

Guide to the taxi industry model

March 2013

This documentation is for the version of the model used for the 2012 licence review. This version of the model is no longer current.

As part of our review of the number of new taxi licences to be released in Sydney from July 2013, we asked economic consultants the Centre for International Economics (the CIE) to develop a model of the taxi industry in Sydney. The model shows what happens to taxi trips taken, taxi occupancy rates, waiting times for taxis and annual taxi licence costs if taxi numbers, fares or the costs of providing taxi services change. We will also be using the model to inform our review of taxi fares to apply from July 2013.

The taxi industry model is available on our website as an Excel 2010 file.¹

This guide is structured as follows:

- ▼ Section 1 explains the background to the model
- ▼ Section 2 describes in general terms how the model works
- ▼ Section 3 explains the sources of data and assumptions for the model
- ▼ Section 4 steps users through using the Excel file to model scenarios (including how to install and run the model in both Excel 2010 and Microsoft 2007)
- ▼ Appendices A and B contain more detailed technical descriptions of the model's workings
- ▼ Appendix C sets out a full list of source references for the data and assumptions in the model.

This guide does not discuss any of the specific scenarios we modelled for our taxi reviews. Please read our review reports² for a discussion of specific scenarios.

¹ The model can be found on the Review of Sydney Taxi Licences page at http://www.ipart.nsw.gov.au/files/c32a0c3b-1d34-4dbe-a024-a1ba00c2be71/Taxi_Industry_Model_-_May_2013.xlsm

² Review of Sydney Taxi Licences - draft report (December 2012) and final report (February 2013), and Review of Fares from July 2013 - issues paper (February 2013) are all available on our website.

1 Background to the taxi industry model

All models of markets are estimations of how things actually work. Good models are those that are able to take into account the key features of the market and measure how changes to certain variables (such as the number of taxi licences) affect other variables (such as the waiting time for taxis).

The CIE's model builds on work done by other economists to understand the relationship between number of taxis, the price of taxis, waiting time for taxis and demand for taxis.

Good models also require good data about the features of a market (such as the costs of providing taxi services). Where data is not available for elements of a model, modellers must make assumptions. They make assumptions that are reasonable based on the information that is available.

The data for the taxi industry model comes from data collected from the taxi industry by Transport for NSW, surveys, a sample of taxi meter data and a taxi network. More detailed information on the data sources is provided in section 3 and Appendix C of this guide, as well as more detailed information on the assumptions used in the taxi industry model.

The model uses the costs of providing taxi services, the number of taxi licences, fares, and the customer responsiveness to waiting times and fares to predict:

- ▼ outcomes for passengers – waiting times
- ▼ taxi use and productivity – the total number of trips and taxi occupancy
- ▼ impacts on licence lease costs – costs for operators and income for licence owners.

These are the key outputs of the model.

The taxi industry model is a 'long-run equilibrium' model, which means it shows what the taxi market will look like once it has fully adjusted to a change (such as increasing the number of licences). It does not show what the impacts will be as the market works through the changes – eg, how quickly passengers will react to lower waiting times, whether occupancy will be lower in the short term – or how long the market will take to reach equilibrium.

The model is also a 'real' model, which means it shows changes to values after inflation has been taken into account. That is, it does not matter what the rate of inflation is over the period that the market takes to adjust to the changes: the model shows the changes that occur in addition to the changes due to inflation.

2 How does the taxi industry model work?

The model has a worksheet ('Base model') that characterises the Sydney taxi industry for the year 2012/13, and a worksheet ('Sim_template') where inputs can be changed and the model run to show predicted outputs for that set of inputs. The relationships between the inputs and outputs are set out in equations in the model.

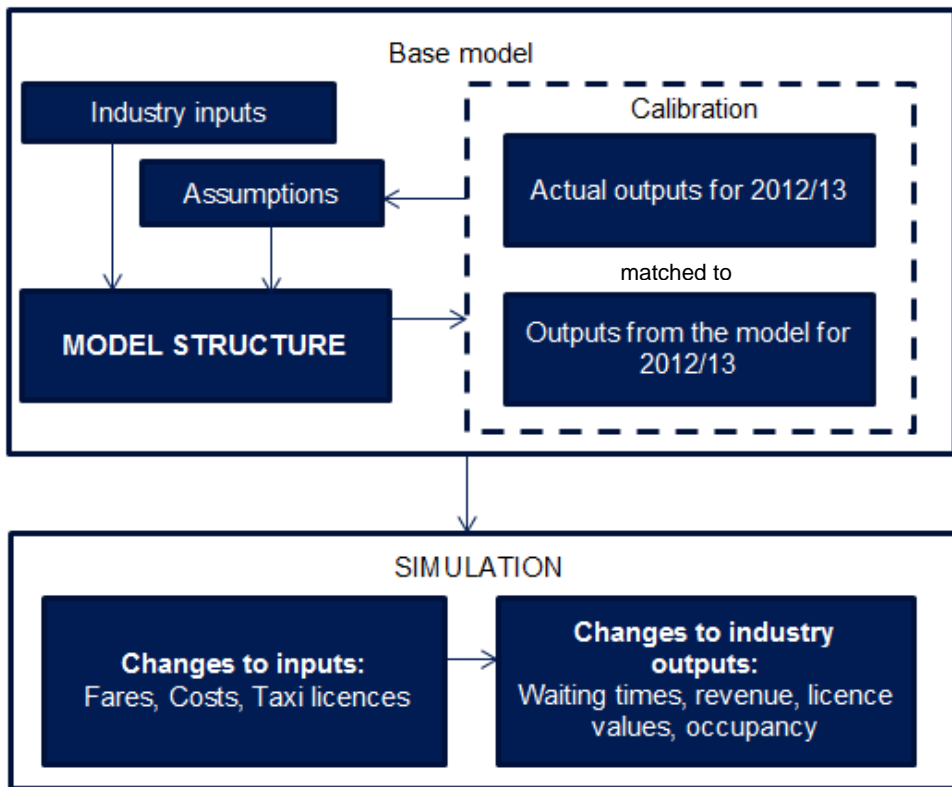
2.1 The base model

The model is set up on the basis of 14 shifts – 2 shifts a day across a week – because the level of demand is different for each shift. This means that for some shifts, the number of taxis on the road is lower than the number of current licences and hence increasing the number of licences will have a smaller impact, or potentially even no impact, on taxis on the road during these shifts.

The base model has been calibrated so that the inputs and outputs most closely match the data that we have for the Sydney taxi industry in 2012/13. However, if users have a different view on the inputs and assumptions (for example, the value of time for a taxi passenger) that were used to calibrate the base model, they can change them on the 'Inputs and assumptions' worksheet. The base model then needs to be run again (and possibly re-calibrated, if the outputs no longer match the 2012/13 data). See section 4 of this guide for more detail on running the base model.

Figure 2.1 below shows the relationship between inputs, assumptions, calibration and simulation.

Figure 2.1 Taxi model



2.2 The equations

The taxi industry model incorporates a demand equation, a waiting time equation and a taxi entry condition, which are discussed briefly below and explained in more detail in Appendix A.

2.2.1 A demand equation

A demand equation links passenger demand for taxi services to the price of the services and waiting time. If prices are higher, demand is lower; if waiting times are higher, demand is lower. The degree to which demand responds to changes in price or waiting time is called the elasticity of demand. Elasticity of demand can be estimated by looking at data (if available) to see how demand actually changed when there was a change in price or waiting time, or by surveying people about how they would change their behaviour in response to specified changes.

The demand equation in our taxi industry model assumes that for a 1% increase in taxi fares, demand would fall by 0.8% (and if there is a 1% decrease in taxi fares, demand would increase by 0.8%). This value of price elasticity is an assumption based on a range of sources from other taxi industry models. The model also puts a 'price' on waiting time by using an assumption of \$30 per hour of the value of that waiting time to the person waiting, and then uses the same price elasticity of demand as for fares. This means that with a \$30 per hour value of time, if waiting time reduced by one minute, the demand response would be the equivalent of the response to a 50 cent saving on the fare.

2.2.2 A waiting time equation

Waiting time for taxis is clearly not constant across shifts (or within a shift). Even when there are a lot of taxis on the road, if there is very high demand, waiting time can be high or a taxi might not come at all, such as on New Year's Eve. Taxis can also be difficult to obtain on Friday and Saturday nights and peak times within a shift.

A waiting time equation links the average waiting time for passengers during a shift to both the number of taxis on the road and the share taken up by demand (so waiting times are generally shorter if there are more taxis on the road, but if there are more people trying to catch taxis at the same time, waiting times are longer for a given number of taxis on the road).

The waiting time equation in the model is based on established taxi waiting time equations from models of the 'cruising' (ie, not booked) taxi market because only around 20% of taxi journeys in Sydney are booked.

The waiting time equation implies that if demand were twice as high for a shift and there were twice as many taxis operating, then waiting times would be lower. Essentially, this is because the (geographical) density of taxis increases as the number of taxis increase, so would-be passengers are on average closer to an empty taxi and waiting times are shorter.

2.2.3 A taxi entry equation

A taxi entry equation links the decision of a taxi operator to put a taxi on the road for a shift to the revenues available and the costs incurred. An operator should put a taxi on the road as long as expected revenues exceed expected costs (not including the plate lease cost). Any return above costs from a shift is 'profit', which when added up over a year is equal to the plate lease cost.

The costs per shift consist of the costs that are only incurred while the taxi is on the road ('variable costs') and a share of the costs that are incurred whether the taxi is on the road or not ('fixed costs'). The 2 types of costs and the way that the share of fixed costs per shift is allocated are discussed in more detail in the sections below.

Variable costs

Costs that are only incurred while the taxi is on the road include driver labour costs, fuel costs, cleaning costs, operator cost of organising the driver, and potentially some insurance and maintenance costs (where these differ depending on kilometres driven). Costs also include GST payable.

Variable costs are relatively simple to observe and calculate on a per shift basis. The model allows for the possibility that higher earnings are required to get more taxi drivers to take on a shift as the number of taxis increases and the demand for taxi drivers grows.

Fixed costs

Costs that are incurred whether or not the taxi is on the road include network fees, vehicle costs and some part of insurance and maintenance costs. The decision to put a taxi on the road for a particular shift involves allocating these fixed costs in some way.

Fixed costs can be allocated in many different ways. In some industries, such as retail, fixed costs are allocated using margins. That is, every item sold has a mark-up over its wholesale cost that contributes to the fixed costs of selling it (such as the shop rent, electricity etc).

In other industries, such as gas-fired electricity generation, fixed costs are recovered by peak pricing. That is, when demand for electricity is low, the price only covers the variable costs of producing it. When demand for electricity is highest, the price charged by generators includes a portion that goes towards the fixed costs of producing it (such as the cost of building the power station).

The way that operators allocate fixed costs in the taxi industry is critical for understanding the impact of changes in licence plate numbers and therefore on lease costs. In our model, we assume that operators recover a similar share of fixed costs in all shifts. We allow for a moderate increase in costs allocated to peakier shifts, which allows for recovering more fixed costs for these shifts, as well as higher returns to encourage more drivers into these shifts.

A limit on number of taxis on the road

The model also includes a limit on the number of taxis operating in each shift. Obviously, the number of taxis operating cannot exceed the number of taxi licences. However, the actual number of taxis operating is lower than the number of licences because some vehicles are being repaired or serviced at any one time. Some operators may also prefer to drive a cab one-out and not put it on the road for every shift. Thus, even at peak times not all cabs are on the road.

For our model, the CIE used data from networks on number of cars on the road compared to the number of licences to make an assumption about the theoretical maximum number of cars on the road in relation to number of licences.

Plate lease values

The limit on the number of taxis on the roads means that there can be a gap between the revenues that taxis can earn and the costs of providing this service. In a competitive market this gap would lead to additional entry of taxis, which would lower revenue until revenues and costs were equal. However, in the taxi market, additional taxis are unable to enter because of the restriction on taxi licences.

The gap between revenues and costs is a rent that accrues to the owners of taxi licence plates. The annual rent is the sum of rents across all the shifts allocated on the basis of the average number of shifts a taxi undertakes. Because of this annual rent, taxi licences have considerable value for sale.

2.3 Modelling results

Waiting time, demand for taxis, number of taxis and the price of taxis are all inter-related. For example, if more taxis are available to be on the road, waiting times will come down, and demand for taxis will increase, which will increase the amount of revenue a cab can earn in a shift, which will induce more operators to put their cabs on the road, which will shorten waiting times again, and so on.

The situation where all these factors (demand, number of taxis and waiting time) are in balance and not leading to further change is called equilibrium. The following sections explain the 'Base model' equilibrium and the 'Sim_template' equilibrium. Appendix B provides a technical explanation of how the model calculates the equilibrium point.

2.3.1 The base model equilibrium

The base model is calibrated to match the Sydney taxi industry in 2012/13 as closely as possible. When the user runs the model on the 'Base model' sheet, the model calculates the equilibrium solution based on the input data and assumptions on the 'Inputs and assumptions' page and the 'Base model' page.

To calibrate the model, the CIE looked at the outputs (in terms of average occupancy, trip kilometres per year, average revenue per taxi per year, and plate lease costs) and checked them against any data that was available. If the outputs were not a good fit, the CIE changed the input assumptions (such as a taxi's average speed) and ran the model again until the outputs were the best fit. (See the 'Model Validation' table between rows 59 and 69 of the 'Base model' worksheet for details of the fit between base model outputs and data.)

2.3.2 The 'Sim_template' equilibrium and outputs

The main purpose of the taxi model is to understand how changes in areas that we make recommendations on to the NSW Government will impact on the taxi industry. This includes changes in fares and changes in the number of taxi licences available. The model also allows for consideration of the impact of other changes that may occur that are outside of areas where we make recommendations. This includes changes in the costs of providing taxi services and changes in demand due to population growth and rising incomes. We allow for these changes so that we can understand the likely impact of our recommendations after allowing for these other changes.

The 'Sim_template' worksheet allows users to consider changes to the industry in terms of:

- ▼ changes to the level and structure of fares
- ▼ changes in the number of peak and standard licences
- ▼ changes in the costs of operating taxi services, and
- ▼ changes in demand.

When the user changes the inputs (the green cells on the 'Sim_template' worksheet) and runs the model, the model calculates a new equilibrium using the changed inputs and the equations set out above. The 'Detailed model outputs' table at row 43 on the 'Sim_template' worksheet shows the estimated outputs on a per shift basis.

The 'Sim_template' worksheet also compares the outputs from the simulated scenario to the outputs from the base model in the 'Comparison of simulation to baseline' table at row 59.

Note that a user cannot change inputs between the base model and the simulation other than those specified in the green cells. This is because the other inputs aim to understand the underlying structure of the industry, which should not vary through time. The user also cannot change industry outputs, such as occupancy rates, as this is a result of the model rather than an input into the model.

3 Data and assumptions

The model is built with a combination of data (where we have a reliable source of information about a particular input) and assumptions (where we do not have data).

The sections below set out the sources for and explanations of the data and key assumptions in the 'Inputs and assumptions' and the 'Base model' worksheets. (A detailed list of the data and assumption sources is at Appendix C.)

In the Excel model, input data cells are coloured pink, while input assumption cells are coloured blue.

3.1 Sources for the data

In October 2011, the CIE, on behalf of IPART, conducted a survey of taxi drivers and operators to assess the costs of providing taxi services in NSW. The final survey report is available on IPART's website.³ Data from this survey included in the taxi model include operator costs (operator administration, insurance, vehicle maintenance, vehicle lease, plate lease, network fees), driver costs (fuel, labour, cleaning), revenue per taxi, hours per shift, total kilometres travelled per taxi per year and number of shifts per taxi per year.

As part of the CIE's review of the costs of providing taxi services, they obtained data from Combined Communications Network (CCN) on the number of taxis logged on to the network throughout the year.

The Australian Taxi Drivers' Association also provided meter data to the CIE from 14 taxis operating from Alexandria, covering more than 5 million kilometres of operation. This data is used in the model to estimate the share of revenue charged at the distance rate, the share of revenue charged at the waiting time rate, the share of paid kilometres and the share of unpaid kilometres. We also cross-checked this data with data from the Victorian taxi industry inquiry.

³ Ref to CIE's final survey report.

As the principal regulator of the taxi industry, Transport for NSW collects data from taxi networks, operators, drivers and licence holders. Transport for NSW provides some information to IPART for our reviews about licence numbers, ownership patterns, and network performance standards. No individuals are identified in any data provided by Transport for NSW. Transport for NSW provided data for the model on numbers and types of taxi licences as at 1 January 2013.

The Bureau of Transport Statistics (BTS), a division of Transport for NSW, conducts an ongoing household travel survey in which some data is collected about taxi travel. We used BTS data on the average taxi trip length in the model.

In November 2012 we commissioned Taverner Research to conduct an on-line survey of 2000 Sydney residents about their taxi use (or non-use). We used waiting time data from this survey in the model.

3.2 Key assumptions in the model

Price elasticity of demand (Row 5 of the inputs and assumptions sheet): is an assumption based on a review of the literature about taxi markets:

- ▼ The Victorian Taxi Industry Inquiry's draft and final reports suggested a price elasticity of demand for taxi services of around -1 for Melbourne, based on survey work undertaken in Melbourne by The Hensher Group.⁴
- ▼ Booz Allen Hamilton, in a report for IPART in 2003, noted that the majority of international studies reported a demand elasticity of -0.2 to -1.0.⁵

The CIE adopted a value of -0.8 for the taxi industry model. It is above the midpoint of the Booz Allen Hamilton review, but closer to the value used in Melbourne.

Value of time (C 48 in the inputs and assumptions sheet): An assumption about the value of a passenger's time spent waiting for a taxi, used to convert waiting times to dollar values. Putting a value on time is standard practice for cost/benefit analysis, and the dollar value is usually lower than the \$30 per hour proposed by the CIE; on the other hand, the Victorian Taxi Industry Inquiry used a value of \$60 per hour.

Constraint on number of cabs (C49 to C51 in the inputs and assumptions sheet): These are assumptions about the effective maximum number of cabs that can be on the road (as opposed to the notional maximum, which is the number of licences).

⁴ Victorian Taxi Industry Inquiry, Final report, October 2012, p 260.

⁵ Booz Allen Hamilton 2003, *Appraisal of taxi fare structure issues*, p 10.

The assumptions are based on data from networks about the share of cabs on the road per shift. The effective maximum is determined not only by the percentage of cars which are off the road for maintenance or repair, but by the behaviour of operators, some of whom prefer to drive their cab one-out and leave it off the road some of the time rather than engage a bailee driver, and it varies by shift depending on demand. This reflects that one-out operators are most likely to choose higher demand shifts to operate their taxi.

4 How to use the model

4.1 Instructions to install the model (for Excel 2010 and Microsoft 2007)

The taxi model uses an Excel function called Solver, through a Microsoft Excel macro. To run the model, Solver should be installed and macros enabled. To do this, use the following steps:

4.1.1 For Excel 2010

Install Solver

- ▼ Go to the File menu and scroll down to Options.
- ▼ Choose Add-ins
- ▼ Choose Manage: Excel Add-ins and select Go...
- ▼ Check if Solver Add-in is ticked. If not tick it. Select OK.

Enable macros

- ▼ Go to the File menu and scroll down to Options
- ▼ Choose Trust Center
- ▼ Select Trust Center Settings
- ▼ Select Macro Settings
- ▼ Choose Enable all macros. Select OK.

4.1.2 For Microsoft 2007:

- ▼ Select the Office button  and choose Excel options.
- ▼ Then proceed from 'Choose Add-ins' step onwards above.

4.2 How to run a simulation using the model

1. Go to the sim_template worksheet.
2. The cells coloured green are the input cells which can be changed by the user.

You can:

- Change individual fare components (cells B1 to B5). These are nominal changes – that is, not adjusted for inflation – so just put in whatever face value you expect fares to have.
 - Add unrestricted licences and/or Peak Availability Licences (cells D1 and D2); these are the additional licences only (ie, do not include licences that are released as replacement for existing licences that are no longer operated).
 - Input an estimate of demand growth from external sources (cell D3).
 - Input an estimate of CPI (cell D4) – this adjusts the fare revenue and the value of time to allow for the rate of inflation.
 - Input the nominal change to the value of time. If you expect the value of time to grow at the same rate as inflation, you will need to put the same figure in here as you did for CPI.
 - Input a change to the real cost (that is, the change in costs after allowing for inflation) of providing taxi services (for example, if driver labour costs go up as a result of an Industrial Relations Commission decision, real costs of providing taxi services could increase. If costs just increase at the rate of inflation, the real cost has not changed and this cell can be left as 0).
3. Click on the 'run model' button. The model will take a few seconds to run.
 4. Look at the detailed model outputs (rows 43 to 57) or the comparison of the outcomes from the simulation you just ran to the base model outcomes (rows 59 to 71).
 5. It is possible to change the assumptions and the inputs in the assumptions and input sheets (we did this when testing the sensitivity of the model to some of our assumptions), but if you do change these, you need to run the base model by clicking on the 'run model' button on the 'base model' sheet BEFORE running the model on the 'sim_template' sheet. You may also have to recalibrate the model if the model validation section no longer shows a close match between the data and the model outputs, by adjusting some of the assumptions on the inputs and assumptions sheet and re-running the base model until the outputs are a good fit.

A Detailed discussion of the equations in the model

The taxi industry model incorporates a demand equation, a waiting time equation and a taxi entry condition, which are explained in more detail below.

A.1.1 A demand equation

A demand equation links passenger demand for taxi services to the price of the services and waiting time. If prices are higher, demand is lower; if waiting times are higher, demand is lower. The degree to which demand responds to changes in price or waiting time is called the elasticity of demand. Demand is measured in trip kilometres.

The demand equation is

$$Q = C_Q - \beta_P \cdot P - \beta_W \cdot W$$

Where

- ▼ Q is the quantity of taxi services demanded (in trip kilometres)
- ▼ P is price per passenger kilometre. This is a combination of the flag fall, distance rate and waiting time. The CIE used meter data to estimate an average price per kilometre. Note that the price can differ across shifts if prices are different at different times of the day, as is currently the case for Sydney taxi fares (due to the night time surcharge on the distance rate).
- ▼ W is waiting time for the taxi user (in minutes).
- ▼ C_Q is a constant set to the estimated trip kilometres for each shift in 2012/13 (based on data from the CIE's 2011 cost survey of operators and drivers).
- ▼ β_P is the elasticity of demand to price and β_W is the elasticity of demand to waiting time. The value of the elasticity of demand to price parameter (β_P) is an assumption made by the CIE after a review of the literature. Rather than trying to estimate a separate elasticity of demand to waiting time, the CIE converted waiting time to a dollar value by making an assumption about passengers' value of time, then applied the same elasticity of demand.

A.1.2 A waiting time equation

A waiting time equation links the average waiting time for passengers during a shift to the number of taxis on the road and the share taken up by demand (so if there are more taxis on the road, waiting times are shorter).

Waiting time for taxis is clearly not constant across shifts (or within a shift). When there is very high demand, waiting time can increase substantially or a taxi might not come at all, such as on New Year's Eve. Taxis can also be difficult to obtain on Friday and Saturday nights and peak times within a shift.

Two models of taxi waiting time have been developed in the academic literature. The first is a model of a cruising (or 'rank and hail') taxi market and the second of a dispatch (booked) market.

The waiting time formulae for these models suggest that areas of higher demand density, such as Sydney, are likely to be best represented by the cruising taxi market model as this type of taxi behaviour leads to lower waiting times than relying on dispatch. This aligns well with evidence across Australia, with reliance on bookings being far higher in smaller towns. In Sydney, only around 20% of taxi trips are booked through a dispatch system.

Therefore the model assumes a cruising market, which means that taxis are caught by rank and hail, rather than being booked.

The cruising market waiting time equation is as follows.

$$W = \frac{C_W}{V}$$

Where W is waiting time, C_W is a waiting time constant and V is the number of vacant taxis.

The number of vacant taxis can be rewritten as

$$V = T - Q \cdot C_V$$

Where T is the total number of taxis working the shift, Q is total demand and C_V is a constant that relates kilometres of demand to hours of demand as a share of hours available for the shift.

The waiting time equation implies that if demand were twice as high for a shift and there were twice as many taxis operating, then waiting times would be lower. Essentially, this is because the (geographical) density of taxis increases as the number of taxis increase, so would-be passengers are on average closer to an empty taxi and waiting times are shorter.

A.1.3 A taxi entry condition

A taxi entry condition links the decision of a taxi operator to put a taxi on the road for a shift to the revenues available. The return above costs from a shift, which should be equal to the plate lease cost, is as follows.

$$\pi = P \cdot q - Cost$$

Where P is price per kilometre, q is the number of trip kilometres (quantity of service) travelled by a particular taxi and cost is all costs *associated* with the shift. Cost has a variable component and a fixed component and can be rewritten as:

$$Cost = c \cdot q + F$$

Where c is the cost per kilometre and F is fixed costs.

We assume that each taxi receives an equal share of total demand.

$$q = Q/T$$

Where T is the number of taxis.

We also allow for the possibility that there is an upward-sloping labour supply curve: that is, higher earnings are required to get more taxi drivers to take on a shift as the number of taxis increases:

$$F = F_0 + F_1 \cdot T$$

Where F_0 is a fixed cost component and F_1 is a constant relating the number of taxis to costs.

The return above the costs of the shift can then be written as follows.

$$\pi = \frac{(P - c) \cdot Q}{T} - F_0 - F_1 \cdot T$$

Entry occurs to the point at which revenues and costs are equal ($\pi = 0$). Hence T is determined by solution of the following quadratic.⁶

$$0 = F_1 \cdot T^2 + F_0 \cdot T - (P - c) \cdot Q$$

This has one positive solution at

$$T = \frac{-F_0 + [(F_0)^2 + 4 \cdot F_1 \cdot (P - c) \cdot Q]^{1/2}}{2 \cdot F_1}$$

There is also a limit to the number of possible taxis operating (\bar{T}). Obviously, the number of taxis operating cannot exceed the number of taxi licences. However, the limit on the number of taxis operating (\bar{T}) is actually lower than the number of licences because some vehicles are being repaired or serviced at any one time. Some operators may also prefer to drive a cab one-out and not put it on the road for every shift. Thus, even at peak times not all cabs can be on the road for any shift. The CIE used data from networks on number of cars on the road compared

⁶ Note that if the fixed costs are not related to T then F_1 is equal to zero and the solution is straightforward.

to the number of licences to make an assumption about the theoretical maximum number of cars on the road in relation to number of licences.

If T is greater than \bar{T} then the number of taxis actually operating is \bar{T} . In this case π , which is the annual plate lease cost, is greater than zero.

Costs associated with a shift include labour costs to make it worthwhile for the driver to work the shift, fuel costs, cleaning costs, operator administration cost of organising the driver and potentially insurance costs (where these differ depending on kilometres driven). A break-down of costs for an example mid-week night shift is shown in Table A.1 below.

Table A.1 Example of costs and revenue for a Wednesday night shift (\$2012/13, excluding GST)

	\$/vehicle/shift (excluding GST)
Driver costs	
Labour costs	104
Cleaning costs	11
Fuel	28
Total	143
Operator variable costs	
Operator maintenance – variable component	7
Operator insurance – variable component	13
Total	20
Total shift variable costs	
163	
Operator fixed costs recovered from shift (including network fees, operator administration, vehicle lease and the fixed component of maintenance and insurance)	
66	
Total costs recovered from shift	
229	
Revenue (excluding GST)	
281	
Plate lease residual	
52	

Note: This allocates 50% of insurance and maintenance costs as variable costs.

Source: The CIE 2012, drawn from the CIE 2011 taxi cost survey of operators and drivers.

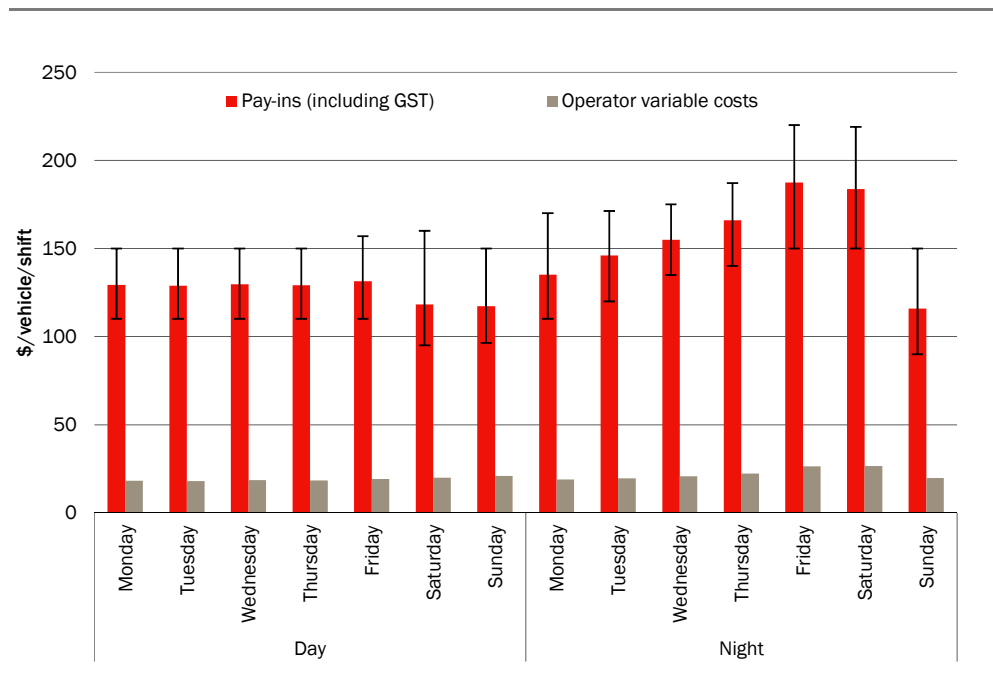
For all shifts, it appears that taxis earn a substantial margin on their variable (or shift-specific) costs.

The fixed costs of operating a taxi include network fees, vehicle costs and some part of insurance and maintenance costs. (Plate lease costs are the residual between revenue and costs of providing services.)

Fixed costs can be allocated in many different ways. The way that operators allocate fixed costs in taxis is critical for understanding the impact of changes in licence plate numbers and lease costs.

Figure A.1 shows the pay-ins for each shift. The difference between the pay-ins and operator variable costs is an allocation by an operator of fixed costs and plate leases to each shift. It shows that operators recover some share of fixed costs/plate leases on all shifts and a higher share of fixed costs/plate leases on peak shifts. The increase in pay-ins for higher demand shifts is likely to mostly reflect higher plate lease rents to shifts where the restriction on taxi licence numbers is more important.

Figure A.1 Operator pay-ins and variable costs




Note: Error bars are the 10th percentile and the 90th percentile.

Data source: The CIE 2012. Pay-ins for day shifts are very similar across the weekdays, while the pay-ins for night shifts vary far more. This may reflect the permanent employment structure of day drivers versus a more casual basis for night drivers.

In interpreting these pay-in figures, there are a number of factors that are relevant.

- ▼ The maximum pay-ins that can be charged by an operator is set by the NSW Industrial Relations Tribunal. At the time that information was collected, this ranged from \$171.92 for all day shifts to \$261.84 for Friday and Saturday night shifts (including GST). This is higher than the actual pay-ins observed. In reality there is less variation between shifts pay-ins because bailee drivers are often on the basis of rosters of shifts rather than single shifts.
- ▼ Operators may be unwilling to reduce pay-ins for specific shifts because of the precedent that this sets. Instead, operators may leave their cars off the road rather than reducing pay-ins.

This documentation is for the version of the model used for the 2012 licence review. This version of the model is no longer current.



- ▼ The pay-in numbers provided by drivers are for one week and there is some variations in pay-ins. The 10th percentile and 90th percentile from the survey of taxi operators and drivers is shown in Figure A.1.

B How the model calculates equilibrium – technical explanation

The equilibrium of the taxi industry model is the number of taxis (T), the waiting time (W) and the demand, expressed in number of trip kilometres (Q), which solve the following equations. (2A is with an upward-sloping labour supply curve – that is, drivers have to be paid more to work when there are more taxis available – and 2B with a flat labour supply curve – that is, drivers do not require more to work when there are more taxis.)

$$Q = C_Q - \beta_P \cdot P - \beta_W \cdot W \quad (1)$$

$$T = \frac{-F_0 + [(F_0)^2 + 4 \cdot F_1 \cdot (P - c) \cdot Q]^{1/2}}{2 \cdot F_1} \quad (2A) \text{ or}$$

$$T = \frac{(P - c) \cdot Q}{F} \quad (2B)$$


$$W = \frac{C_W}{T - Q \cdot C_V} \quad (3)$$

Combining these equations gives a quadratic in Q for the flat labour supply curve, whose solutions are the equilibria. For the upward-sloping labour supply curve the solution has to be found numerically. For all models there are 2 possible equilibria – one with very high waiting times and lower demand and one with shorter waiting times and higher demand. Only the second equilibrium is stable, in the sense that an additional taxi entering would earn below its costs.⁷

The charts below show the equilibrium in quantity and waiting time. This involves substitution of T into the waiting time equation. In Figure B.1 we show the free entry equilibrium. In this case, higher demand induces more cabs into service and waiting time is lower in high demand shifts. (Waiting time would still be lower within high demand periods of a shift.)

⁷ A taxi entering in the high waiting time equilibrium would earn revenue above its costs and hence this equilibrium is not stable.

Figure B.1 Free entry equilibrium




Note: Slopes are not for the final model parameters.

Data source: The CIE.

In Figure B.2 we show the comparable waiting time equation for a fixed number of cabs. As would be expected, waiting time worsens as quantity demanded increases. The Sydney market can be seen as a combination of a free entry market up to a point until the restriction on plate leases becomes binding. This is shown in Figure B.3. When demand is low, the equilibrium is based on a free entry model of taxis. When demand is high, the equilibrium is based on a fixed number of taxis.


Figure B.2 Fixed number of taxis



Note: Slopes are not for the final model parameters.

Data source: The CIE.

Figure B.3 Free entry and capped taxi numbers

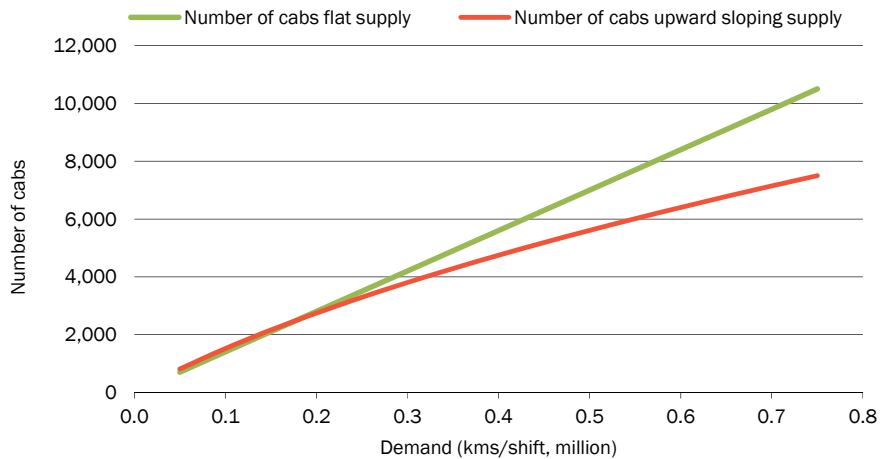


Note: Slopes are not for the final model parameters.

Data source: The CIE.

Figure B.3 is based on an upward-sloping labour supply curve for taxi drivers. If this was flat then the number of taxis responds more quickly to additional demand (Figure B.4).

Figure B.4 Upward-sloping versus flat labour supply curve



Note: Slopes are not for the final model parameters.

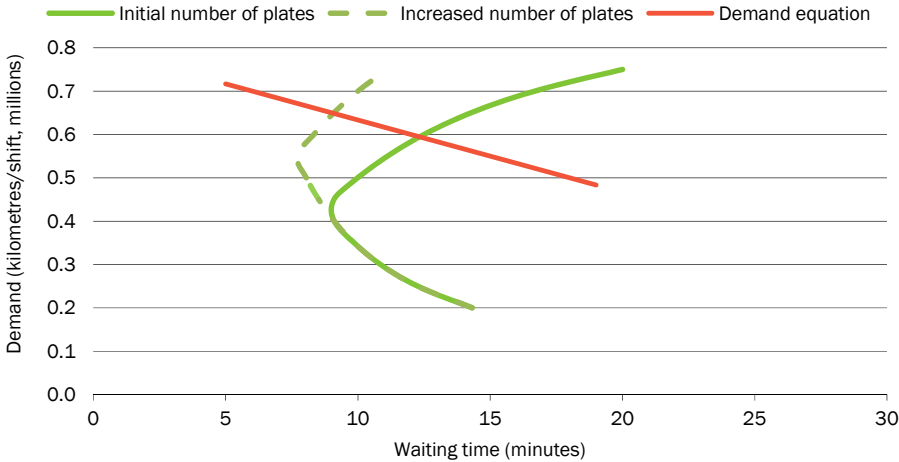
Data source: The CIE.

An increase in taxi licence numbers changes the part of the waiting time function where the number of licences is currently binding (Figure B.5). This means that it leads to higher demand and lower waiting times in shifts where the number of taxi licences is currently constrained.

There may also be flow-through changes to low demand shifts, given the way that operators allocate their fixed costs across shifts. However, these would be expected to be minor and there is no logical basis for expecting more taxis to be on the road when the number of taxis is not currently constrained by licence numbers on those shifts.

Even though the number of taxis on the road in low demand periods may be unchanged, the actual taxis doing this work probably will change. For example, it may be the case that there would be a shrinking of the number of shifts undertaken by most taxis rather than taxis using new plates only working at peak demand shifts.

Figure B.5 Increasing the cap on taxi numbers



C Full list of the sources of inputs and assumptions

C.1 Inputs and assumptions worksheet

Price elasticity of demand (Row 5): is an assumption based on a review of the literature about taxi markets. The Victorian Taxi Inquiry's draft and final reports suggested a price elasticity of demand for taxi services of around -1 for Melbourne. Booz Allen Hamilton, in a report for IPART in 2003, noted that the majority of international studies reported a demand elasticity of -0.2 to -1.0.⁸ The CIE adopted a value of -0.8. It is above the midpoint of the Booz Allen Hamilton review, but closer to the value used in Melbourne.

Does the weekend labour cost premium apply? (Row 6): reflects data from the CIE cost survey that showed that taxi drivers receive higher earnings per hour for shifts from Friday night to Sunday night (whether the shifts are high or low demand). The premium is approximately \$2 per hour (see cell H 30).

Revenue per taxi (Row 7): data from the CIE survey conducted in October 2011 (including GST).

Share of standard taxis on the road (Row 8): Data provided by Combined Communications Network for IPART's 2012 fare review.

Share of trips at standard distance rate (Row 9): An assumption about the share of trips in a shift for which the night time surcharge does not apply. Required because the hours at which the surcharge applies (10pm to 6am) do not match the shift hours (3am to 3 pm for day and 3pm to 3 am for night). Used to calculate the average price per trip kilometre across a shift.

Baseline waiting time (Row 10): Data from the Taverner survey of Sydney residents about average amount of time people have to wait to catch a taxi.

Fares (C 13 to C 17): Data about fare components for 2012/13.

Taxi fare increase July 2012 (C 18): data about the urban Taxi Cost Index that IPART calculated for the taxi fare review 2012, estimating the change in taxi costs for the year to April 2012. Used to index the costs measured by the CIE in October 2011 to 2012/13 levels for the base model.

Operator costs (H 22 to H 26) and driver costs (H 29 to H 32): Data from the CIE taxi cost survey October 2011.

⁸ Booz Allen Hamilton 2003, *Appraisal of taxi fare structure issues*, p 10.

Operator variable costs (C 24): Calculated per kilometre from the data from the CIE taxi cost survey October 2011, assuming 50% of maintenance and 50% of insurance costs are variable. (That is, depend on the number of kilometres driven.)

Operator fixed costs (C 25): Data from the CIE taxi cost survey in October 2011, less the variable costs.

Fuel cost (C 26): Calculated per kilometre from the CIE taxi cost survey in October 2011.

For calculation of peak licence value, annual fixed costs to be apportioned over fewer shifts (C 29): An adjustment for PALs because their annual fixed costs must be recovered over fewer shifts than for an unrestricted licence.

Current taxi licence numbers (C 32 to C 35): Data from Transport for NSW; taxi licence numbers as at 1 January 2013.

Hours per shift (C 41): Data from the CIE taxi cost survey 2011.

Average trip length (C 42): Average distance in kilometres for taxi trips in Sydney (average weekday) 2010/11. Data provided from the Household Travel Survey by the Bureau of Transport Statistics.

Total km travelled per taxi per year (C 43): Data from the CIE 2011 survey of taxi drivers and operators.

Number of shifts per taxi per year (C 44): Data from the CIE taxi cost survey in October 2011. Average for all standard taxis (including PALs).

Share with booking fee (C 45): An assumption about the proportion of trips for which a booking fee is charged. IPART has used this assumption in our calculation of the average fare each year for fare reviews, based on a review of a range of sources.⁹ Data from the CIE taxi cost survey 2011 also supported this assumption.

Speed when on waiting time rate (C 46): An assumption that the average speed when on the waiting time rate is half the waiting time rate threshold (26 km/hr).

Average speed (C 47): An assumption used to calibrate the base model so it matches 2012/13 outcomes.

Revenue at waiting time rate/revenue at distance rate (C 48): Data from the Australian Taxi Drivers Association. Used to calculate the average minutes of waiting time per taxi trip.

⁹ IPART, *Review of fares for taxis from July 2008 - Final Report*, June 2008, p 57.

Minutes on waiting time rate per taxi trip (C 49): An assumption about the average number of minutes of each taxi trip charged at the waiting time rate rather than the distance rate. Calculated using data about revenue split between waiting time and distance rate (C 48) and average trip length (C 42), and an assumption about the average waiting time speed (C 46).

Kilometres on distance rate (C 50): An assumption about the average number of kilometres of each taxi trip charged at the distance rate rather than the waiting time rate. Calculated using data about revenue split between waiting time and distance rate (C 27) and average trip length (C 42), and an assumption about the average waiting time speed (C 45).

Multiplier for occupied to total km for a trip (C 51): The inverse of the occupancy rate (per kilometre). Data from the Australian Taxi Drivers Association.

Value of time (C 52): An assumption about the value of a passenger's time spent waiting for a taxi, used to convert waiting times to dollar values. Putting a value on time is standard practice for cost/benefit analysis, and the dollar value is usually lower than the \$30 per hour proposed by the CIE; on the other hand, the Victorian Taxi Inquiry used a value of \$60 per hour.

Constraint on number of cabs (C53 to C55): Assumptions about the effective maximum number of cabs that can be on the road (as opposed to the notional maximum, which is the number of licences). Based on data from networks about the share of cabs on the road per shift. The effective maximum is determined not only by the percentage of cars which are off the road for maintenance or repair, but by the behaviour of operators, some of whom prefer to drive their cab one-out and leave it off the road some of the time rather than engage a bailee driver, and it varies by shift.

Slope of the supply curve (C 56): The slope is defined as the change in fixed costs per shift per additional 5000 cabs on the road to roughly relate to the difference in fixed costs from a very quiet to very busy shift.

Standard deviation of expected shift revenue (C 57): An assumption that accounts for the variation in demand for taxi services and therefore revenue over a year. Set to calibrate the base model to 2012/13 conditions.

Calibration constant to account for rising marginal costs (C 58): This is applied if 'upward sloping cost curve' is selected at cell C 59 to ensure that the balance between fixed and variable costs per shift is correct.

Flat or upward sloping cost curve (C 59): This is an assumption about the shift costs for lower and higher demand shifts. Selecting 'upward-sloping' at this point introduces an assumption that costs will get higher the more taxis are on the road (which captures both that the demand for driver labour is higher and earnings must be higher to attract more drivers to the shift and that operators might recover somewhat more of their fixed costs from higher demand shifts). A flat cost curve assumes a sufficient supply of drivers such that additional taxis can find drivers at the same earnings rate (per hour) and that operators recover fixed costs equally across shifts.

What revenue premium can peak licences obtain by working across best part of shifts? (C 61): An assumption used to calibrate the model to match 2012/13 conditions.

Of shifts available, what per cent do peak licences do? (%) (C 62): An assumption to allow for peak availability taxis to be off the road for maintenance and repairs or for other reasons for a portion of shifts.