WORKING PAPER
ITS-WP-03-12

Using values of travel time savings for toll roads: Avoiding some common errors.

By

David A Hensher and Phil Goodwin*

January 2003

*ESRC Transport Studies Unit
University College London
London WC1E 6BT UK

ISSN 1440-3501

NUMBER: Working Paper ITS-WP-03-12

TITLE: Using values of travel time savings for toll roads: Avoiding some common errors

ABSTRACT:

There are many empirical studies on the estimation of values of travel time savings (VTTS), with varying degrees of rigour and relevance, mostly based on the observation that travellers are prepared to spend money to save time. These values are applied both to forecasting the effects of speed changes on behaviour, and also to estimation of the social benefit of such savings, in order to calculate value for money of spending public funds on transport investments. The sources of empirical information on such values are not always compatible with the models and software within which the results are used. In recent years, an increasingly important application has been to calculate the potential revenue from tolled roads, and networks with user charges, which offer high speeds at a higher price: here the important issue is not hypothetical willingness to pay, but the actual money which will be handed over. This changes the focus from hypothetical to *bankable* values of travel time savings. It is shown that some common practices risk substantial error in calculation, affecting the sharing of risk between public and private sectors. A particularly important case is where an average value is taken as representative of a skewed distribution of values – in these circumstances there will be a tendency to overestimate the revenue, and underestimate the traffic impact, of a charge, because for a given mean VTTS, there will be a smaller number of individuals who are prepared to pay the toll. To correct this bias, the main tasks are: establishing a relevant set of trip-purpose specific VTTS distributions and selecting a way of handling the distributions in patronage forecasting, growing VTTS through time, treating the VTTS of car passengers, and establishing an appropriate set of rules for converting disaggregated (or heterogeneous) components of travel time values into a single trip value appropriate to the project being evaluated. Other related problems of the use of values of time relate to the assumption that these values grow in proportion to income, and the extent to which they are confounded with other effects.

One troublesome feature is that most, and perhaps all, of the problems discussed tend to produce biases in the same direction, namely to risk overestimating revenue, in the short and long run. This produces a tendency to appraisal bias, which can distort the contractual confidence between partners. Overall, it is likely that current assumptions are underestimating the degree of toll-avoiding behaviour, and overestimating the financial viability of projects.

KEY WORDS: Tolls, Patronage Over-estimates, Travel Time Values.

AUTHORS: David A Hensher & Phil Goodwin

CONTACT:

Institute of Transport Studies (Sydney & Monash)
The Australian Key Centre in Transport Management, C37
The University of Sydney NSW 2006, Australia

Telephone: +61 9351 0071
Facsimile: +61 9351 0088
Email: itsinfo@its.usyd.edu.au
Internet: http://www.its.usyd.edu.au

DATE: January 2003

For a Special Issue of Transport Policy on the Valuation of Travel Time Savings

1. Introduction

Public agencies everywhere spend public funds on transport infrastructure in the hope of providing opportunities for faster, more efficient movement, and therefore the amount of resource worth spending to make a unit saving in travel time has always been an implicit or explicit issue in transport policy. As discussed elsewhere in this volume, in the past half century there has been a large number of studies using ingenious methods of discovering how much such time savings are worth to the travellers themselves, as judged by their willingness to spend extra money to achieve them. Over the same period, there has been an increasingly complex structure of models using these values to assess how many people will change their route of travel, or their mode, destination, trip-making etc, when faced by various different alternatives which are possessed of different relative advantage, especially as between fast expensive, and slow cheap, alternatives.

It seems to be frequently (perhaps always) the case that the models used for forecasts and appraisal have had some or many differences in theoretical base, assumptions, or algorithms of convenience as compared with the empirical studies of willingness to pay. We have, for example, seen many cases such as:

- traffic assignment models only concerned with change of route, using a generalised cost framework using values of travel time savings (VTTS) derived from empirical studies of choice of mode;
- apparent differences of VTTS calculated in a way which is confounded with differences in comfort, convenience, status, effort spent or stress, applied to choices where those attributes are quite different;
- VTTS calculated from stated preference methods which must logically be based on very short term (immediate) preference structures, applied to equilibrium models which implicitly deal with behavioural response which takes some years to evolve;
- relationships between VTTS and other influencing variables (notably income) assumed to develop over time in ways, which are inconsistent with other evidence on such relationships, eg direct demand aggregate estimations.

Such issues have been subjects of concern for each of the authors over some years (eg Hensher and Goodwin 1978, Hensher 2001a,b and Goodwin 1976, 1998). In the recent period, however, the authors have found themselves working in a new context where *willingness* to pay is not a hypothetical issue, but a financial imperative. What happens when notional charges are replaced by real ones? The two most important applications of this are (a) consideration of construction of a new road, probably by a public-private partnership where revenue from charges provides the reward to justify capital investment; and (b) application of charges to an existing road network for reasons of demand management, congestion relief, or reduction of environmental damage, and often with some form of dedicated use of the revenues for related transport improvements.

In these cases, incorrect use of values of travel time savings may cause serious distortion of investment priorities, and potentially financial stress serious enough to call the viability of a company, or the sustainability of a risk-sharing agreement, into question. An additional dimension is that any errors are likely to become apparent not in thirty years (by which time the issue will be confused and of minority interest) but within the first year or two of operation, with intense public and private interest.

The central question is: will the willingness to pay we have assumed in the model, be converted into cash in the bank? This question converts into a sensitive and growing list of implementation tasks that are necessary to satisfy the private sector ventures preparing bids to be short listed, and subsequently to win the right to enter into a contractual arrangement with government to build, own, operate and maintain infrastructure. Conversely, public agencies themselves need confidence that the risk sharing arrangements, based on market assumptions, will not be in danger of rapid collapse or embarrassing renegotiations, or public discontent about unexpected fortunes and accusations of monopoly profits. Trujillo et al (2002) provide a very useful overview of these issues in the context of strategically (ie intentionally!) over- and under-shooting travel demand.

This paper brings together a set of experiences on the processes and challenges faced by analysts charged with taking a set of VTTS and integrating them into a patronage forecasting framework that must past muster with construction companies and financial institutions who will use the patronage evidence to raise equity and debt – as well as affording confidence that public bodies are responsible custodians of public resources. The selection of themes has been driven primarily by what we have identified as the important, if not controversial, issues that often add hours and days of debate and work in the molding of a set of patronage forecasts that are acceptable to key stakeholders and which are defensible when reviewed by government.

We acknowledge that the issues discussed here are only part of a potentially much wider set, focussing first on those for which a most immediate solution seems both necessary and possible. Of these, one in particular is the focus of this paper – the issue of use of an average value to represent a distribution. Prima facie, this is a technical issue of textbook (and rather elementary) statistical theory. In practice, it is problematic, often controversial, and with substantial implications.

The paper is organised as follows. The main sections that follow focus on the following themes: the role of trip purpose-specific VTTS distributions and how information in the distributions is handled as an alternative to a simple average; the treatment of passengers travelling in cars (especially for toll road projects); how to grow VTTS through time with especial consideration of the escalation criteria; the concern of analysts using VTTS to calibrate patronage models; the extent to which quality bonuses are incorporated in heterogeneous measures of VTTS; and the appropriate way of establishing a single overall VTTS for each trip purpose when data provides disaggregated values for time components such as free flow and congestion-related travel time. The latter theme focuses particularly on whether one establishes the weighted average process across travel time components in respect of the time composition associated with the alternative an individual is switching *from* or the alternative they are switching *to*.

2. VTTS Distributions

Consider the case where a faster road is available for a moderately expensive toll, and a slower one is free. Then some drivers will choose one, and some will choose the other. This is partly because the actual size of the time saving to be made will vary according to where one joins the road, how many other people are making the same choice, etc, and partly because each driver will accord a different personal importance to the price and the time saved. The former issue can and should be taken into account by consideration of how many drivers are making a big time saving, and how many are making a small saving, rather than by assuming all travellers are faced with exactly the same attributes, and how to do so is not the focus of this paper. We will assume that the distribution of incidence of travel time savings and costs has been correctly calculated. Even after this is done, however, there remains an importance distribution of differences among the population.

We take it as axiomatic that for a population or subgroup of individuals or journeys, there will be a distribution of values of travel time savings. This arises directly from any attempt to state a theoretical basis for such values in terms of utility maximisation (or indeed any other theory), and in any case is so intuitively obvious hardly to need justification. We know that there are some occasions when time saving is important, and others not; some individuals who are under time pressure and others not; some travellers with plenty of money and others not. We know logically that such considerations will vary according to a huge range of potential differences in circumstances.

We also know empirically that every study produces slightly or substantially different results, and every study can easily find dimensions of segmentation of its own data base for which there are systematic, sensible variations in the values found (by, for example, income, employment status, journey purpose etc). By extension, when such studies find a different representative or average value for any particular segment, we can be confident that within that segment all the members are not perfectly homogeneous: there is always the logical possibility of a finer disaggregation that would show further differences, until the final atomistic stage of a single journey by a single individual in a single context on a single day.

For practical purposes, the most popular way of acknowledging such variation has been (within a specific trip purpose) to segment a sample based on some exogenous criteria such as income, trip length and time of day (especially peak and off-peak). This segmentation is achieved through estimating separate models for each segment or by interacting the travel time attribute (s) with the exogenous criteria (eg travel time*personal income). In both segmentation strategies a specific set of potential influences as explanations of the variation in VTTS are imposed.

At this stage we interrupt the argument to note that (as will become important below) in practice the selection of the number and dimensions of discrimination is not usually driven by questions of statistical diagnostics, research hypotheses and evidence. It is constrained by the specific properties of the forecasting and appraisal models within which the empirical values will be used. Even if – for example – gender were to emerge as a powerful statistical reason for VTTS to differ, that information would not be easy to use within a model where gender has no role, or an assessment in which there was no way of forecasting future travel behaviour by gender. The design of data collection often, and the usable analysis always, is constrained by the properties of the model to be used. Most important, all models known to us in common practical use will at some level of segmentation use the average value for that segment to represent the behaviour of the people or trips in it.

We now propose an alternative, more general approach, well justified by theory and evidence, without initially worrying about how easily it might be usable.

Consider that for whatever segmentation is justified by systematic difference, we then allow for further differences within the segments, by expressing the underlying parameters that represent preference weights for travel time and/or cost as random parameters (in contrast to point estimates) such that a distribution can be obtained for the preference weights for one or more attributes used in the derivation of VTTS, and hence for the VTTS itself. These distributions represent the preference heterogeneity of a sampled population or segment. One can even allow these random parameters to be some function of exogenous criteria as a way of establishing whether the segmentation criteria commonly used (such as personal income) do indeed systematically vary with the real behavioural variation in VTTS.

Then the proportion of a population, P, who will choose to pay a toll t is given by the proportion is value of the time saved is greater than t, ie

$$P_t = \int_t^{\infty} f(V)$$

The analyst, according to taste, convenience and internal evidence, will select among a number of appropriate analytical distributions (eg normal, lognormal, gamma, triangular, or a non-parametric set of bars), in order to find a satisfactory representation of the ‘true’ empirical distribution. The essential issue is then shown in the two diagrams of figure 1. (Attached at end)

The area to the right is the measure of the number of people whose value of time savings exceed the toll charged, who will therefore pay it. This is then the measure of revenue to be received by the charging agency. In the case of a symmetric distribution, eg normal, in general representing the distribution by its mean will be able to produce the correct revenue. In the case of a substantially skewed distribution (eg lognormal) the average will not be in the centre of the distribution, and in the case as drawn in Figure 1 (lower graph), there will be fewer people in the population actually ready to pay the toll. In this situation the mean is greater than the median.

Since the distribution, in most circumstances, will logically be bounded by zero, it will tend logically to be skewed in the direction shown. Thus the *general* case, we argue, is that representation of the distribution by an average is likely to give over optimistic projections of revenue (and consequentially underestimates of the extent of behavioural adaptation, hence impacts on traffic congestion etc).

This paper is not focussed on how we establish VTTS distributions (see, for example, Hensher and Greene (2003) for details and other papers in this special issue) but on what follows from the logical starting point. Empirically, the question is whether the effect is big or small. In circumstances where it is big enough to be important, the analytical question is what to do with the information from such generalised distributions. The main options are: (i) use the full distribution (ii) take a number of points on the distribution as representative of the distribution (iii) take areas of the distribution and convert to a single weighted average VTTS, ensuring that all areas sum to the total area and (iv) use the unweighted average or median[1]. The implications of choosing one or other of the options is profound, with the implications on patronage forecasts being greater as the distribution become more skewed. For example, if the distribution was skewed to the left, then we would have relatively more sampled individuals not prepared to pay the toll. Thus if we use the unweighted average we are likely to overestimate the number of toll payers.

2.1 Some practical experience

The following commentary is based on experience by the authors of discussions with local or national authorities, and funding agencies such as banks etc, of how such issues are discussed. Note that we have made the locations and agencies anonymous, and in some cases extended the logical implications – no specific criticism is made of any particular body or individual.

Empirically, there is growing evidence of a left skew in the distribution of VTTS. For example, in an Australian study carried out by Hensher (2002) in a toll vs free road setting, the results[2] for car commuters are summarised in table 2, and suggests that there are a disproportionately large number of individuals with relatively low VTTS who are not prepared to pay a toll to save travel time; in contrast the proportionately smaller number of individuals with a high VTTS are more than prepared to pay a toll to save time.

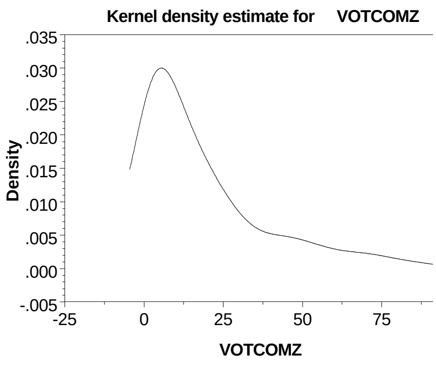

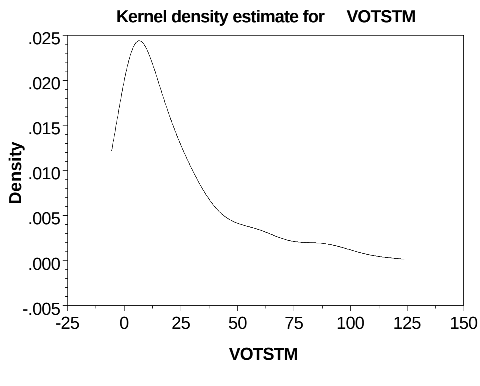

---

[1] If it is impractical to use the entire distribution in making policy, for whatever reason, one should use the median rather than the mean as their benchmark. This is however still only an approximation because what we really need is the value of the percentile that will be making the marginal decision.

[2] We also developed VTTS distributions for car non-commuter, light commercial vehicles and heavy trucks. The distributions were also skewed to the left.

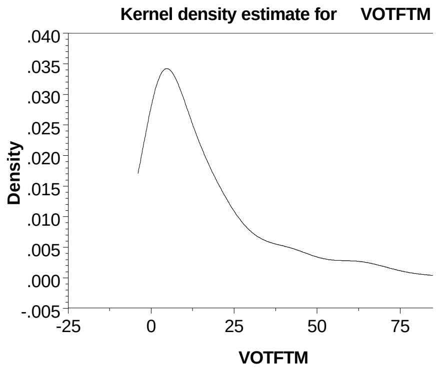

Figure 2 Distributions of VTTS for Car Commuters

Notes:

X-axis is the VTTS in $ per person hour. VOTCOMZ = VTTS for total time,

VOTSTM = VTTS for slowed-down time, VOTFTM = VTTS for free flow time.

"Density" means the % of experiment population. Separate graphs are preferred to account for the different ranges of VTTS.

Financial institutions have two interests in their negotiations with public agencies on a public-private partnership. First, there is an interest in the best and most reliable possible estimate of the expected revenue. Second, there is interest in figures which strengthen their bargaining position in relation to the case for the scheme to go ahead at all, and on what basis of risk apportionment.

Consider the case where there is a well-established convention, used by the public agency for many years, to represent the distribution of values of travel time savings by the average, partly for reasons of adequacy for purpose in previous applications, and partly because the models and consultants available find it convenient to do so. Then estimates made using the average, other things being equal, will tend to overestimate the revenue.

In this case, the financial agency has the choice to go along with the standard procedure, or to 'rock the boat' by suggesting using a distribution. The effect of doing so may well put the whole project at risk. So the perceived best interests of the agency are served by accepting the standard practice, which strengthens the case for the project, but (suspecting that it overestimates the revenue) finding a risk sharing agreement, explicit or implicit, which cushions them against the likely result.

Conversely, the public agency's perceived best interests are served by using the standard practice, since this will increase the probability of raising the funding, anticipating that the public benefits in terms (for example) of congestion and pollution relief will be higher than calculated, and seek to ensure that the risk will be wholly born by the funders.

The paradoxical case is that each will be better served by using the distribution themselves, for internal, confidential reasons, but using the average (or preferably the median) value for public discussion, and hoping that the other party believes it. But this is not a long-term solution, since it is almost bound to lead to later disputes, attempts to renegotiate, or collapse of confidence in such deals. There are signs that this can happen. The dilemma is obvious – will the financial advisers prefer to go with an overestimate to secure patronage and the contract (in a bid setting) knowing the likelihood (from previous contractual arrangements) that the risk can be transferred to government, or act as good corporate citizens and promote the more appropriate VTTS across the distribution.

In practice, this question is either ignored, or not expressed in this language (though accepting the underlying significance). The great majority of patronage studies around the world use simple averages for VTTS, so this provides an almost unquestioned benchmark as an always available fallback position, and a handy defensive (but not necessarily defensible) instrument.

From a practical perspective, acceptance of the distribution rather than the mean (or median) for VTTS may still involve some compromise. Practitioners are unlikely to import the full continuous distribution into their patronage forecasting models unless such models are driven by either synthetic households or sample enumeration (both with weights up to the population). The majority of patronage modelling styles use zone-to-zone averaging of service levels (including travel times), limiting the use of the full distribution. There is however practical merit in the compromise of establishing a manageable number of values of travel time savings points to *represent* the continuous distribution. When this happens it appears most common for analysts to take three or four points on the distribution and to impose some rule as to what percentage of the population each VTTS should be assigned to. While this might be acceptable if one can assume a normal distribution around that point without overlaps between the points, so that the implied average that the point represents is capturing the distribution in some sense, this is questionable where there exists skewness in the distribution. A preferred approach is to select areas under the VTTS distribution curve and calculate the weighted average. From a practical stand, this can be achieved by using a frequency distribution with agreed bandwidths (eg 50 cents) and simply establishing a frequency of incidence in each band and calculating a weighted average VTTS for that area. If one selects three VTTS then each area might represent 33.3% of the distribution.

3. Treatment of Passengers in Cars

Toll road patronage studies are interested in the vehicle and not the occupant since the toll is per vehicle. However the behavioural response of switching to the toll road is a decision of an individual, typically the driver of a car (or in the case of trucks and some light commercial vehicles, a mix of the driver and the person(s) in an organisation responsible for transport services). Patronage forecasts of toll road use typically use the VTTS for the car driver only. For cars it is often suggested that occupant's other than the driver might play a role in the establishment of an appropriate VTTS for the vehicle trip as a whole. A search of the published literature reveals a notable dearth of consideration of this issue. The exception is a recent study on VTTS in the UK (Accent and HCG 1999) that established car driver VTTS in the presence of passengers as well as a number of passenger values. However this study does not discuss the possibility of double counting or how such values might be treated in a car-based project setting. The authors indicate (on page 169) that they had a relatively small amount of passenger data which by implication precluded any serious assessment of the passengers VTTS and the role of the passenger in influencing the driver's VTTS.

Precisely, should the VTTS for a car trip be based on the VTTS of the driver or the VTTS of the travelling party? Given the absence of any empirical evidence, we draw on a series of intuitive arguments to establish a series of options, all of which might be the basis of future research. What we have is a framework within which individuals might act as independent or interdependent agents. Does the car passenger(s) influence the time-cost trade-off of the driver and to what extent is the decision to use the toll road influenced by the joint evaluation of the time-cost trade-off? Another way of stating this is: would the driver's time-cost trade off and hence VTTS be different in the presence and absence of passenger(s)? The UK study did find differences, but to what extent are they truly differences due to the occupancy of the vehicle or whether the occupancy is acting as a correlate for some other influences on preferences (for example a parent taking a son or daughter to university en route to work may not be relevant at all). Only by making the driver's VTTS a function of the occupancy might this be established. Simply segmenting by occupancy may not provide useful information. An interactive agency choice experiment that tests for the endogeneity of preferences amongst vehicle occupants (see Hensher 2001) would be interesting. Strictly, all that we need to know for VTTS is the values added by different types of passengers; however the broader research framework has value in establishing the role of each agent in determining the existence and even the mode chosen for a trip.

One appealing position is that the driver's decision to use the toll road is unrelated to the presence or absence of passengers in circumstances where the driver pays the toll. But what if the passengers are household members? Is there some sense of sharing the cost? In contrast to a free road the cost to the driver is unaffected by the presence or absence of passengers (the toll still applies as does the change in operating costs and time savings). In a modal choice context it might be different because the switch from public transport is a saving in fares per person and not per vehicle. The story is starting to get complicated; and so for the remainder of the discussion we will assume that all switchers are moving from another (free) route and remaining with the car. This is the usual setting for VTTS derivations in toll road studies.

Another appealing position is that the presence of passengers conditions the driver's time-cost trade-off. Imagine the situation where the driver talks a great deal with a passenger, which tends to pass the time quicker (and may make the slower free road more tolerable). Also there may be a feeling that the toll is yielding a benefit to more than one person and so, regardless of who is paying, there is a greater benefit to all occupants than to the driver. Thus the time-cost trade-off may involve a reduced marginal utility from a time saving but an increased marginal utility for the toll paid. These adjustments would tend to lead to an increase or decrease in VTTS depending on the relative change in the respective marginal utilities.

Another way of looking at the VTTS associated with the car passenger, assuming it has no impact on the car driver or that in any sample of drivers the incidence of passengers is somehow internalised in the driver VTTS (without knowing its contribution in the upwards or downwards direction), is to treat it their VTTS as a positive contribution to toll road time savings benefits. This is essentially the implicit outcome of most procedures adopted by toll road patronage forecasting studies.

The empirical study carried out by MVA et al (1987) made considerable attempts to separate values of time for passengers and drivers, or to establish values related to vehicle occupancy, and came to the conclusion that most of the evidence suggested that the values of time of passengers were discounted by drivers (who in effect were making the choices). Although passengers might indeed be valuing their own time savings, there seemed to be little evidence of a 'market' which allowed these fully to enter the choice process. One explanation was that car sharers might be a special group of the population with lower than average VTTS, and another was that application of economic willingness to pay ideas did not represent the sociology of car sharing. The authors speculated that 'If tolls were charged on the basis of occupancy rather than per vehicle, some more explicit trading might be done'. The study by Accent Marketing et al (1999) reported that their model results indicated that 'driver's value of time increases as the number of passengers…increases', but less than proportionately – ie their results were also consistent with the idea that passengers' values were discounted.

Our conclusion at this stage is that to assume values per vehicle would be proportional to occupancy, would give an overoptimistic assessment of revenue, and correspondingly underestimated assessments of toll-avoiding behaviour.

4. Growing VTTS over Time

The literature on how to treat changes in VTTS over time has existed for over 30 years. In the earlier years it was assumed (based on economic theory) that the mean VTTS was a function of the average gross personal income (or the average wage rate) and the percentage change over time in the average wage rate was used to adjust the mean VTTS. The adjustment used the exact same percentage for VTTS.

In recent years some research has been accumulating (primarily in the United Kingdom as part of a number of UK Value of Time Studies commissioned by the Department for Transport[3] on whether this assumption of proportionality is appropriate. There have been four strands of work.

Theoretically, it has been argued (notably in MVA et al, 1987) that there is no prior reason for expecting any particular reason for proportionality, or indeed any monotonic relationship. The reason for this is that the value of travel time savings is shorthand for the ratio of two distinct quantities – the marginal utility of the time, and the marginal utility of money. There is a strong expectation that the marginal utility of money decreases as (disposable) income increases, but the corresponding statement for time would be an

[3] The current name of the responsible ministry, previously MoT, DoE, DETR, DLTR.

expectation that the marginal utility of time savings decreases as the availability of (disposable) time increases. Both are confounded by changes in tastes, leisure activities, education, and opportunities or choice set open to people of different incomes. Overall, there probably is a reason to expect that willingness to pay for time savings increases with income, largely because of the money effect, but this does not translate into utility, and need not be proportional.

Empirically, there is now available a large set of studies of the ratio of utilities (though not there separate variation), in which the resulting values of time savings have been compared with income within the studies, or can be compared across studies. Both the MVA study mentioned, and a subsequent one by Accent et al (1999) came to a similar conclusion, using cross-sectional studies, namely that there was evidence of an increasing relationship, probably monotonic, but less than proportional. Their recommendations were to assume that values of time savings would grow over time, but at a rate less than the increase in income expected.

This comparison of 1985 and 1991 VTTS results, and a further 1995 study in which VTTS was formulated as a function of gross personal income, produced a series of income elasticities to approximate the impact on the average VTTS of overall income changes over time. For car drivers (and passengers) they recommended income elasticities of 0.45 (business travel), 0.65 (commuting) and 0.35 (other travel) – in other words, VTTS would grew roundly at about half the pace of income, for personal travel. There is less empirical support for similar effects for commercial vehicles which include light commercials and heavy vehicles), and the same study recommended the use of real GDP growth per capita as a proxy for growth in spending power and thus approximately in the long run for growth in the value of goods transported. The implication would be a secular growth in the weight afforded to goods travel in project evaluation over time, as compared with personal travel, for which there is little supporting evidence, and not a strong obvious rationale.

Similar evidence to support such elasticities is provided by Steer Davies Gleave in a recent study in Sydney (unpublished) where they plotted the relationship between mean VTTS (all in $US) and GDP per capita for 14 data points. The implied elasticity was reported as 0.5. Thus a 1% increase in GDP per capita produces a 0.5% increase in the mean VTTS (holding everything else constant).

A third approach has been developed by Wardman (1998a, b, 2001) applying formal meta-analysis techniques to around 1,000 data points drawn from UK studies for urban and interurban travel choices. His early results suggested a 0.5 elasticity of VTTS with respect to income, and the later study noted a wide range of different influences depending on methodology, differences between time series and cross section studies, etc. The range of results was large, and there was scope for considerable judgement in interpreting the results. Wardman concluded that a rather higher elasticity of 0.72 with a 95% confidence interval of ±43%.

A fourth strand of evidence relates to empirical work in a different tradition, the estimation of price elasticities mostly using econometric methods on aggregate data. There is an important connection with value of time savings studies in this connection, both because the ratio of price and travel time elasticities reflect the ratios of the marginal utilities when using a generalised cost approach, and also because the price elasticity is an alternative, and directly relevant, approach to estimating revenues from tolls.

A common practice, in UK and some other countries, has been to assume VTTS proportional to income, and price and time combined within generalised cost. Taken together, this implied that, other things being equal, the price elasticity will tend to be inversely proportional to income, with a strong expectation for price elasticity to decline over time. (The result being that it would be progressively easier over time to raise large revenues, but more difficult to influence traffic, from toll or other charging systems). However, a recent literature review and meta analysis of price elasticity results carried out by Hanly et al (2002) shows a puzzling result – there is no sign of any systematic decline in price elasticities in studies over the last 30 years, nor from re-analysis of specific data series divided by time period. This applied to a wide range of different price elasticities, eg fuel prices, vehicle prices, public transport fares, etc. Indeed, there were some signs of the elasticities increasing over time, though this was not well established. They argued that if the strong assumption for VTTS to increase with income, and price elasticity correspondingly to decline were well founded, than the effect should be big enough to be able to see some signs of it happening over the last thirty years, which was not the case.

In summary, the theory suggests that the utility of time savings is not necessarily related to income in any specific direction, but the willingness to pay for them should increase with income. Empirical value of time savings studies suggest the willingness to pay has increased over time, but less than proportionally, somewhere between a quarter and three quarters of the rate of income increase. Price elasticity studies do not show any sign of price elasticity declining over time in a way which would be expected if values of time increased with income.

Overall it seems reasonable to conclude that revenue calculations – especially where year-by-year[4] cash flow is of interest – will not be safely made by assuming values of travel time savings will grow proportionally to income. The VTTS benefit will grow less than this[5], and/or the resistance to price increases will decline less. Thus to assume that values of time will increase proportionally to income is essentially to assume that the market for time savings is strongly buoyant over time, and even if the early revenues are risky, in future years revenue growth will be strong: this assumption will tend to be overoptimistic on revenue, and potentially underestimate behavioural response, ie the same direction as the

[4] All cash-flow calculations rely heavily on year-by-year build-up or decline of the market, but this issue is the least well treated of any issue in travel demand forecasting which nearly always focus on end-states, not on paths over time. The main exception is price elasticity studies, which mostly show short run (one year) effects being rather less than half as great as long run (5-10 year) effects. Ignoring such demand effects will make a big whole in the early revenues, which may be practically more important than the other issues discussed in this paper.

[5] Ken Small raises a very important issue: if the elasticity is constant over long periods, travel time will become completely unimportant relative to other considerations in just another few decades of growth. Furthermore, if it has been constant for the last century, travel time must have been enormously important (relatively) a century ago, which does not appear to square with common observations.

distribution issues discussed above, and therefore tending to reinforce the problem rather than offset it.

5. Using VTTS to Calibrate Patronage Demand Models

It is not uncommon for transport consultants to use the value of travel time savings as a calibration parameter to reproduce base traffic levels on a link and a network. While this might have some appeal in situations where it might be argued that travel time is all that drives the patronage forecasts, there is the real risk that the behavioural meaning of VTTS is being threatened (if not destroyed) by this strategy. With additional influences on (route) choice excluded from a choice model and treated as part of the random error component of a utility expression for each alternative (see Hensher and Rose 2003), the components of time and cost run the risk of being confounded with other excluded attributes. Thus the focus should be on establishing the real behavioural valuation of time savings by careful specification of all the statistically significant influences on choice and then calibrating on the alternative-specific constant.

In practical terms, the importance of this reservation will depend on what other specific attributes are being confused with ‘pure’ time. There is limited strong evidence on this, but two aspects, which are of manifest and frequent interest, are those issues connected with ‘comfort’, and those connected with ‘reliability’. Toll roads, being faster, should also be more reliable, but might also be more comfortable though not necessarily so, and to the extent that both are behaviourally important, the equilibration process in route choice would tend to reduce the relative advantage anyway. So there might be a case for calibrating on a quality premium but there is a caveat – to what extent is this premium already captured in the decomposition of travel time used in the computation of the overall VTTS? We address this issue in the next section.

6. Quality Premiums and VTTS

Individuals choose toll roads for a number of reasons, in particular the travel time savings and the quality of such savings in respect of the changing mix of free flow and congested travel conditions. A toll road premium is often introduced to account for the differences in the quality of the traffic environment offered by a toll road in contrast to an alternative free road. This quality difference is additional to the amount of travel time saved in using a toll road.

The literature often makes mention of the quality difference as being primarily the avoidance of traffic congestion, such that a *given amount of travel time difference* involves a change in the mix of free flow and congestion-related time. Hensher (2002) in a Sydney study identified the ratio of slowed down time (SDT) to free flow time (FFT) for each travel segment and used this as a starting point for the determination of a quality bonus. These ratios are summarised in Table 1. Appropriate weighted average estimates of the ratios across all segments can be derived by the application of weights to represent each segment in the travel population.

Table 1. Ratio of Slowed Down to Free Flow time and Total Time

Segment	SDT/FFT	SDT/Total Time
Car Commuter	1.40	1.23
Car non-commuter	1.38	1.23
Light Commercial	1.00	1.00
Heavy Vehicle	1.77	1.37

If the application uses the mixture of free flow and slowed down travel times, appropriately weighted by their incidence in total travel time, we would argue that travel-time related trip quality as a 'bonus' or 'premium' is already taken into account.

If we assume that a toll road offers free flow under most circumstances (and certainly one might reasonably expect this is the perception of potential users), then a recognition of savings in travel time based on the mix of free flow and slowed down time on the current non-tolled routes would yield a substantial benefit in the form of elimination of slowed down time for a given amount of time saved.

If however there are additional benefits not related to travel time, then they would need to be added in. Examples might include the perceived safety of a modern toll road and the clear definition of where the road is going (which may be especially useful for travellers unfamiliar with travel in a specific geographical setting). The question is – what should this extra premium be? The only evidence we could find to account for differences in driving quality is derived from a study undertaken by Hensher and Sullivan (2003) in 2000 in New Zealand which looked at the additional benefits of upgrading a 2-lane road to 4 lanes (both with and without a median).

The New Zealand study found, after controlling for travel time differences (defined in terms of free flow and non-free flow time), that the driving benefit of the higher service quality road was on average 5.276 c/km for cars travel (all purposes) and 16.8 c/km for trucks (in Australian dollars based on an exchange rate of $0.813 New Zealand Dollars to an Australian dollar). To convert these unit rates to equivalent hourly measures of willingness to pay (WTP), we have to assume a difference in average speed between the two classes of service quality roads. For example, if the mean difference were 30 kph, the equivalent WTP per hour would be $1.58 for cars and $5.04 for trucks. Thus is we were to contrast these estimates with the VTTS for cars and trucks where the VTTS is derived from a given mix of free flow and slowed down time for the tolled and non-tolled routes, the toll road premium (or service quality bonus) is respectively 9.2% of car commuter mean VTTS and 20.3% of heavy vehicle mean VTTS.

Whether one would wish to take this premium, converted to equivalent VTTS, and use it to calibrate the patronage model (on the assumption that it is the only additional source of utility) is a matter of judgement.

7. Switching From or Switching To?

One of the most interesting practical issues in VTTS implementation is what we call the 'to or from' determination. The issue is not the VTTS per se but how one handles the decomposition of VTTS when the intent (as is so common) is to use a single overall (weighted) VTTS in the patronage model. It reflects the limitations of software as much as the concern about complexity in model application. The challenge is in deciding what weights to apply to represent the mix of travel times. Should we be using the mixture of say free flow and congested travel time associated with the toll road that someone might choose to use or the (free) route that currently exists? Let us explain the argument.

A typical study from which the empirical measures of VTTS are to be derived involves a sampled individual evaluating the levels of service offered by an existing (free) route and a proposed tolled route. A stated choice (SC) experiment is used in which an individual is asked to compare the levels of times and costs of their current route with the levels that might be offered on the new toll route. Assuming an unlabelled SC design in which the alternatives are nothing more than bundles of attributes, we can derive a generic VTTS. It is generic in that we are establishing a VTTS based on parameter estimates for each specific attribute that are the same across the alternatives. It makes no sense to treat the parameter estimates as alternative-specific, simply because the alternatives have no labelled meaning. The toll cost however would relate only to the tolled route simply because it does not exist on the free (or current) route; although if the SC study permitted the current route to be an existing tolled facility, we would have a generic toll parameter across all alternatives.

This is all fine and meaningful and essentially unambiguous. If a single VTTS were derived for *total* travel time it would also be unambiguous as to what VTTS should be used in applications. Increasingly, however, VTTS studies using SC designs disaggregate travel time into its constituent heterogeneous components such as free flow time and the additional travel time caused by a range of factors such as traffic congestion, random incidents etc. (see Hensher 2001a,b for examples). The composition of the overall travel time is likely to be different on the current route and the proposed new tolled route. If one were to take the separate VTTS for each time component (which are themselves generic VTTS) and implement this information as a weighted average VTTS for the entire trip, we need to establish the mix of type of travel time. The dilemma now faced is apparent. Do we use the time composition associated with the route someone might transfer from (in this example it is the free route) or the route someone might transfer to (in this case the tolled route)?

Our experience has been that the practitioners tend to take the mix from the route an individual is switching *from*, on the grounds that this is known, whereas it is an unknown on the *to* route until the switching patronage (for a given capacity) is predicted. Even though there is an argument that the actual switching traffic will also have an impact on the predicted time composition on the *from* route (after switchers have moved on), there appears to be a view that the best information we have on the weights attached to VTTS is from current experience which involves ex ante weights. Simply put, ex post weights are unknown and problematic.

If, as is expected, the toll road will offer a higher component of free flow time to other (congestion-related) time, and given that free flow VTTS is lower than non free-flow VTTS, we would expect that the weighted average VTTS would be more for the free route than the tolled route. This also seems problematic because it amounts to a loss of potential benefit simply by a quirk of calculation. It may be that the equilibrium models used in transport forecasting, in which generally all changes are non path-dependent, do not provide the most powerful framework for taking such inconsistencies forward: an explicit calculation of benefit with a time-dependent trajectory from a specific starting point should resolve the problem (and provide other useful advantages such as cash flows), but this is outside the scope of the paper.

8. An Assessment and Conclusions

Some convenient practices and simplifications, which have been used for many years and treated as good practice, are re-examined sharply when what is at stake is real cash, received, or not received, by real companies. This is not to say that the traditional questions of public investment for public benefit are any less ‘real’ – but they rarely have to stand up to rapid scrutiny: a road justified by a social cost-benefit analysis with a thirty year time frame does not lend itself to early retrospective assessment. If the forecasts are badly wrong, it will be many years before anybody notices, there will be a wide range of different explanations, and in any case there will be few of the original analysts around to take the blame. Only in cases of very extreme error (eg the assumption that road construction has no induced traffic effect) will the early experience be strong enough to be noticed with a sufficient degree of professional agreement.

In tolled and charged regimes, however, and especially where there are complex questions of risk-sharing, it is much more difficult to evade or delay consideration of the error. The Channel Tunnel connecting the UK and France made assumptions about the number of people prepared to pay a premium price for the time advantage, and these assumptions were found so badly wanting that its market estimates were undermined from the very first year – combined with cost over-run, many people lost a lot of money, and the whole exercise came close to bankruptcy. The project was saved in the short run, in effect, by writing off a substantial part of the capital debt, and in the longer term by a buoyant market in the context of European trade, which saved both it and its ferry competitors from a fight to the ruin of one or other.

Thus we do not claim that questions dealt with in this paper are not necessarily more important than the traditional arguments about assumptions and methods of transport policy appraisal, but we do claim that they are very much more immediate, and errors will be more apparent. . One might argue that the problems we have discussed were always weaknesses of established appraisal practice. The difference is that they are now very much more difficult to evade.

Acknowledgements

Ken Small read an earlier draft and made very useful comments to assist in the revision of the paper.

References

Accent Marketing and Research and Hague Consulting Group (1999) The Value of Travel Time on UK Roads – 1994, report prepared for the Department of Transport, London.

Armstrong, P., Garrido, R. and Ortuzar, J. de D (2001) Confidence intervals to bound the value of time, *Transportation Research* 37E, 143-161.

Goodwin, P.B. (1976) Human effort and the value of travel time, J Transport Economics and Policy, X (1) January

Goodwin, P.B. (1998) The end of equilibrium, in Theoretical foundations of travel choice modelling, Garling, T., Laitila, T. and Westin, K. (eds) Pergamon, Elsevier, Oxford, 103-132.

Gunn, H. F. (2000), An Introduction to the Valuation of Travel-Time Savings and Losses, in *Handbook of Transport Modelling*, Vol. 1., D. A. Hensher and K. J. Button, ed., Elsevier Science Ltd.

Hanly, P., Dargay, J. and Goodwin. P. (2002) Review of Income and Price Elasticities in the Demand for Road Traffic, Final Report to Department for Transport, ESRC Transport Studies Unit, University College London.

Hensher, D.A. (1977) *Valuation of Business Travel Time*, Pergamon Press, Oxford.

Hensher, D.A. (2001a) Measurement of the valuation of travel time savings, *Journal of Transport Economics and Policy* (Special Issue in Honour of Michael Beesley), 35, 71-98.

Hensher, D.A. (2001b) The valuation of commuter travel time savings for car drivers in New Zealand: evaluating alternative model specifications, *Transportation*, 28, 101-118.

Hensher, D.A. (2002) *Valuation of Travel Time Savings: New Evidence for Sydney*, Report prepared for Transfield, Sydney.

Hensher, D.A. and Goodwin, P.B. (1978) "The Transport Determinants of Travel Choices: An Overview" in Hensher, D.A. and Dalvi, M.Q. (Eds.), The Determinants of Travel Choices, Farnborough, England, Teakfield Saxon House Studies, March; 1-65.

Hensher, D.A. and Greene, W.H. (in press) The mixed logit model: the state of practice, T*ransportation*.

Hensher, D.A. and Sullivan, C. (2003) Willingness to Pay for Road Curviness and Road Type, Transportation Research D.

The MVA Consultancy, Institute of Transport Studies University of Leeds, & Transport Studies Unit University of Oxford (1997) *The Value of Travel Time Savings*, Policy Journals, Newbury.

Paag, H., Daly, A. and Rohr, C (2001) Predicting use of the Copenhagen harbour tunnel in Hensher, D.A. (ed). *Travel Behaviour Research: The Leading Edge,* Pergamon Press, Oxford, Chapter 36.

Trujillo, L., Quinet, E. and Estache, A. (2002) Dealing with demand forecasting games in transport privatisation, *Transport Policy*, 9(4), 325-334.

Wardman, M. (1998) The value of travel time: a review of British Evidence, *Journal of Transport Economics and Policy*, 32 (3), September 285-316.

Wardman, M. (1998b) A Review of British Evidence on the Valuations of Time and Service Quality. Working Paper 525, Institute for Transport Studies, University of Leeds.

Wardman, M. (2001) UK Value of Time, Working Paper 4, Institute for Transport Studies, University of Leeds.

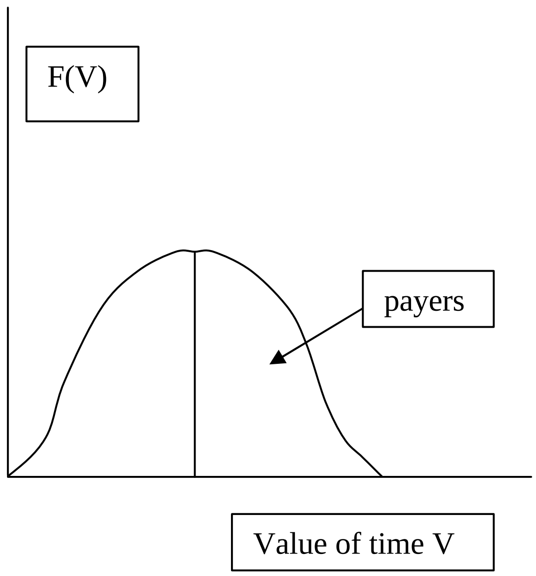

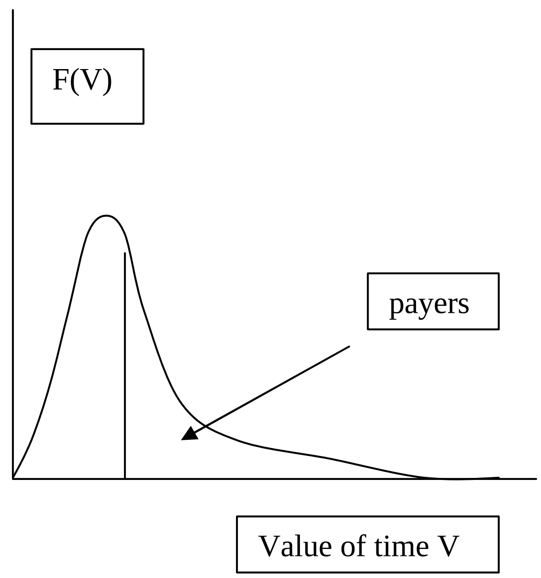
F(V)
payers
Value of time V