, somewhat offsetting the scale economy effects.

Both effects are statistically significant, with the elasticity of average fares to the number of flights approximately 1.6 (i.e. a 1.6 per cent increase in average fares for every one per cent increase in the number of flights) and the elasticity of average fares and average aircraft size 1.4.

Table 1: Estimation results – best discount airfares by route

Dependent variable Independent variables	log (Fare per km)	
	OLS	GLS
Constant	-185.743*** (11.771)	-195.785*** (10.654)
Log (Route distance)	-0.540*** (0.035)	-0.574*** (0.027)
Route load factor	0.033*** (0.003)	0.026*** (0.002)
Log (Avg. aircraft size)	2.011*** (0.188)	1.427*** (0.149)
Log (Flights)	2.199*** (0.188)	1.583*** (0.149)
Log (No operators)	-0.362*** (0.012)	-0.308*** (0.012)
Log (Route passengers)	-2.315*** (0.187)	-1.751*** (0.148)
Log (Trend)	24.695*** (1.553)	26.141*** (1.404)
Log (Oil price _t)	-0.136*** (0.033)	-0.174*** (0.025)
Log (Oil price _{t-1})	0.543*** (0.206)	0.956*** (0.161)
Log (Oil price _{t-2})	-0.282 (0.207)	-0.701*** (0.161)
Log (Oil price _{t-3})	0.405** (0.202)	0.390** (0.157)
Log (Oil price _{t-4})	-0.657*** (0.201)	-0.205 (0.155)
Log (Oil price _{t-5})	0.638*** (0.200)	0.391** (0.153)
Log (Oil price _{t-6})	-0.588*** (0.136)	-0.463*** (0.104)
Summary statistics		
Observations	11,010	11,010
R ²	0.817	0.901
Adjusted R ²	0.816	0.900
Residual Std. Error (df = 10914)	0.256	0.081
F Statistic (df = 95; 10914)	514.418***	1,043.847***

Note: Significance levels: *p<0.1, **p<0.05, ***p<0.01. Standard errors in parentheses.

5.1.5 Oil prices

The preferred model specification includes lagged oil price terms up to 6 months from the current period, all of which, with the exception of the two-period lag parameter, are statistically significant. While the significance of multiple lag effects appear consistent with the dynamic and potentially varying hedging strategies applied by different airlines, the estimated lag effects themselves alternate in sign, somewhat cancelling out over the entire period. The combined oil price impact (i.e. the sum of all current and lagged oil price terms) is +0.195, which implies that average air fares broadly increase on average with increases in world oil prices. Moreover the value of 0.195 is also broadly consistent with fuel's share (about 20 per cent) of airline operating costs (at least for full service airlines).

5.1.6 Load factor

Average route load factor is also statistically significant and positive, implying that average fares increase as average aircraft occupancy increases, which accords with expectations. Load factor is a derived variable (passengers divided by seats), that captures the average proportion of occupied seats on flights across a route. Other things equal, where load factors are low airlines will have incentives to price competitively to fill empty seats, and vice versa when load factors are high.

5.1.7 Route-specific fixed effects


Route-specific fixed effects (dummy variables) are statistically significant and imply there are systematic differences in average fares across routes, over and above the impacts of distance, oil prices, passenger numbers, competition, etc. outlined above. Figure 2, below, shows the route-specific fixed-effects parameter estimates (natural logarithm airfare - \$/km), together with the two-standard error confidence interval, for each route.

Immediately apparent from the figure is the large spread of values and that only for a handful of routes are fixed effects not significantly different from zero. Examples of the latter include Sydney–Tamworth, Port Macquarie–Sydney, Darwin–Melbourne and Cairns–Townsville. Routes where average fares are systematically below average (i.e. where the route-specific constant is less than zero) include Hervey Bay–Sydney, Melbourne–Newcastle, Newcastle–Gold Coast and Brisbane–Newcastle—these routes have the largest negative fixed effects parameters. Conversely, routes where average fares are systematically above average (i.e. a route-specific constant greater than zero) include Perth–Port Hedland, Karratha–Perth, Brisbane–Mount Isa and Canberra–Melbourne—these routes are among the largest positive fixed effects parameter estimates.

We also consider the route-specific constants against route distance in Figure 3, below, which yields some additional insights into how airfares vary across different commercial routes. In particular, we overlay the plot with several contour lines, with slope equal to the estimated distance elasticity, and group routes into three broad groups (highlighted by the shaded ellipses).

The juxtaposition of the route-specific constants and distance elasticity contour lines suggests not only is there an inverse relationship between the average fare (per kilometre) and route distance, there is also an inverse relationship between the average fare mark-up (i.e. route-specific fixed effect) and route distance, such that the shorter the route distance, the higher the average fare mark-up.

Figure 2: Route specific fixed effects - Best discount fares by route




Moreover, there appear to be at least three, and possibly more, different groups of routes, which we characterise as follows:

- High-mark-up routes – which feature Perth, Darwin and several remote mineral industry routes.
- Mid-tier mark-up routes – which include most of the higher volume (trunk) domestic commercial routes and several other smaller-distance routes.

- Low-mark-up routes – which include Bass Strait routes and some short regional or tourist routes, that are all shorter distance routes for which car, or ferry in the case of the Tasmanian routes, is a more significant competitor.

Figure 3: Route specific fixed effects - Best discount fares by route and distance



High-end mark-up routes include:

- Sydney/Melbourne/Brisbane/Adelaide-Perth
- Perth-Newman/Karratha/Broome/Port Hedland
- Sydney/Melbourne/Brisbane/Adelaide-Darwin
- Perth-Darwin
- Alice Springs-Darwin
- Sydney/Melbourne-Cairns
- Brisbane-Mount Isa

These routes are distinguished by either being longer-distance routes to/from Perth/Darwin/Alice Springs, or they are routes serving more remote mining centres (e.g. Newman, Karratha, Port Hedland, and Mount Isa).

Mid-tier mark-up routes, include many of the trunk domestic commercial routes:

- Sydney–Melbourne/Brisbane/Adelaide
- Melbourne–Sydney/Brisbane/Canberra
- Brisbane–Sydney/Melbourne/Adelaide/Canberra
- Adelaide–Sydney/Brisbane/Canberra
- Canberra–Melbourne/Brisbane/Adelaide (and Canberra–Sydney marginally)
- Gold Coast–Sydney/Melbourne

and some shorter-distance regional routes and longer-distance tourist routes, such as:

- Perth–Geraldton/Kalgoorlie
- Sydney/Brisbane–Hamilton Island
- Brisbane–Mackay/Emerald/Moranbah
- Hobart–Sydney/Brisbane

Thirdly, low-end mark-up routes include many shorter-distance regional and tourist routes, which appear to be more susceptible to greater competition from other modes, particularly private car travel. Such routes include:


- Melbourne–Adelaide
- Melbourne–Burnie / Devonport / Launceston / Hobart (ferry competition)
- Sydney–Albury / Ballina / Coffs Harbour / Hervey Bay / Armidale / Port Macquarie / Tamworth / Dubbo
- Brisbane–Bundaberg / Proserpine / Rockhampton
- Melbourne–Mildura

As already noted, data for many smaller-volume, regional routes is not available and hence it's not possible to draw any conclusions directly about relative fares for many smaller regional routes. (In Section 6, recent fares on smaller-volume, regional routes are compared with model-predicted fare levels.)

5.1.8 Seasonal effects

Seasonal effects are also statistically significant in fares. Figure 4 shows the seasonal (monthly) dummy variable estimates for fares by route. All monthly dummy variables are relative to January average fares. The results show a clear seasonal pattern, with fares on average higher in months including school/seasonal holiday periods—March, June, September and December. The vertical lines around each point estimate show the two-standard error confidence intervals for each estimate, and these show that the March, September and December fares are, on average, significantly above average fares in all other months. Similarly, fares for flights booked in January are significantly below average fares booked in all other months.

Figure 4: Monthly dummy variable estimates



5.2 Route–airline and route–fare class effects

BITRE’s fares collection also includes the best available monthly fares by route and airline, and by route and fare class, and we used this data to also model variation in fares across routes and airlines and routes and fare class, using the same explanatory variables used in analysing best discount fares by route (above):

- Including airlines in the regression allows for inclusion of more complex interactive effects, such as differences in how fares vary with distance across airlines. We tested the statistical significance of airline-specific effects with respect to distance, total route passengers, average route load factors, and the number of operators on the route.
- Jointly modelling all fare classes in the one specification allowed for estimation of cross-fare effects, such as differences in the behaviour of fares across different fare classes. We tested the statistical significance of fare class-specific effects with respect to distance, total route passengers, average route load factors, number of route operators, oil prices and seasonal differences.

While space constraints prevents presentation of the full results here, the results imply there are statistically significant route–airline specific fixed effects—i.e. statistically significant differences in the average fare level across different routes and across different airlines—and also statistically significant route–fare class effects.

6 Comparing fares on lower-volume regional routes

In order to assess the relative level of fares on lower-volume regional routes, BITRE collected a wider sample of fares as part of the July 2018 fares collection. The collected fares represent the cheapest fare available, by airline, across all city pairs with domestic passenger traffic recorded in April 2018. The fares were for a prospective departure date of 26 July 2018 and return flight on 9 August 2018. The

expanded collection covered over 280 separate domestic air routes, yielding useable fare information for approximately 245 routes.

We then calculated the difference between the best discount average fare (divided by route distance) across each route and the best discount average fare predicted by the preferred model specification (outlined in Section 5, above), excluding the modelled route-specific factors. The resulting difference between the actual and modelled average fares represents a measure of the relative route-specific mark-up after taking into account all other relevant factors (i.e. distance, passenger volumes, number of operators, load factor and seasonal factors).

Because the estimates are based on a one-month snapshot of fares, the results may not reflect longer-term trend differences in fares across different routes. For example, unseasonably or unusually high bookings on any single route may result in higher than average available fares at the time of collection and, conversely, unusually low confirmed bookings on any route at the time of collection, may result in below average quoted fares. Also, the magnitude of the relationship between airfares and the various modelled factors (i.e. distance, passenger volume, load factor, aircraft size, etc.) may differ between the 'top-70' routes and lower-volume regional routes. Hence, the model may not predict fares on lower-volume regional routes as accurately as for the 'top-70' routes.

6.1 Interstate and intrastate aviation governance arrangements

Commonwealth, state and territory governance arrangements also appear to influence relative fares across different routes. Where Commonwealth legislation ensures operation of interstate aviation services are unregulated, varying legislative and regulatory regimes apply to intrastate services in some states and territories, which may systematically affect fares across different routes. Space prevents discussion of these various schemes here, but the results presented in Section 6.2, grouped by the various intrastate regulatory regimes and interstate services, appear to suggest these schemes have a material effect on prevailing fare levels.

6.2 Route-specific difference between actual and modelled fares

Figure 5, below, shows the estimated difference between actual and modelled average return airfares, for the month of July 2018, grouped according to the jurisdiction and regulatory regime governing each route. As previously noted, the estimated difference between the actual and modelled fares provides an implicit indicator of the relative fare mark-up across different routes. For the purposes of comparison, Figure 8 also includes several average fare–distance contour lines, based on the modelled relationship between average fares and route distance, which show lines along which average fares are equivalent for different length routes.

The results imply some apparent systematic differences in average fares for some routes and groups of routes, but that for the broad majority of routes, including lower-volume regional routes, estimated differences between actual and modelled fares are within the range of variation of that for major routes.

Notable results include:

- Queensland Local Fare Scheme (LFS) routes tend to be among the lowest fare mark-up routes. This is exhibited by the 'convex hull' of the difference between actual and modelled fares for these routes (grey-shaded region, upper-right

panel, Figure 5) lying to the lower left of all other route groups and the prevalence of Queensland LFS routes in the lower left portion of this area.

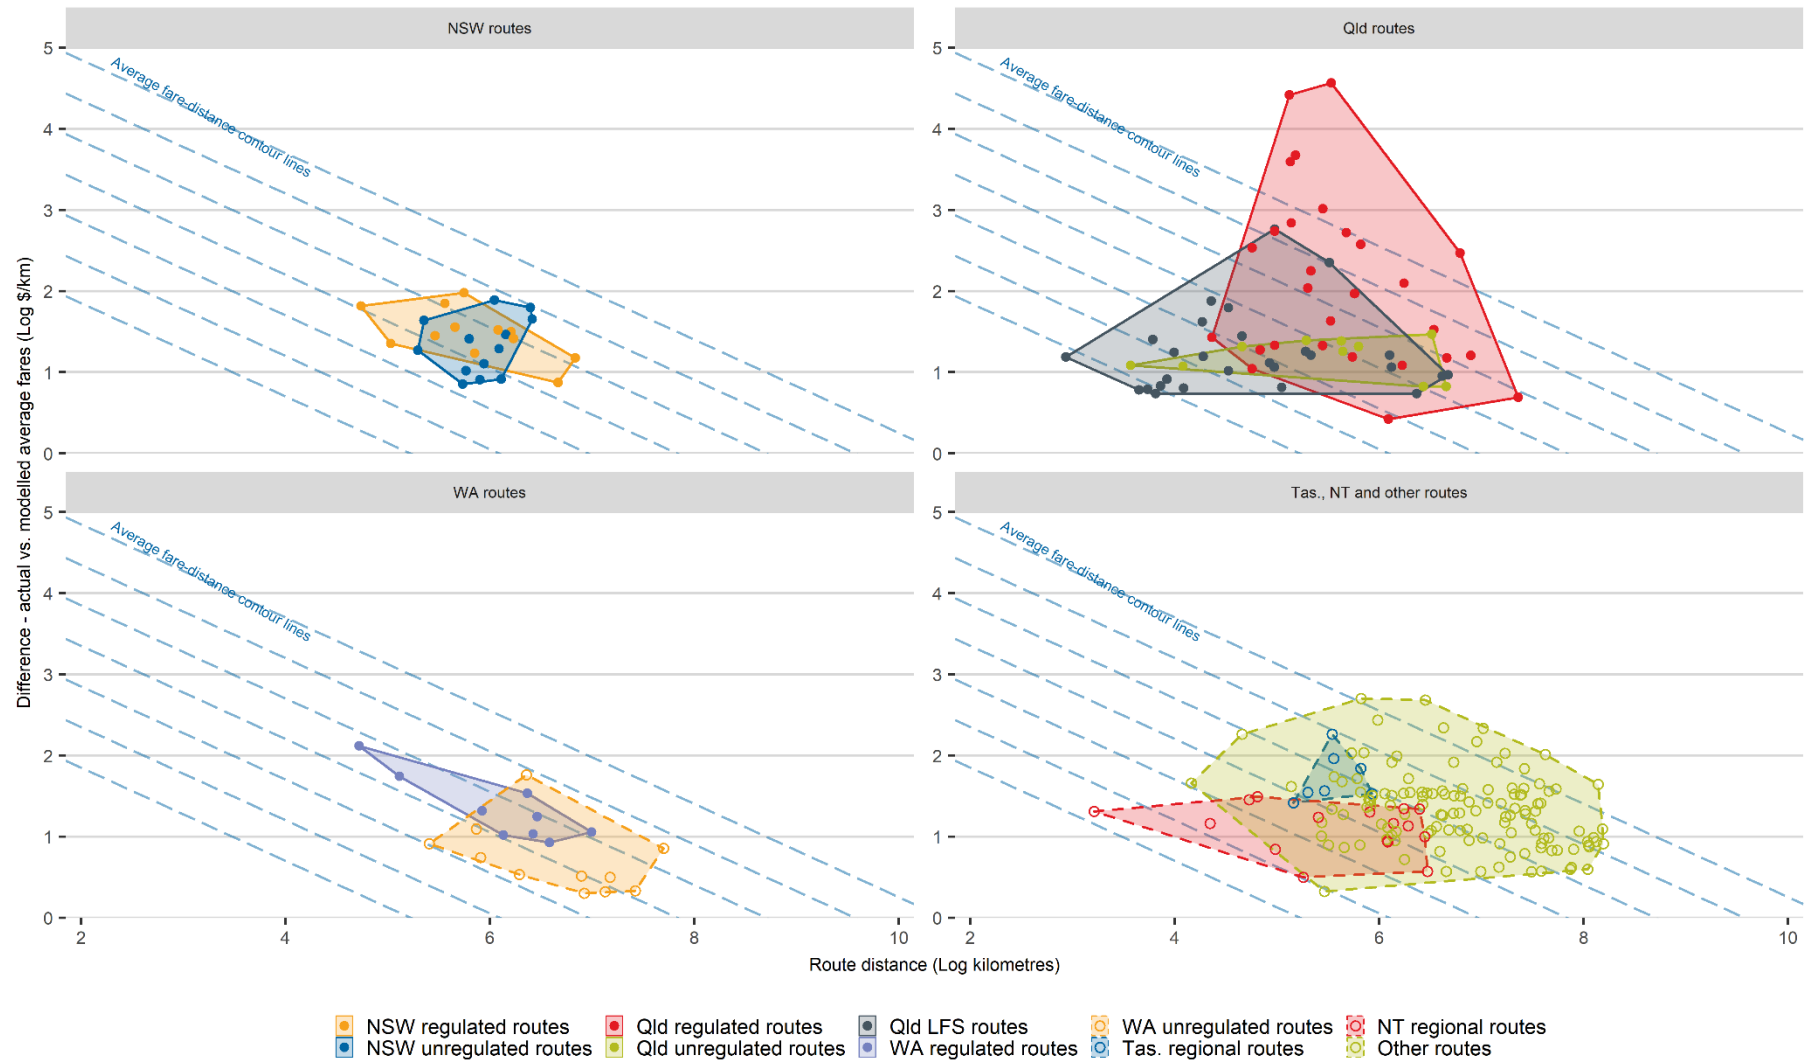
- In contrast, Queensland regulated routes include some of the highest apparent fare mark-up routes across Australia—exhibited by the convex hull of fares for these routes (red-shaded region, upper-right panel, Figure 5).
- Northern Territory (NT) regional routes also feature among the lower-end fare mark-up routes, exhibited by the convex hull of these routes (red dashed-line and shaded region, lower-right panel, Figure 5) also lying towards the lower-left quadrant.
- Fares on NSW regulated and unregulated routes are adjacently clustered, broadly parallel to the average fare–distance contours, suggesting average fares are broadly similar across regulated and unregulated routes in that jurisdiction (top-left panel, Figure 5).
- Fares on services to and from Bass Strait islands (King and Flinders Islands), grouped as Tasmanian regional routes in Figure 6 (bottom-right), appear to be around average fares on a per route kilometre basis.
- Like NSW, fares on WA regulated and unregulated routes are also adjacently clustered, also broadly parallel to the average fare–distance contours, suggesting average fares are broadly similar across regulated and unregulated routes in that jurisdiction (bottom-left panel, Figure 5).

The results also provide some similarities and contrasts with the results reported in Section 5.1. For example:

- The implied mark-up (i.e. divergence between actual and modelled fares) for long-distance routes between Perth to northern Western Australian airports are much less than that implied by the model results. Examples include: Perth–Port Hedland, Perth–Broome, Perth–Karratha, Perth–Paraburdoo, Perth–Kalgoorlie, and Perth–Kununurra.
- On the other hand, a number of the longer-distance routes, still appear among the higher mark-up routes. For example, Brisbane–Darwin, Brisbane–Port Hedland, Cairns–Perth, Perth–Sydney, Canberra–Perth, Brisbane–Perth, Darwin–Sydney, Darwin–Melbourne, Melbourne–Townsville appear among the higher fare mark-up routes.
- Routes to and from Avalon Airport also feature among higher mark-up routes, including Avalon–Sydney, Avalon–Adelaide, Avalon–Gold Coast.

Again, it is important to remember that these results are based on a one-month snapshot of fares, and hence may not reflect longer-term systemic differences in fares across routes. While these initial results may make intuitive sense—e.g. below-average average fares on subsidised routes, higher-than-average average fares on some regulated routes—a longer-term time series set of observations would be necessary to more conclusively estimate systematic differences in average fares across routes

Figure 5: Grouped difference between actual and modelled domestic return airfares, by route distance July 2018



Note The average fare–distance contours (dashed lines) represent lines along which average fares are 'equivalent' for different distance routes.

7 Implications and concluding remarks

The analysis presented in this paper suggests that regional airfares respond largely as expected in a competitive market. In particular, the analysis found:

- significant distance-based scale economies in the provision of aviation services, hence average fares generally decline with increasing route length
- strong market-based scale economies in aviation services, with average fares strongly declining with increasing market size (as measured by the number of passengers)
- competition (i.e. number of operators) has a significant impact on average fares, with fares declining significantly with increasing numbers of route operators.
- oil prices have a statistically significant, but relatively small impact on average airfares, which appears to reflect that fuel costs are typically around 20 per cent of total operating costs of major airlines and also airline fuel hedging practices.
- the number of flights has a positive impact on average fares, presumably reflecting the increased costs of adding flights, and some dilution of the scale economies from increasing capacity utilisation.

Moreover, the analysis also suggests there are statistically significant differences in average fare levels across routes, which also appear correlated with the degree of competition and also the availability of alternative transport options for travellers. BITRE identified three broad groups of routes:

- High-mark-up routes (i.e. above average fare routes) – which feature predominantly longer-distance trunk routes to/from Perth, Darwin and Alice Springs, and also several routes services remote mineral industry locations (e.g. Karratha, Port Hedland, Newman, Weipa). For these routes, there are no time-competitive alternatives to air transport and demand may be less responsive to price changes, enabling airlines to price accordingly.
- Mid-tier mark-up routes – which include most of the higher volume (trunk) domestic commercial routes and several smaller-distance routes.
- Low-mark-up routes (i.e. below average fare routes) – which predominantly comprise shorter distance routes, such as Melbourne-Devonport, Melbourne-Burnie, Coffs Harbour-Sydney, Hervey Bay-Sydney and Melbourne-Launceston. These routes are all routes for which car, or ferry in the case of the Tasmanian routes, is a more significant competitor.

Examination of actual and modelled fares across a broader selection of routes, including lower-volume regional routes, imply that while there are some apparent systematic differences in average fares for some routes or groups of routes—e.g. below-average fares on some subsidised routes and above-average fares on some regulated routes—for the broad majority of routes the estimated difference between actual and modelled fares are generally within the range of variation exhibited by major trunk routes. However, as these results are based on a one-month only snapshot of fares (July 2018), they may reflect unaccounted for one-off effect, and should not be treated as conclusive evidence of systematic differences in pricing across different routes.

Overall, the results of this analysis show that measures that:

- increase effective competition on routes
- increase competition from other transport modes
- increase passenger volumes through airports

are most likely to put downward pressure on regional airfares.

8 References

BITRE 2018, *Domestic Air Fare Indexes*, Unpublished data, BITRE, Canberra.

South Australian Department of Planning, Transport and Infrastructure (DPTI) 2017, *Annual Report 2016–17*, South Australian Government, Adelaide.
URL: dpti.sa.gov.au/annual_report.

Qantas Group 2017, *Annual Report 2016–17*, Qantas Group, Sydney.
URL: investor.qantas.com/investors/?page=annual-reports, Accessed: 10 August 2018.

Virgin Australia Airlines 2017, *Annual Report 2016–17*.
URL: www.virginaustralia.com/au/en/about-us/company-overview/investor-information/annual-reports/, Accessed: 10 August 2018.

World Bank 2018, *Historical Monthly Commodity Prices*,
URL: www.worldbank.org/en/research/commodity-markets, Accessed: 20 August 2018.

Appendix A – BITRE best discount fares by route

Figure A.1 Average best discount fares, by route



Figure A.1 Average best discount fares, by route (continued)



Figure A.1 Average best discount fares, by route (continued)



Figure A.1 Average best discount fares, by route (continued)



Figure A.1 Average best discount fares, by route (continued)




Figure A.1 Average best discount fares, by route (continued)



Source BITRE (2018).