A New Approach to Calculating the Benefits Associated with Infrastructure Investment

Simon Sagerer and Nick Wills-Johnson

Abstract

It has long been standard practice in economic theory and practice to assess the benefits associated with the consumption of a good or service based upon the consumer (and producer) surplus associated with consuming the good or service. To do this requires the estimation of a demand curve. However, benefit cost analysis of infrastructure projects in Australia is commonly based upon point estimates of demand and the consequences of the forecast level of demand, rather than an assessment of demand curves and consumer surplus. Under a set of very restrictive conditions, this approximates consumer surplus, but in general, it does not. In this paper, we outline a methodology for using information very similar to that used by policymakers in Australia in the context of point forecasts and their consequences to derive a demand curve for a given mode of transport, and thus to estimate consumer surplus. The result is a more robust cost benefit analysis framework.

Key Words: transport, infrastructure, cost-benefit analysis, consumer surplus.

October 2011

ACIL Tasman

Economics Policy Strategy

Introduction

This paper presents a novel method, at least in the context of infrastructure policy in Australia, of assessing the benefits associated with infrastructure investment. From the perspective of basic economic principles, the benefit associated with the consumption of any good or service is the sum of the consumer and producer surplus. However, this is not, in general, how assessments are made by Australian policymakers. Instead, forecasts of demand are made (usually through some exogenous modelling process) and the various societal benefits associated with that level of demand (reductions in pollution, congestion changes, economic activity generated and so on) are then calculated. The basic problem with the Australian "standard approach" is that it approximates true benefits only under a very limited set of circumstances; otherwise, it is simply wrong.

The authors were faced with this issue in a recent project in Hobart, which aimed to establish the benefits associated with the development of a new rail service through the Northern suburbs of that city. Dissatisfied with the "standard approach", we developed the novel approach outlined in this paper, which estimates consumer and producer surplus directly. Although the particular application is to an item of transport infrastructure, in principle, the approach could be applied to any form of infrastructure; all that is required is some way to estimate the total monetary and non-monetary costs associated with consumption of the service, and some way to differentiate consumers in order to simulate sufficient point estimates of demand under different scenarios to derive a demand curve.

Section Two of this paper outlines the flaw associated with the "standard approach". Section Three outline how we came to our approach, as the path followed was not direct. In this section, we also explain, in general terms, how the approach works. Since it is difficult to fully appreciate the operation of our approach in abstract, we present a case-study, drawn from our work in Hobart in Section Four. Section Five draws some conclusions for future work.

The flaw of using point estimates

The benefits associated with the consumption of a good or service, from the perspective of society as a whole, are the sum of the consumer and producer surplus associated with that good or service; the area between the demand and supply curves. By way of a simple example, consider a consumer who would have paid $15 for a bottle of wine, but finds it for sale in a bottle-shop for only $10. The consumer receives a benefit of $5 from being able to buy the wine for less than she would have been prepared to pay. If the retailer would be prepared to sell the wine for any price above $8, then receives $2 benefit from a sale price of $10. The overall benefit to society of the wine is the sum of the benefits that consumers receive because the wine costs less than they would have been willing to pay, and the benefits that the retailer receives because she has received more than the cost of supplying the wine. The actual price paid is not relevant, as this is just a transfer between two parties; the purchaser and the bottle-shop owner.

Infrastructure is intrinsically no different from any other good or service in terms of the use of consumer surplus to assess its benefits. This is clear in the academic transport economics literature (see Winston & Maheshri, 2007, Harford, 2006 & Bagliano et al 2007 for three examples). The only major difference in respect of public transport infrastructure compared to other goods and services is the shape of the relevant supply curve; it is usually flat (or at least stepped) given the economies of scale of infrastructure production, and may in fact intersect the demand curve at a point above that where government wishes prices to be. This then requires a subsidy be paid. Both the general and the infrastructure-specific situations are described in Figure 1. The left hand side is the standard approach in a competitive market (like the wine example above) and the right hand side describes the situation as it exists for infrastructure (for example, a railway). In the railway example, the fact that the price level government deems appropriate is lower than that which would be supplied by the private sector is reflected by the producer loss. In a benefit cost assessment, this would need to be subtracted from the consumer surplus. If the price is not set lower than the cost of delivery, then no subsidy is required. If the first step of the supply curve appears to the left of the intersection with the demand curve (rather than to the right, as shown here) then some producer surplus will need to be added to the consumer surplus.

Figure 1 **Graphical representation of consumer surplus**

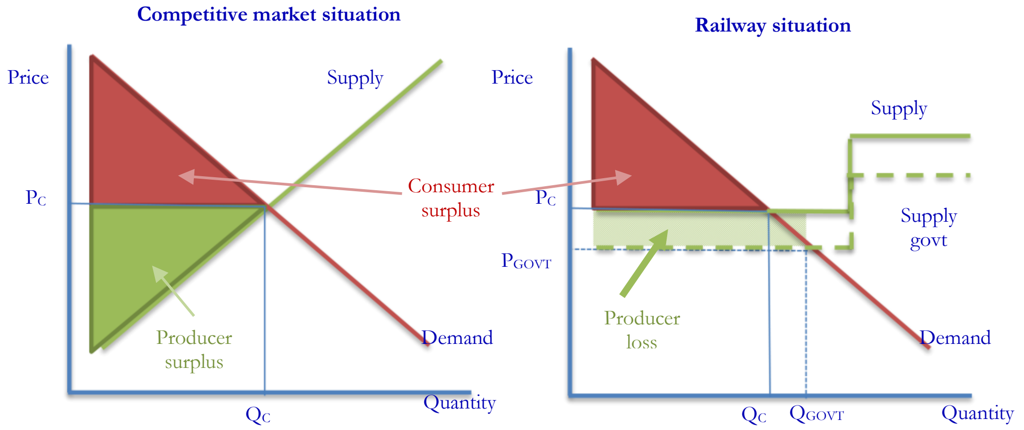

Note that the consumer surplus in the right-hand diagram extends down as far as P_{GOVT} and across as far as Q_{GOVT}, but is obscured by the losses shown for the producer.

As noted above, assessment of benefits of infrastructure in Australia by government agencies (see, for example, the assessments undertaken by Infrastructure Australia – which match those undertaken elsewhere in government - at www.infrastructureaustralia.gov.au/2011_coag/project_assessments.aspx) does not follow an approach based in consumer surplus, but instead relies upon point forecasts of demand, and analysis of the consequences of a given level of demand; rather than a demand curve and associated consumer surplus.[1] Since the bases of the two approaches are different, there are difficulties in comparing them directly, but an endeavour to do so is shown in Figure 2. Since the consequences of choosing a point estimate over a demand curve as the basis of analysis differ depending upon the shape of the demand curve, two separate cases are shown; one with a linear and one with a curved demand curve. In each case, the (flat) supply curve has been normalised to zero for ease of presentation; the basic results would not change with a flat supply curve at a different level, or an upward-sloping supply curve.

[1] Note that other jurisdictions do use consumer surplus. One example is the TUBA model used in the UK (see www.dft.gov.uk/publications/tuba-downloads-and-user-manuals)

Figure 2 **Point estimates and consumer surplus**

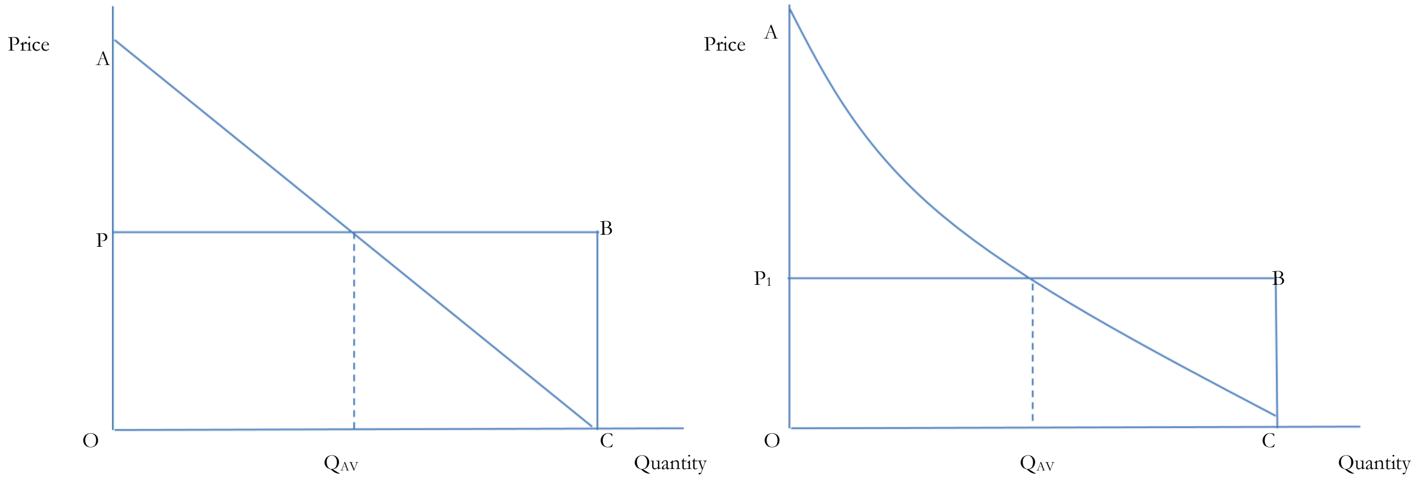

In each case, the consumption level of the average consumer is the same Q_{AV}. In the left-hand diagram, the demand curve is linear. This means that, geometrically, the areas of the rectangle *OPBC* and the triangle *OAC* are equal. The area of the triangle *OAC* is, of course, the consumer surplus associated with the consumption of this particular infrastructure service. The area of the rectangle *OPBC* does not have a simple economic interpretation, in terms of consumer surplus. The distance *OP* is equal to the consumer surplus enjoyed from consumption of the relevant infrastructure service by the "average person" in the marketplace. Thus, the area of the rectangle *OPBC* is the consumer surplus which would be enjoyed if every consumer in the relevant market (that is, along *OC*) was exactly at the average; if consumers were all exactly homogenous and the point estimate of demand by a "representative consumer" was exactly correct. There is no economics in this equivalence, but it rather results from a very simple piece of geometry. The basic interpretation is that, provided demand curves are linear, and the forecast demand level predicts exactly the demand of an "average consumer", then there are no problems associated with using point estimates of demand to conduct cost-benefit analyses.

Predictions associated with the "average consumer" are, however, fraught with difficulty. Take for example, the issue of the value of travel time; as Douglas & Karpouzis (2011) show, there is little agreement on what the appropriate level is. Moreover, there can be important equity considerations if different averages are used for different regions, but if the "lowest common denominator" is used, then too little transport infrastructure is built. This problem of averages, however, is avoided in an approach based on estimating demand curves, as the different values of time are just one aspect of consumer heterogeneity.

Consider now the right-hand case, where the demand curve is now curved. Since it is curved downwards P_1 is less than *P*.[2] Now, the area OP_1BC is not the same as the integral under the demand curve *OAC*. The same would be true of *OPBC*, if it were shown. This implies that the consumer surplus result for the average consumer can no longer be extrapolated to the whole population as it could be for a linear demand curve even if the point estimate is accurate. This means that the point estimate approach will always be

[2] We could have kept price the same and adjusted quantity; the interpretation would be the same.

wrong.[3] It will over or under-estimates actual consumer surplus depending upon whether demand curves are concave or convex to the origin.

There is a further complication. If the demand curve is linear, elasticity of demand will vary all the way along the demand curve. This means that any estimate of elasticity calculated within the context of a demand modelling exercise which involves point estimates of demand will only be valid at the point along the demand curve for which they are estimated. By contrast, if demand curves are iso-elastic (they have the same elasticity along their entire length), they must be curved. The practical up-shot of this is that one cannot accurately estimate elasticity and consumer surplus jointly using a point estimate of demand for a representative consumer and extrapolating for the population as a whole.

The genesis of our approach

Our approach is grounded in the notion of the consumer surplus representing the benefits to society (or welfare consequences) of a particular transport infrastructure project being developed. More specifically, we look at the modal options before and after the introduction of the new service, calculate the demand curve for each move before and after the service is introduced, assess the consumer surplus associated with each demand curve and sum the consumer surpluses before and after to ascertain how consumer surplus has changed. Once the demand curves are estimated, consumer surplus estimation is a relatively simple application of integral calculus. Thus, our main focus below is on explaining how we derive demand curves.

Ordinarily, economists estimate demand curves by looking at data on consumption choices made by consumers. The problem in estimating the demand curve for a new piece of infrastructure is that, since the service does not yet exist at the time when the analysis is undertaken, there are no data on demand from which to estimate the demand curve. We solve this problem by looking at the characteristics of the population which would use the infrastructure being proposed, and, based on the characteristics of the population and the service being proposed, estimate how much of it the population would demand, and what other modal shifts this would drive.

The logical starting position for any consideration of demand is utility; from a utility function, a demand curve can be derived, and thus consumer surplus estimated. We therefore began our consideration of a method of calculating consumer surplus with the model of Parry & Small (2009). This model considers utility to be associated with consumption of a numeraire good and with transport. Transport utility is then a function of aspects of the trip including its cash cost, the time taken, waiting times, congestion of the vehicle and the service frequency of public transport. From this relatively simple model, Parry & Small (2009) are able to derive demand curves and examine questions associated with urban public transport subsidies.

[3] There will always be a point *i* along *OC* such that the rectangle OP_iBC equals the integral *OAC*, but the point is that this will not be at the average consumer point. Rather, the exact "representative consumer" *i* will depend upon the shape of the demand curve. Thus, it is still possible in-principle, to use a "representative consumer" approach, but since one needs to know the exact shape of the demand curve to identify the representative consumer and the shape of the demand curve is enough to estimate consumer surplus, it is pointless to do so.

The model is useful in a context where one has data about actual public transport use (which Parry & Small have), as the relevant parameters can then be derived through suitable regression analysis. However, our requirement is to develop a model to examine the societal benefits of developing new services which, by definition, do not yet have data on their use which can be fed into an econometric regression to derive the relevant parameters. In theory, one could take the parameters directly from Parry & Small (2009), but this would require an underlying assumption that the characteristics of transport in the city being examined are like those in Parry & Small's (2009) sample. Since Parry & Small (2009) examine London, Washington and Los Angeles and our application was to Hobart, this is unlikely to be true. Thus, Parry & Small's (2009) model, although it provided useful elements (see discussion below) could not be used in its entirety.

Our next approach was to adopt a Becker (1981) type household production function approach, wherein households are presumed to use transport and other services to produce utility. However, unless we adopt very restrictive assumptions about the shape of the production function, we are faced with essentially the same problem; econometric estimation would be relatively simple, but because a new service is being estimated, we cannot undertake it.

The result of our work with utility and production functions was to develop a much simpler approach; one based upon cost. Put simply, each consumer has a number of modes from which to choose for a given trip. Each mode has both different combinations of elements in the cost of travel, and different levels. Costs are monetary (bus fares, fuel prices, parking fees and so on), time related (transit time, waiting time) and related to other aspects of the trip, such as its perceived safety, greenhouse emissions or some other non-monetary, non-time related cost. The various costs can be combined in a cost function for travel by each consumer using each of the available modes, and thus we are able to derive a "price" for a given trip, which can be combined with a quantity measure such as the number of kilometres in that trip.

We then make an assumption about each consumer; from the available choices of mode, she chooses the lowest-cost option, taking into account the monetary, time and other costs each mode would impose upon her when making a given trip. We then record, for each consumer, the choice made; the price-quantity combination involved in the bus, train or car (say) trip that she chooses to make because it is the lowest-cost way of getting from where she is to where she wants to be.

If all consumers were entirely homogenous, and making exactly the same trip, we would obtain just one point forecast, and not improve upon existing approaches at all. We thus introduce heterogeneity in order to create a demand curve.[4] This can be done in several ways:

- Firstly, the decision process could be repeated for a consumer with different characteristics (different co-efficients for the same elements in his cost function) but located in the same physical location (say along a proposed railway line). Such a consumer might be richer, or poorer, and thus place a different value on time savings vs monetary costs, or he might have greater or lesser concerns about the environment, and thus place a different value on the green elements of a trip. Different types of trips can be addressed in this way; a young man travelling to work has different characteristics compared to

[4] This is what we mean when we suggest that we make use of much of the information currently used by policymakers in Australia (in fact we use more, considering the variance around estimates as well as their mean) but put the information together in a different way.

travelling home from the pub on a Friday night, and thus makes different modal choices, even when the origin and destination are the same.

- Secondly, we can examine consumers who are exactly the same, but located at different places along the line; an executive in her mid forties living near the centre of town and travelling to town for work might make different modal decisions because of where she lives compared to her sister living further out.
- Thirdly, we can take into account the fact that the elements in the cost function are not exactly known, but vary. For example, fuel prices vary considerably, and exactly the same person, travelling from the same origin to the same destination, might change her modal choice if fuel prices increase to $1.50 from a base of $1.00 per litre. Variance in travel time costs shown in Davis & Karouzis (2011) could be included here.

In essence, adding heterogeneity in the manner outlined by the three points above creates different ways of asking the same question; "what choice would this type of person make for this kind of trip, given this set of elements in a cost function?" and thus a different answer (in terms of the mode chosen) each time the parameters in the question change. We record the answers in a scatter plot, one for each mode, which shows the quantity (distance travelled) and price (total cost of that travel). Each time the question is asked, a new point is plotted on the scatter plot of one of the modes. By asking the question several hundred thousand times, and writing down the price-quantity combination that each cost-minimising choice represents, we create very dense scatter plots for each mode of travel in the population being examined.

Having created the scatter plots, we fit demand curves using regression analysis. These curves may be linear, but there is no requirement for them to be so; we make our choice on the functional form which best fits the actual data. Once each demand curve is known, since the supply curve is assumed to be a straight line, the consumer surplus is simply the integral under the demand curve from the origin to its intersection with the supply curve.[5] To establish the net societal benefit from the introduction of the relevant service, we simply sum the consumer surpluses for each mode before and after the service has been introduced in the model, and compare the difference.[6] Conceptually, we can take as many "after" snapshots as we like, to capture aspects such as dynamic interaction when a rail service (say) decreases congestion, making cars more attractive and thus bringing people back onto the roads, or the land use reaction as developers respond to a new rail line by creating denser transit-oriented development near stations. The important point is that, because we examine the consumer surplus for all modes, we can pick up benefits such as the benefits for those who remain on the road receive when some of their fellow car-drivers begin to use the train and leave the road less crowded.

Finally, because there is a "before" snapshot, the model can be at least partially validated; the modal choices it predicts for consumers of different characteristics or in different locations should match those which a

---

[5] As flat supply curve is a reasonable assumption to make for public transport infrastructure. However, the methodology does not require a flat supply curve to work; it is very simple using calculus to establish the area between two curves. Conceptually, an approach similar to that we use to derive a demand curve (that is, recording supply choices; the difference being that they would be profit maximising, not cost minimising) could be used to derive a supply curve. Alternatively, the underlying engineering aspects of the infrastructure could be used to derive a cost function for its supply. There are many alternatives which might be used.

[6] Note that we include changes in the costs of subsidies to support public transport as well; including the cost of taxes required to provide the subsidies.

travel survey of the city (say) shows actually occurs. Since the process of deriving the demand curves results in recording exactly which consumer under exactly which situation made what modal choice, the model generates a lot of data with which to make such comparisons.[7] Moreover, since it is the set of available modal choices which changes as new transport infrastructure is introduced, and not the underlying characteristics of the consumers themselves, if the model predicts choice accurately before the new service is introduced, there is some confidence that it will also do so afterwards.

In the following section, we provide a case study of the approach outlined above, which helps show how it operates in practice

A case study example: Hobart light rail

In this section, we provide an overview of our approach in operation, in the context of an application of the model to estimating the benefits associated with the development of a passenger rail service in Hobart. The full report from which this section of the paper is drawn can be found at www.transport.tas.gov.au/miscellaneous/northern_suburbs_to_hobart_cbd_light_rail_business_case.

The rail service being assessed in the case study was to run from Hobart to Claremont, a distance of roughly 15 km. The line itself is shown in Figure 3.

The study examined a number of different patronage scenarios formed by changing parameters in the model associated with the operating characteristics of the rail system. It found that in all but a handful of cases, the benefits of the rail system were far outweighed by the costs; in most cases, the benefits were not even sufficient to cover its operating subsidies. This is because the configuration of the track means that journey times by train would have been longer relative to cars, even in peak periods, for most users. This means that few consumers choose rail, and those that do receive small benefits from doing so. The only scenario we examined where the benefit cost ratio was greater than one was one where we artificially increased the speed of the train in the model to be faster than a car (something which is physically impossible given the track configuration – which would be very costly to alter) to proxy the intrinsic attractiveness which many observers suggest trains have over other forms of public transport, and which might cause people to favour the train, even though it is a little slower than cars for commuter traffic. However, even this scenario achieved a benefit cost ratio only slightly in excess of one, and this declined rapidly with only small reductions in patronage.

---

[7] It also generates data which can be used in more formal regression equations to answer questions such as "if the price of fuel goes up by ten percent, how many people will switch to bus?". This can be very useful in decomposing sources of demand, though of course, one would need to adequately deal with the errors made in the forecasts of modal choice used in deriving the demand curves when undertaking subsequent econometric analysis. However, bootstrap regressions (which are not all that different conceptually to what we do here) provide suitable bases for doing these regression analyses with confidence.

Figure 3 **Proposed alignment of rail service in Hobart**

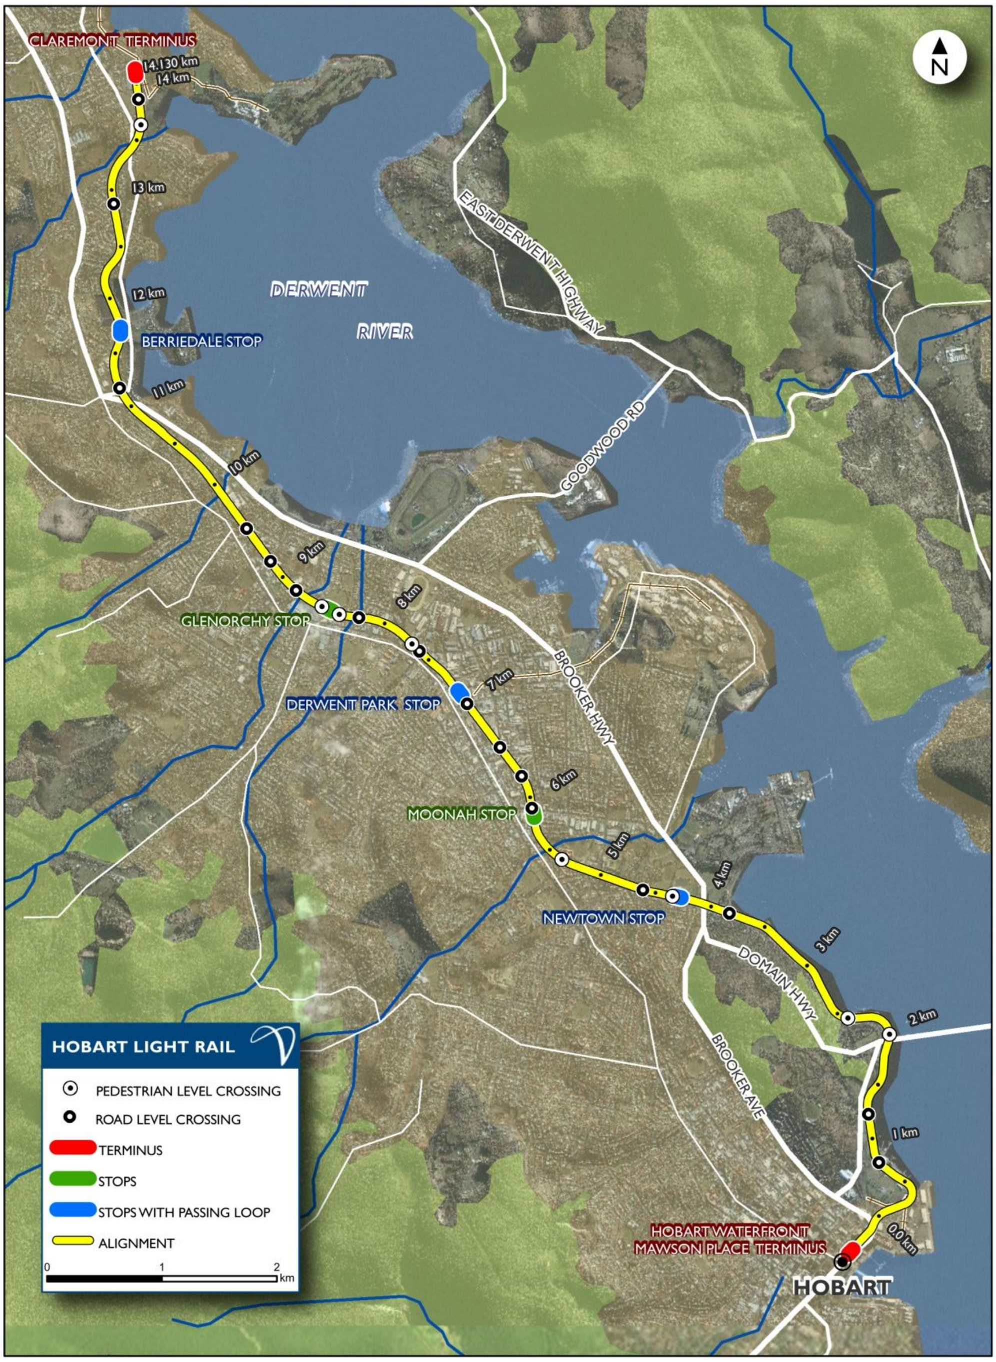

Source: Hyder: available from: www.transport.tas.gov.au/miscellaneous/northern_suburbs_to_hobart_cbd_light_rail_business_case/what_is_stage_2

The line would have drawn patrons from a much wider geographic area than shown in Figure 3; a large park 'n ride at the terminus (as well as numerous smaller facilities along the route) was designed to capture patrons who would drive to the terminus from areas north of Claremont, and an assumption was made that monies saved from removing competing bus services (mostly between Glenorchy and Hobart) would be put into improving feeder-bus services from Bridgewater and Brighton to the north of the terminus, further driving demand for public transport; it was the change in the transport system rather than just the addition of a single line that was modelled.

Most of the traffic along the route would have likely been commuter traffic, which is mostly directed towards either Hobart City or Glenorchy (both key loci of employment in Hobart), and some data were available on proportions of commuters travelling to each from various regions within the catchment area. However, significant demand would also likely have come from non-commuter traffic, which is much more diffuse in its destinations. For the purposes of assessing demand, we therefore had to divide passengers into those using the service for commuting (directed travel), and those using it for non-commuting purposes (non-directed travel). Amongst the directed travellers, we differentiated between those using it for work, and those using it to get to school or higher education. We also differentiated on the basis of location, choosing a "representative consumer" from each of 204 statistical collection districts from the Australian Bureau of Statistics Census. We used these because they are the smallest unit of measure for data from the Census, and provide a rich set of socio-economic data with which to populate the model.

We assume in the model that five modes of travel are available:

- Drive a car.
- Walk to the bus stop and catch a bus.
- Walk to a train station and catch a train.
- Walk to a bus stop and catch a feeder bus before catching a train.[8]
- Drive to a train station and catch a train via park 'n ride.

Prior to the development of the railway service, only the first two options are possible, and there are thus four demand curves to derive and four consumer surpluses to calculate (directed and non-directed travel by each mode). Subsequent to its development, all five modes are possible, and there are thus ten demand curves and ten consumer surpluses. The sum of the four initial consumer surpluses are subtracted from the sum of the ten subsequent consumer surpluses to obtain the change in consumer surplus which is then used as the benefit component of the cost benefit analysis. In the project itself, we considered two different cases of land use post the development of the rail; one immediate and one after five years when transit oriented development had occurred, which we proxied by increasing the density around each railway station and decreasing it at the periphery of the study area (taking into account overall population growth). This had the practical effect of increasing the number of people taking the third option at the expense of all other options.

Given this very basic background, the sub-sections below provide detail on the way in which the model operates in the context of assessing the benefits of the rail service, specifically:

- The formation of the cost functions.
- The assumptions underpinning directed and non-directed travel.
- The formation of the demand curves and the calculation of consumer surpluses.

[8] We assume in our analysis that feeder bus timetables match the rail timetable perfectly at every rail station so there is zero waiting time between bus and train.

Formation of the cost functions

As noted above, the "cost" of travel has a number of different cash, non-cash and time components. That is, the time taken for a given trip is considered costly by consumers; all else being equal, they will choose the mode that allows them to make a given trip in a shorter time. The costs we consider are thus:

- The costs of bus and train tickets.
- The costs of fuel and vehicle operation.
- The costs of parking.
- The value of time spent travelling.
- The societal cost of environmental pollution.
- The costs of road crashes.
- The costs of social exclusion.

Other costs (perceptions of personal safety, for example) could have easily been included, but these were the main costs which were important in the context of the client for this study. These elements are combined in cost functions for each of the five trips outlined above which incorporate three basic elements:

- A cash cost in the form of ticket prices or fuel and parking costs.
- A travel time and (except for cars) waiting time cost.
- A per kilometre value for pollution and road crashes.

For some consumers who are deemed socially excluded (see discussion below), a countervailing benefit associated with the alleviation of social exclusion through better (public) transport options is included.

The resultant cost functions are shown below. The cost functions used are based on a utility function developed by Parry & Small (2009). In this framework the cost of travel is determined by the direct monetary expenses as well as the time spent travelling, waiting and accessing each mode and external effects such as pollution. Furthermore we have introduced two speed zones in order to take urban congestion into account. The variables and parameters of the cost function are shown in Table 1.

Table 1 **Variables of the cost functions**

Description	Variable name	Affected mode
Distance to work car	*DC*	Car
Distance to work bus	*DB*	Bus
Distance to work train	*DT*	Train
Distance to closest train station feeder	*DF*	Feeder (train)
Distance to closest P&R	*DPR*	P&R (train)
Distance to closest bus stop	*WD_B*	Bus
Distance to closest train station	*WD_T*	Train
Distance to closest feeder stop	*WD_F*	Feeder (train)
Distance non-directed trip	*DL*	All
Average speed car outer sector	*CSo*	Car
Average speed car inner sector	*CSi*	Car
Average speed bus outer sector	*BSo*	Bus
Average speed bus inner sector	*BSi*	Bus
Average speed train	*TS*	Train (feeder and P&R)
Average speed walking	*WS*	Bus
Waiting time	*WT*	Bus
Fuel price	*FP*	Car
Average fuel consumption	*FC*	Car
Ticket price	*TP*	Bus
Value of travel time	*VTT*	Car and Bus
Inner distance	*ID*	Car and Bus
Parking cost	*PC*	Car
Pollution cost	*EPC*	Car
Road crash cost	*ERC*	Car

The cost functions (C) by mode and journey type are:

$$C_{directed}\ Car = DC * FC * FP + \frac{DC-ID}{CSo} * VTT + \frac{ID}{CSi} * VTT + EPC * DC + ERC * DC + PC$$

$$C_{directed}\ Bus = TP + \frac{WD_B}{WS} * VTT + WT * VTT + \frac{DB-ID}{BSo} * VTT + \frac{ID}{BSi} * VTT$$

$$C_{directed}\ Train = TP + \frac{WD_T}{WS} * VTT + WT * VTT + \frac{DT}{TS} * VTT$$

$$C_{directed}\ Feeder = TP + \frac{WD_F}{WS} * VTT + WT * VTT + \frac{DF}{BSo} * VTT + \frac{DT}{TS} * VTT$$

$$C_{directed}\ P\&R = TP + \frac{DPR}{BSo} * VTT + WT * VTT + \frac{DT}{TS} * VTT$$

$$C_{non-directed}\ Car = DL * FC * FP + \frac{DL}{CSo} * VTT + EPC * DL + ERC * DL \ \ /DL$$

$$C_{non-directed}\ Bus = TP + \frac{WD_B}{WS} * VTT + WT * VTT + \frac{DL-WD_B}{BSo} * VTT \ \ /DL$$

$$C_{non-directed}\ Train = TP + \frac{WD_T}{WS} * VTT + WT * VTT + \frac{DL-WD_T}{TS} * VTT \ \ /DL$$

$$C_{directed}\ Feeder = TP + \frac{WD_F}{WS} * VTT + WT * VTT + \frac{DF}{BSo} * VTT + \frac{DL-WD_F-DF}{TS} * VTT \ \ /DL$$

$$C_{non-directed}\ P\&R \ = TP + \frac{DPR}{BSo} * VTT + WT * VTT + \frac{DL-DPR}{TS} * VTT \ \ /DL$$

We now discuss the parameters of each of these cost functions in more detail.

Cost function parameters

In this section, we provide an overview of each of the parameters used in the cost functions outlined above. Before doing so, it is important to digress slightly and explore a parameter not explicitly mentioned in each cost function; the area from which each representative consumer comes.

Most of our parameters are associated with vehicle speeds, as it is travel times which drive much of the resource costs. Data on vehicle speeds have come from studies by DIER (2011) and DIER unpublished data in respect to bus speeds. We assume that speeds drop from free-flow speeds to congested speeds at a point just north of Glenorchy (DIER, 2011), and we allow all speeds (except those of the rail service) to vary in the creation of demand functions. Fuel prices (which we also allow to vary) are from the Australian Institute of Petroleum and data on pollution costs come from the Australian Transport Council (2006) while those on road crashes come from BITRE (2009). These are the latest figures available in each case. Ticket prices (for all forms of public transport in Hobart) and parking costs were based on information provided by the client. Other parameter information is provided below.

Waiting time

A detailed analysis of average waiting time was conducted by Parry & Small (2009). They find that for frequencies (*H*) of 15 minutes or less, passengers arrive randomly. For larger headways travellers tend to plan their trip. These results lead to the following waiting time functions:

$$WT = \begin{matrix} [H/2\,]/60 & for\ H \leq 15 \\ [6 + 0.2 * H/2]/60 & for\ H > 15 \end{matrix}$$

Average fuel consumption

In 2002 the Australian Bureau of Statistics (ABS) conducted a detailed survey on fuel consumption in Australia in which it reports total fuel consumption in this year (ABS, 2002). BITRE (2009), in turn, provides figures on the total number of driven kilometres per year, state and vehicle class. From these datasets we calculated an average consumption per kilometre. The resulting value is 7.2 litres per 100km. We applied a 25 per cent mark-up to account for stop-and-go driving when roads are congested.

Value of travel time

The cost of travel time is generally valued as a fraction of the average gross hourly wage. According to the ABS (2011) the average gross hourly wage in Tasmania is about $28. We assumed a travelling time discount factor for Australia of 60 per cent, which is half way between the average net to gross income ratio (70 per cent) and similar discount factors estimated in the USA by Parry and Small (2009) of 50 per cent.

Assumptions underpinning directed and non-directed travel

The cost functions outlined in the previous section are applied to directed and non-directed travel. Directed travel is travel between a given Census collection district and a given place of work or education, whilst undirected travel is travel from a given Census collection district on a trip of uniform length, but no particular direction. In this section, we detail further the assumptions underpinning these trips.

Directed travel

As noted previously, directed travel essentially involves travel from home to work or school (and back). The only relevant choice is overall resource cost, which representative consumers are assumed to minimise.

Work trips are undertaken by all employed people (based on Census data) in each collection district who are over the age of 15. We are constricted somewhat by age bands in the Census data in this respect, and we thus assume that all people between the ages of 15 and 19 who travel for work do so by public transport. This likely over-estimates public transport travel somewhat, as some of these people will have a licence and a car, and will drive to work. However, the numbers of people are not large. All people 20 and above have the full choice between the five types of trip outlined previously.

Work locations are either the centre of Hobart (70 percent) or the centre of Glenorchy (30 percent), based upon Census data on the numbers of people employed in these two local government areas. We recognise that employment is spread more widely than this, but the level of data required to fully specify work origins and destinations does not exist at present, and it was beyond the scope of the analysis to collect it. However, since the main focus of the analysis is to estimate the benefits associated with the rail system, the problem is not overly large; we will over-estimate the consumer surplus associated with each mode by making trips longer than they actually are, but our main focus is the change wrought in consumer surplus as the rail system is introduced, and this is not affected if all trip types are affected equally.

School trips are estimated based upon those classified in the Census as students. Those less than 15 are assumed to walk or be driven to school. Those between 15 and 19 are assumed to take public transport if the nearest stop is closer than 800m (which it usually is). Those over 19 have the full choice of modes. We assume schools are located in Hobart, rather than along the line, which has the effect of increasing the length of a train ride, but not the number of rides. As with workers above, the effects are small.

Non-directed travel specific assumptions

Non-directed trips include trips for leisure, trips to visit tourist attractions and special events, trips for shopping and other similar trips without an easily definable destination; essentially, everything except regular trips to work or education. . Since the range of possible destinations is so large, rather than trying to assign particular destinations for these kinds of trips (such as MONA or the shopping centre at Glenorchy) we assume that each representative consumer makes undirected trips of uniform length until the marginal utility from doing so is exceeded by the cost of the final trip. The uniform length of these trips is 11 km, which is the average for "leisure" trips in the Greater Hobart Travel Survey (DIER, 2010).

The major difficulty in our approach is that measuring utility between people is widely considered to be impossible. We therefore make use of the cost functions, the Greater Hobart Travel Survey and the theory of revealed preference (see, for example, Varian, 2006).

The survey reports that each person undertakes an average of 2.7 trips on a weekday and 4.4 trips on weekends. Assuming that 2 of the weekday trips are directed (e.g. to work), each individual makes about 1.28 non-directed trips per day. From the cost model, we know that a car trip (the dominant mode) costs $0.52 per kilometre, so if each person travels until marginal cost equals marginal revenue, the marginal utility of the last trip made by the marginal consumer should be $0.52.

Another finding from the survey is the standard deviation from the average number of trips which amounts to 0.36 (DIER, unpublished data from DIER 2010). Assuming normal distribution, this means that 95 per cent of the travellers undertake less than 2.02 non-directed trips per day. Therefore we set the marginal utility of every trip additional to 2.02 to zero.

This provides us with two data points upon which to fit a curve. Economic theory generally assumes utility functions are curved, so we make use of a logarithmic function; implicitly assuming that marginal utility decreases at a fixed rate. The resulting marginal utility function is shown in Figure 4.

Figure 4 **Marginal utility of non-directed trips**

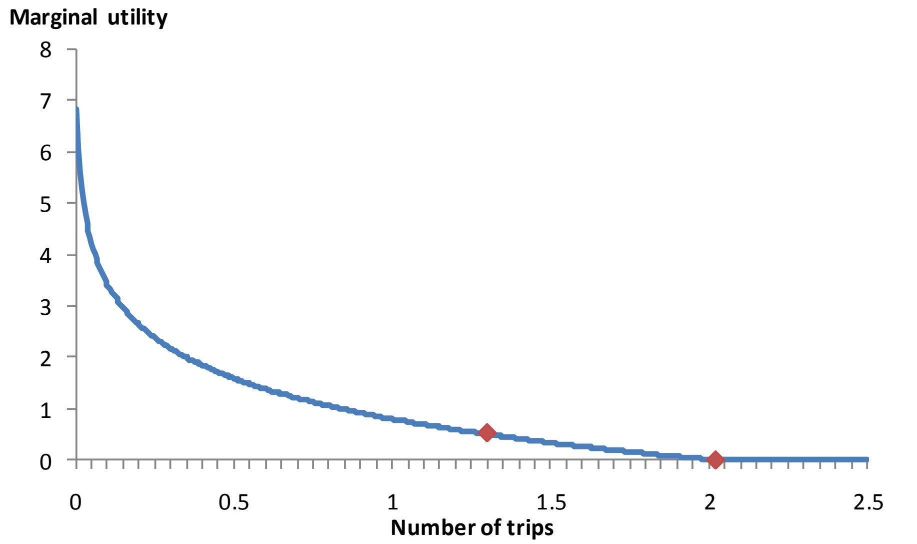

Source: ACIL Tasman modelling based on the Greater Hobart Household Travel survey *(Department of Infrastructure, Energy and Resources, 2010)*

We make use of this utility function by allowing every consumer to keep making leisure trips until the cost of the last trip she makes exceeds the relevant point on the curve in Figure 4 above.

Social inclusion and exclusion

One important aspect of non-directed trips in the context of the Hobart case study is social inclusion and exclusion. Social inclusion is a term that refers to whether a person has the resources, opportunity and capability to learn, work, engage (connect with people, use local services and participate in local, cultural, civic and recreational activities), and have a voice; influence decisions that affect them (Department of

Prime Minister and Cabinet, Social Inclusion Unit, 2008). Social exclusion occurs when constraints prevent adequate participation in these activities.

For socially excluded people, the ability to access transport is crucially important as it provides them with an opportunity to engage with the wider community. The client therefore wanted to find a way to include it in the analysis. One study of the value of transport to the socially excluded (Currie et al, 2010) suggests that it can be as high as $19.30 per trip, and that it declines with income.

The Census data do not allow us to delineate socially excluded people exactly as defined above (Currie, et al, 2010 made use of survey to do so), so we instead take as socially excluded all those in each collection district who are unemployed and earn less than $250 per week. For these people, we subtract from their (public) transport costs a social exclusion amelioration benefit derived from Currie, et al (2010, p13). We then take a weighted average of this cost, and the cost of travel for the non-excluded representative consumer in that collection district, and use this as the cost of (public transport) travel for that particular collection district. The net result is that there are more non-directed trips because of the presence of socially excluded people in an area, and that the number increases as the number of socially excluded people increases. As the public transport system improves (such as with the introduction of the NSLRS) and becomes more popular with people in general, its utility to the socially excluded increases still further because of the adjustment process outlined above.

Demand functions and the estimation of consumer surplus

As noted previously, scatter plots of modal choices made by different consumers (by location and characteristic) are used to derive demand curves. However, since this provides only limited variation and, more importantly, since the coefficients in the cost functions are themselves either measured with error or subject to variation, we expand the range of variation by considering different coefficient values. We do this by allowing several parameters to range across a set of values, and then drawing numbers from each range to simulate scenarios. This expands the set of situations we examine dramatically; to several hundred thousand choice scenarios. The elements of the cost functions we allow to vary, as well as their range, are shown in Table 2. Note that we could also have allowed other elements to vary (conceptually, any of them), but these elements gave us useful results.

Table 2 **Parameters and their variation in the model**

Variable	Min	Max
Average speed car outer (km/h)	53.0	57.0
Average speed car inner (km/h)	36.0	40.0
Average speed bus (km/h)	43.0	70.0
Average speed bus inner (km/h)	30.0	60.0
Fuel price ($)	1.2	2.0
Ticket price ($)	0.1	3.0

The result of asking the question of what modal choice would be made across a wide range of consumers, geographic locations and values of the parameters above, and recording the results in a scatter plot for each mode resulted in very dense scatter plots with quite distinctive shapes. Regression analysis was used to fit a demand curve to each scatter plot, and we found that an iso-elastic curve fitted the data best. An example of this, for buses, is shown in Figure 5. Note that the errors of the regression for each demand curve can be (and in this study were) used in sensitivity analysis around the benefit cost ratios that were eventually calculated. This is again much more robust that simply taking arbitrary extreme values, as usually occurs in cost benefit analysis for infrastructure in Australia.

Figure 5 **Bus travel scatter plot and fitted demand curve**

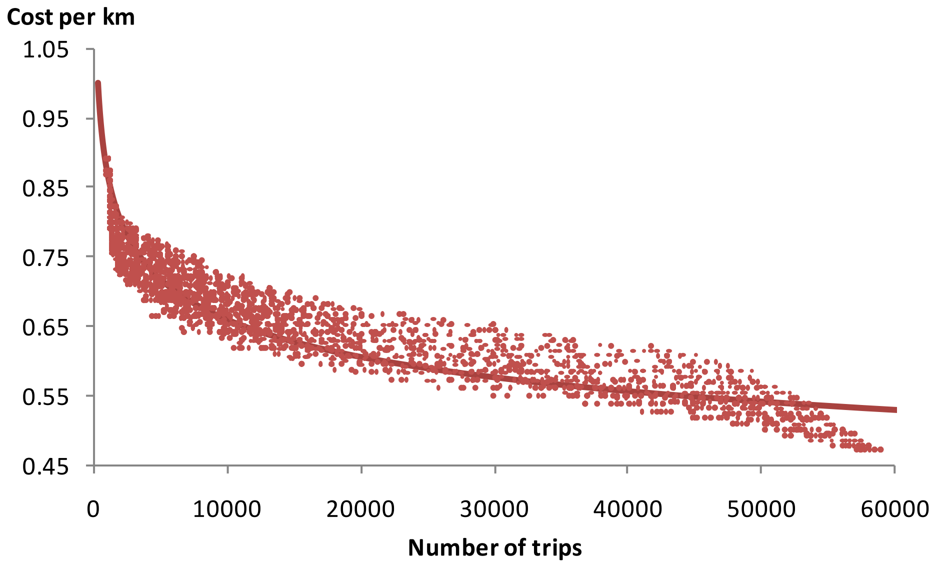

The consumer surplus is the area under each demand curve at the prevailing price. We approximate consumer surplus as the sum of rectangles made up by trip increments and the cost associated with each increment up to the expected number of trips.[9] From the result we subtract the product of the expected number of trips and the associated expected cost.

The final step is to calculate the change in consumer surplus. This is the sum of the consumer surplus for the each of the post rail system cases minus the sum of the pre rail system cases. Note that we examine two post rail system cases; immediately following its introduction, and five years after it has been introduced, when Hobart has "reacted" by creating four TODs. In principle, any number of "after" snapshots could be accommodated, to capture various different dynamic reactions to the infrastructure being put in place.

Conclusions

The approach outlined in this paper and illustrated with the Hobart case-study example was borne out of dissatisfaction with prevailing approaches used by practitioners in Australia for establishing benefits in cost benefit analyses of infrastructure in Australia, which do not follow best-practice in terms of the transport economics literature, and thus do not robustly estimate the true benefits of infrastructure service provision,

[9] That is, a series of rectangles under the demand curve. This is a common way of approximating the integral (and indeed, how it was discovered) and is used because MS Excel does not calculate integrals.

except under a very narrow set of circumstances. This is important in the context of infrastructure whose cost often runs into billions of dollars. It is surprising given that other jurisdictions make use of methods grounded in consumer surplus (see, for example, the UK TUBA model), and the academic literature provides both ample evidence of how to do this, and ample evidence of why it is appropriate.

To ground an approach of estimating benefits in consumer surplus, one first needs a demand curve. There are many ways in which demand curves can be derived, which might fit different circumstances, and availability of data. In this instance, given a lack of data about likely demand for a service which did not yet exist, we chose a relatively simple approach. We first looked at the total costs of travel for different modes, before and after the introduction of the new service. In these costs, we included cash costs, time costs and other costs such as externalities. We then stratified the overall population into several hundred different "representative consumers" based upon location and demographic characteristics, and we asked how each such consumer would choose amongst the different modes available for each trip they might make, if their aim was to minimise their overall costs of travel. We then changed the parameters of the cost functions slightly (to reflect uncertainty around their values), and asked the question again. We repeated this process several hundred thousand times, recording on a scatter plot for each mode, each time it was chosen and the quantity (in terms of distance) and price (in terms of total cost) that prevailed in each of these instances. Our demand curve for each mode was then fitted to this scatter plot via regression analysis, and the consumer surplus found through integral calculus. The model results, in terms of modal shares amongst different groups in the population, was then compared with actual travel survey data to validate model results.

We then repeated the process again, with the new transport service added into the set of modal choices. We took the sum of consumer surpluses from this "after snapshot" and subtracted from them the sum of consumer surpluses calculated in the snapshot before the new infrastructure service was added. We in fact repeated this process several times, to allow for different potential dynamic reactions in Hobart (such as the construction of transit oriented development around each station) to the rail system being assessed.

We do not claim that our approach is the only way in which demand curves can be derived. However, we do contend that it is important for infrastructure policymaking in Australia to move away from point estimates of demand and assessments of the consequences of these levels of demand, and towards estimation of benefits grounded in consumer surplus. We would also suggest that the approach outlined in this paper provides way of grounding benefit cost analysis in consumer surplus which makes use of much of the same information as is currently used by policymakers, making it easy to apply in an Australian context.

Bibliography

Baglino, A, Harrington, W, Lipman, A, Nelson, P & Safirova, E, 2007, "Transit in Washington, DC: Current benefits and optimal level of provision", Journal of Urban Economics, 62, 231–51.

Becker G 1981 *A Treatise on the Family.* Harvard University Press

Bureau of Infrastructure, Transport & Regional Economics (BITRE), 2009, *Australian Transport Statistics Yearbook,* available from www.btre.gov.au/info.aspx?ResourceId=710&NodeId=50.

Bureau of Infrastructure, Transport & Regional Economics. (BITRE), 2009, *Cost of Road Crashes Australia,* available from www.btre.gov.au/info.aspx?NodeId=26&ResourceId=47.

Department of Infrastructure, Energy and Resources (DIER), 2010, *Greater Hobart Household Travel Survey,* available from www.transport.tas.gov.au/miscellaneous/greater_hobart_household_travel_survey.

Department of Infrastructure, Energy and Resources (DIER), 2010, *Tasmanian Urban Passenger Transport Framework*, available from www.dier.tas.gov.au/passenger_transport/passengertransportframework.

Department of Infrastructure, Energy and Resources (DIER), 2011, *Congestion in Greater Hobart*, mimeo.

Department of Infrastructure, Energy and Resources, (DIER), 2009, *Review of Passenger Travel Demand Measures Greater Hobart*, available from www.dier.tas.gov.au/passenger_transport/hobart_passenger_transport_case_study

Douglas, N & Karpouzis, G, 2011, Tracking the Value of Rail Time Over Time Australasian Transport Research Forum 2011, 28 - 30 September 2011, Adelaide, available from www.patrec.org/atrf.aspx

Harford, JD, 2006, "Congestion, Pollution, and Benefit-to-Cost Ratios of US Public Transit Systems", *Transportation Research D*, 11, 45–58

Parry, IW, & Small, KA, 2009, "Should Urban Transit Subsidies Be Reduced?", *American Economic Review* , 99(3), 700-24.

Stanley, J., Hensher, D. A., Stenley, J., Currie, G., Greene, W. H., & Vella-Brodrick, D. (2010). *Social Exclusion and the Value of Moblity*, ITLS Working Paper 10-14, available from sydney.edu.au/business/itls/research/publications/working_papers/2010

Varian, HR, 2006, "Revealed Preference", in Szenberg, M, Ramrattan, L & Gottesman, AA (eds), *Samuelsonian Economics and the 21st Century*, New York: Oxford University Press, 99-115

Winston, C & Maheshri, V, 2007, "On the Social Desirability of Urban Rail Transit Systems", *Journal of Urban Economics,* 62, 362-82.