JOHN BATES SERVICES

1 Pickler's Hill. ABINGDON OX14 2BA Tel: Abingdon (01235)-526702

Stock-take of Travel Time Variability

produced for ITEA Division, Department for Transport

November 2010

Foreword

This is the Final Report relating to a request in September 2009 from ITEA Division. I was asked to "deliver a document of 10-15 pages that will describe the findings of the Department's research into modelling and appraising travel time variability. The document will serve as a summary for those unfamiliar with our research. It will compare and reference the different approaches to each other, and to the modelling and appraisal techniques recommended in the Department's WebTAG guidance."

In an Annex to the original specification, the Department provided a list of studies which it had carried out over the last 15 years. Many of the reports are available on its website http://www.dft.gov.uk/pgr/economics/rdg/ttv/ . In other cases, hard copies have been made available, and in some instances I was able to provide material from my personal sources.

As a by-product of the assignment to which this Report relates, I have separately prepared an Appendix entitled "A chronological review of Department's commissioned work on Reliability". This Appendix also contains a bibliography that simply lists, without attempting to assess relevance, all the documents that have been referenced in the various reports commissioned by the Department.

The Final Report has been discussed with ITEA Division, but I remain responsible for any errors of interpretation.

John Bates, 02 November 2010

Contents

1. INTRODUCTION

1.1 Understanding the Issue

An essential part of project and policy appraisal is an estimate of the associated travel time savings. These time savings can be multiplied by the "value of time" and converted to a monetary equivalent. Transport models are designed to work with travel time as a component of "generalised cost". Typically they "balance" two separate effects – **demand**, which predicts the amount of travel at a given cost, and **supply**, which predicts the (generalised) cost of travel at a given demand (relative to capacity), through congestion and crowding etc.

For the most part, however, transport models assume that, given the level of demand etc., travel time is known with certainty. As users of the system are well aware, in practice this is not so. Transport systems are subject to unreliability, and transport decisions cannot be made on the basis of certainty. Indeed, it has been suggested that the uncertainty of travel time can cause more disbenefit than travel time itself.

When discussing reliability in the transport context, the essential concept is the variability of travel time. For this reason, there is some tradition of only using the word "reliability" informally, and referring more accurately to "travel time variability" (TTV). This will be the general approach taken in this Report.

It is essential, at the outset, to be clear at what level we are considering TTV. For example, on the rail network, there may be a pattern of fast and slow services between a particular pair of stations. It would be unreasonable to consider that this contributes to TTV, even though an uninformed traveller might not be aware of the service pattern. Correspondingly, travel time on the highway network is affected by congestion, and the time taken to travel between two points is generally longer during peak periods than it is in the middle of the night.

This leads to the conclusion that TTV should not take account of variations in travel time that are inherently **predictable**, even if particular travellers are not adequately informed. We are essentially concerned with unexpected changes to the normal (expected) travel time, at a particular time of day[1].

Purely as an illustration, the diagram below shows the actual pattern of highway travel time on a particular stretch of road 1.9 Km in length in Stockholm, during the peak of the morning rush hour period (8.45 to 9.00 am), measured during weekdays in the months

[1] This is in line with paras 2.1.1-2.1.2 of the DfT's Guidance TAG unit 3.5.7. There remains an issue as to how to deal with uninformed travellers when the information is in principle available. The approach described here is seen as conservative but fair, with respect to the estimation of benefits. Nonetheless, with respect to rail passengers, a somewhat different approach is taken in the guidance (paras 4.1.1-4.1.2), where for those passengers unaware of **advertised** delay the increase in journey time can be treated as unpredictable variation.

June to October over three years. The figure plots the cumulative distribution – thus the vertical ("Y") axis shows the proportion of all journeys with a time **lower** than the corresponding time on the horizontal ("X") axis.

In the bottom left hand corner, the graph rises more or less vertically to a value of about 0.4, demonstrating that the lowest 40% of journeys have almost no variation in their travel time (the values on the X-axis range from about 1.5 minutes to 2.7 minutes). 50% of all journeys are below 3.7 minutes (this is known as the **median** value, indicated by the pink vertical line).

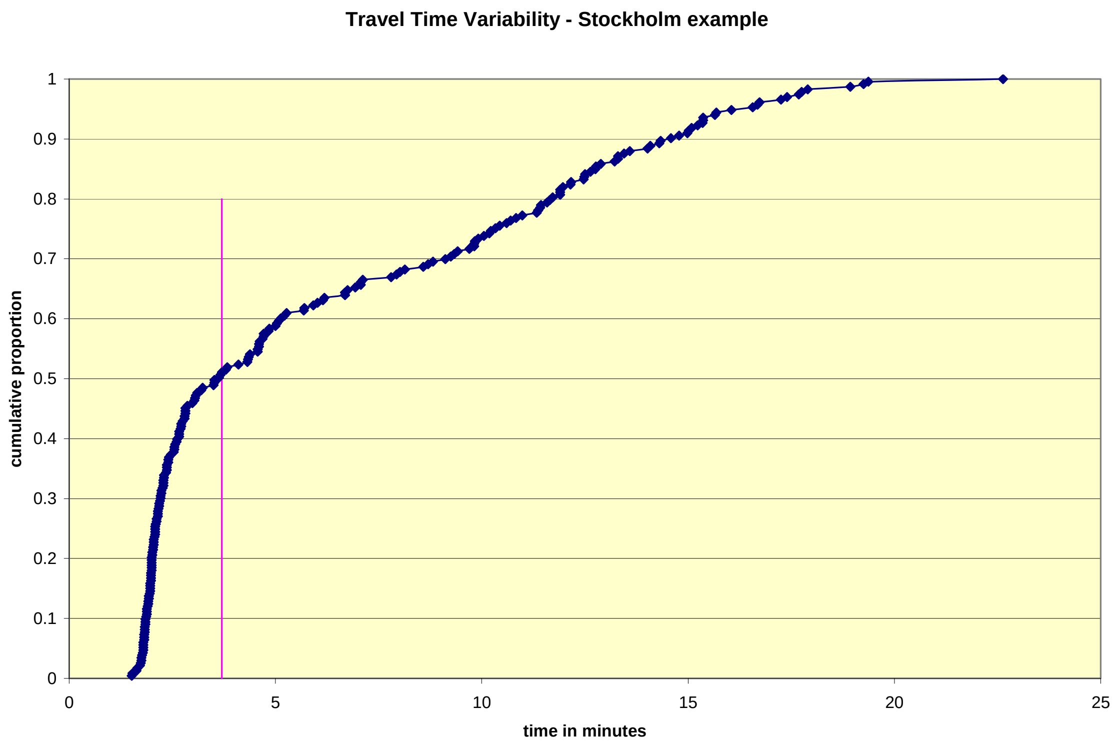

However, to the right of this median value, the graph bends significantly to the right, with ever longer times being observed, the longest time (within the set of observations) being 22.6 minutes. This means that, although half of the time, the time along this stretch of road might be no more than 3.7 minutes, there is an equal chance that it could be somewhere in the range from 3.7 to 22.6 minutes. Clearly, this represents an extremely uncertain travel time (though this particular road is not intended to be a representative case).

What is the cause of this variability? In general, there are three likely reasons. The first is **incidents**. These are randomly occurring events, typically leading to a loss in capacity: on the highway side, they could be traffic-related (eg vehicle breakdown, collision, or simply an illegally parked car or lorry), or they could be infrastructure-related (eg a burst water main, traffic light component failure, etc.). With rail, they could be signal failures, broken rails, faulty locomotives etc.. For all public transport, there are possibilities for crew-related incidents.

The other two reasons fall under the heading of what has been termed "**day-to-day variability**" (DtDV), generally more applicable to highway than to rail: these are random variations that occur between days that would otherwise be expected to be similar – in other words, after making allowance for predictable differences. Following the work of Arup (2002, 2003), there are separate effects on the supply and demand side. On the demand side we have random (unpredictable) variations in demand, which, inasfar as they lead to supply effects, will affect travel times: on the supply side, we have other random operating considerations (eg weather, random vehicle composition) which would lead to different travel times even if the overall level of demand was identical.

Arup expressed this as the following "equations":

> TTV = DtDV + Incident-related variability
> DtDV = Demand-related effects + Capacity-related effects

For rail, the impact of demand DtDV is not normally great, except perhaps for high frequency services such as London underground: this is principally because, in terms of travel time, variations in demand impact mainly on the time taken to board and alight, and with timetabled departures, sufficient slack time is normally built in. By contrast, with bus transport, variations in demand can have quite severe consequences, leading to bunching and general irregularity. For public transport we should note that variability can be an issue for waiting time as well as travel time *per se*.

Of course, both public transport modes can experience **crowding** as a result of demand variability, but in standard practice, when crowding **is** dealt with in models, it applies to a typical average day. Since crowding can occur even with a regular service, it is important to recognise that this is in principle a separate issue from TTV, even if the two may interact.

1.2 The importance of reliability

The Eddington report recognised reliability as a key issue for transport policy and appraisal. In Chapter 1.2, it was noted that: *In addition to the importance of costs and journey time, journey reliability also matters. This is supported by survey evidence and economic analysis. Travellers and business want to know not only whether a particular journey will take 30 minutes on average but whether it will take 30 minutes every day, rather than 30 minutes most days and 60 minutes once a week[…].*

Congestion on the network, failure to maintain assets and incidents can all affect the reliability of a journey. […]. For business, if materials or workers do not arrive on time, this can create bottlenecks and delays to production processes or result in the loss of perishable goods and service contracts. […] As a result, unreliability can also cost business in terms of contingency measures that need to be put in place.

If we wish to integrate reliability into our transport models, we need to know a) how changes in reliability will impact on demand, and b) how, for a given policy, the outturn reliability will be affected by the level of demand. These are the key supply and demand relationships noted earlier, and implicitly require TTV to be treated as an additional component of "generalised cost".

Over the last 20 years, the Department has commissioned a large number of studies relating to TTV, relating to both supply and demand. In 2001, it convened an important seminar[2], in which the research agenda was discussed. This was generally compatible with the long-term aim of integrating TTV into the demand and supply modelling.

Of course, if a particular cost component is not affected by the level of demand relative to capacity, then supply relationships for that component are not needed. However, in the case of TTV, there is no dispute about the fact that travel times are less reliable when the transport system is operating close to capacity, so that as a default we can expect that supply relationships **will** be necessary.

On the demand side, if TTV is worth valuing (as Eddington suggests), then it is likely to affect demand and, depending on the scale of the changes, the DfT's commitment to variable demand modelling (WebTAG units 3.10 and 3.12) suggests that these effects will need to be modelled as well.

There is, nonetheless, a residual concern that the research program is too ambitious. Current practice for accounting for TTV is typically by means of post-processing, ie **after** the modelling has been done – without including TTV. The feasibility of a full incorporation of TTV and demand and supply models is perhaps a question that can be considered in the light of reviewing the work that has been commissioned so far.

The rest of this Report reviews what has been achieved, and what still needs to be done. The aim is to present an overview, rather than a detailed critique, though the Appendix to this Report examines the individual studies in rather more detail. Before doing this, however, it is useful to present a brief overview of the theory which underlies the treatment of TTV.

1.3 The theoretical approach

[*Note: the essence of the approach is described in Annex A ("General Principles") of TAG Unit 3.5.7. A much more extensive review is available in Bates J.J., Polak J.W., Jones P. and Cook A.J., The Valuation of Reliability for Personal Travel, Transportation Research Vol 37E (nos 2-3) pp 191-230, 2001.However, the topic is of growing interest: more recent ideas can be found, for example, in Fosgerau M & Engelson L, The Value of Travel Time Variance, Transportation Research Part B, 2010*]

[2] more details are provided in Section 2.3

The inconvenience of TTV is primarily assumed to be due to the consequences in terms of arriving earlier or later than the ideal time. Note that this implies that if travellers are indifferent to when they arrive then TTV does not convey any “cost”. Of course, it may be argued that there are other aspects – perhaps “stress”-related – due to uncertainty in general.

If travel times are known reliably in advance, then it is assumed that travellers will choose their time of departure so as to arrive when they want to, bearing in mind other factors which might vary with the time of travel. Such models are described as “scheduling” models, and take into account the “cost” of being early or late. If travel times are uncertain, then travellers will typically build in additional “slack” time by departing earlier, particularly if they are more averse to late rather than early arrival. The same models can be used to predict this.

Under reasonable assumptions, it can be shown that proceeding in this way provides some justification for using the standard deviation of travel time as an appropriate indicator of reliability in cases where (as with highway travel) departure times can be more or less varied at will. The same theory can also be applied to the case of scheduled services, as with public transport: in this case, the standard deviation is less important, and there has been more of a focus on mean delay relative to the advertised service.

In deriving a valuation for a unit of standard deviation (analogous conceptually to the value of time for the case of a unit of mean travel time), researchers have developed the concept of the “reliability ratio”. This is a measure of the tradeoff between mean travel time and the standard deviation of travel time.

Thus the derivation based on scheduling theory has been a consistent approach in the theoretical work, because it provides an account of why TTV matters and relates this, with some approximation, to the standard deviation, as well as supplying a connection between the highway and public transport approaches. This is in agreement with the general approach of TAG unit 3.5.7.

While unit 3.5.7 also mentions traveller “stress”, the way in which it is proposed is merely a substitute in the absence of any ability to quantify the standard deviation.

2. STUDIES SPONSORED BY THE DEPARTMENT

2.1 Time-line of DfT research on TTV

Although not all the details of the earlier studies are available, the following table summarises the time-line which relates to the Department's work:

	Reliability Studies	**Demand Data**	**Supply Data**	**Background Events**
1989				NRTF published – capacity limitations acknowledged
1990	TSU Study on Driver Stress			
1991			DfT: PIPSY	London Congestion Charging Study begins
1992				
1993	Cranfield, SDG studies	Cranfield SP	SDG: North London	
1994				SACTRA Report on Induced Traffic
1995	MVA study starts			
1996			MVA: Edinburgh	
1997				
1998	John Fearon: INCIBEN			
1999	TAMTTV (Nov.)			Multi-modal studies begin
2000	Motts Follow-up research (assignment) Oscar Faber study (demand)	O.F. SP		GOMMMS issued
2001	INCA released DfT Seminar Arup Study begins			
2002				
2003	Conclusion of Arup study New Horizons Project (ITS)		M6 Database Arup: Leeds	
2004	TRL/Fearon: INCA			WebTAG launched
2005				
2006	ITS Rail study begins	ITS: Rail SP		
2007	Motts Study (INCA) Hyder Study (Urban) ITS/Polak/Bates (Multi-Modal)		HATRIS Database CJAMS /IT IS	
2008				TAG 3.5.7 Issued
2009				Stock-take

2.2 Early work

An early example of the DfT's interest in reliability is the oft-cited paper by Trevor Knight (1974), which discussed the need for a "safety margin" (ie, allowing additional time for the journey to deal with unexpected delays). The first UK Value of Time study (1980-85) proposed reliability as an important area of investigation, but a full treatment was excluded from the research programme. Probably the earliest study commissioned by DfT which began to address reliability was the 1990 desk study "Uncertainty and Driver Stress – a review" carried out by the Oxford Transport Studies Unit. Chapter 4 ("Uncertainty and the Value of Time") developed the demand response to TTV based on departure time choice modelling ("scheduling"). As discussed in Section 1.3, this is now considered the standard theoretical approach.

In 1993, as part of the London Congestion Charging Programme, two important pieces of work were commissioned by the (then) Department of Transport's London Assessment Division, separately on demand and supply:

> Black I G & Towriss J G (1993), Demand Effects of Travel Time Reliability
>
> Steer Davies Gleave [SDG] (1993), Highway Reliability Supply Effects

We briefly review these two studies below.

Demand Effects of Travel Time Reliability (Cranfield)

The first of these two studies, by Black & Towriss at Cranfield, involved a "stated preference" [SP] study trading travel time, money, and reliability. Reliability was presented as five possible journey times, and the model which was fitted to the data represented TTV by the standard deviation. This study established the key "reliability ratio" (ratio of the coefficient on the standard deviation to the coefficient on mean travel time) with an average value of 0.79.

This may be interpreted as implying that travellers would be happy to accept an increase of 0.79 units in mean travel time if at the same time they obtained a reduction of one unit in the standard deviation of travel time. Such a specification for generalised cost is often referred to as the "mean/variance" approach (even though it typically uses the standard deviation rather than the variance of the distribution).

Highway Reliability Supply Effects (SDG)

The objective of the second study was to predict the levels of variability associated with different traffic conditions. To this aim, the consultants carried out a series of timing measurements on different days along a circular route in NE London, and used this to develop empirical relationships, relating the standard deviation of the **journey** time to a function of the free-flow journey time and the actual journey time. A particular variable that was proposed, and has been used in much of the later work, is the so-called “congestion index” [CI], defined as the ratio of actual journey time to free-flow time.

The combined results of these two studies were implemented in the “APRIL” model used to test different congestion charging proposals. This appears to be the first attempt to include reliability explicitly in a transport modelling system.

Benefits of Reduced Travel Time Variability (MVA)

This was followed, two years later, by an exploratory project carried out by MVA. This confirmed the use of the standard deviation of travel time as an appropriate indicator of TTV, noting, in the process, that this could be justified as an approximation to the “scheduling” approach referred to earlier. An authoritative account of the issues regarding TTV for freight was also produced.

On the supply side, MVA rejected the approach of estimating variability directly on the **journey** (the approach taken by SDG), and emphasised the particular role of junctions in contributing to variability, especially in the urban context. Hence they developed a basic approach to calculate the combination of link and junction (LJ) contributions to the overall journey variability. For each “LJ”, the variances were added on the assumption of independence, but in summing LJs to a complete journey, they concluded that allowance should be made as a minimum for correlation between adjacent links, and proposed a range of values for the correlation. In addition, they developed an incident variability model which provided the foundation for the subsequently developed INCIBEN/INCA software.

It is worth recording that subsequent research has not made use of the “LJ” approach, but the reasons for this are unclear. It seems that in the wake of the MVA Report, DfT concluded that in terms of contributions to potential benefits, variability due to incidents on motorways was particularly important. Other possibilities from the remainder of the MVA work were not taken forward and the work itself appears not to have been widely disseminated

2.3 Discussion – Supply and Demand

While this early work addressed both supply and demand issues, subsequent work has concentrated on one or other of these two aspects. In what follows, we will therefore deal with them separately. In fact, the majority of the work commissioned by DfT has been on the supply side. This is in line with the research agenda set out at a seminar on 2nd July 2001, where the Department’s framework envisaged that ultimately “the generalised cost used in supply and demand modelling would take account of TTV just as it currently takes account of journey times, costs etc”.

This implies that these two elements - demand modelling and the network's TTV performance - need to be handled together, as shown below, to ensure the two are in equilibrium. It implies that the third (appraisal) element of this process should be based on the origin to destination flows and TTV outputs of the combined demand and supply modelling process.

The Figure below (presented at the July 2001 seminar) shows the three strands of knowledge that DTLR believed were needed to achieve this:

strand A - How do people perceive TTV and how do they take it into account in deciding where, when and how to travel?

strand B - How does the network's TTV performance depend on the level of traffic using it, and also how does it depend on the network itself - the standard of road, road geometry, and so on?

strand C – How much are people willing to pay for improvements in TTV?

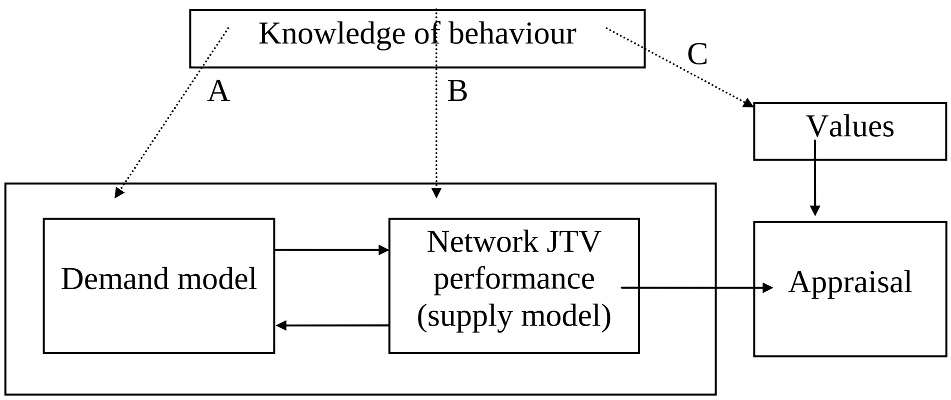

The DTLR Framework for Modelling Journey Time Variability

The general approach is also in agreement with Volume 2 of the Guidance on the Methodology for the Multi-Modal Studies (GOMMMS: DETR, 2000), which implied that some measure of TTV, such as the standard deviation of the journey time distribution, could be included in the utility (or "generalised cost") of travel alternatives. Paragraph 6.3.8 made the supply-side issues clear: "In order to use these results it is necessary to be able to specify how a transport intervention will affect public transport lateness and/or private vehicle travel time variability." Although this is silent on the question of "link vs journey", it is generally accepted that TTV is primarily an issue at the **journey** level (hence the development of the scheduling model relating to time constraints at the origin and destination).

2.4 Demand Work

As noted, both the Cranfield and MVA work took the standard deviation as the indicator of TTV. However, as noted in Section 1.3, the preferred theoretical approach is based on departure time choice (scheduling): in essence, because of the uncertainty of travel time, travellers will tend to leave earlier than they would if travel time was certain, in order to avoid the worst consequences of being late. The inconvenience of this earlier departure can be associated with the disbenefit of TTV.

As noted in Section 1.3, under certain conditions the “scheduling” disbenefit is equal to a factor multiplying the standard deviation of travel time. Strictly speaking, these conditions require that the factor depends on the shape of the travel time distribution, ie on the range of possible outcomes relative to the “normal” travel time. Hence, if this range is significantly affected by the scheme being appraised[3], using a constant factor would not be appropriate. Nevertheless, it appears that in most circumstances, using a constant factor is an acceptable approximation, though further confirmation of this would be useful.

While other studies have been carried out elsewhere on demand aspects of reliability (for a recent review, see the Department-sponsored overview project Multimodal Travel Time Variability by ITS *et al* (2009), which also provides a careful theoretical discussion of the main two approaches), only two further demand projects have been carried out under the auspices of DfT. The first was a project by Oscar Faber (later FaberMaunsell, now AECOM) entitled “Journey Time Variability – Stated Preference Survey”, which was in fact commissioned by the Highways Agency (2001-03). The second, more recent project, was carried out by ITS Leeds, and titled “Revealed Preference Study to Assess Impact of Reliability on Passenger Rail Demand” (2006-09)

Highways Agency work on Journey Time Variability (Oscar Faber)

The Oscar Faber [OF] study used qualitative research to investigate what TTV means to people and attempted to measure the value that people place on TTV using SP data. The research concentrated on the “scheduling” implications of TTV (arriving early or late) and was confined to the journey to work. An important aim was to allow for both a scheduling and a mean/variance model to be estimated and compared. Because the hypothesised equivalence between the two depends on respondents being able to optimise their departure times, the experiment contained a unique graphical display of the consequences of re-scheduling. Unfortunately, after careful examination of the results, it had to be concluded that the optimisation process, despite commendable efforts at presentation, was too difficult for respondents to carry out, and the data could not therefore be considered reliable. It seems unlikely that that further work along these lines could be justified.

Study on Rail Reliability (ITS, Leeds)

This was an attempt to estimate the effects of changes in journey time reliability on passenger rail demand. The main work was in two parts – a time-series analysis using ticketing data, and an SP exercise.

---

[3] for example, a scheme might be designed to remove the incidence of very long delays, while increasing the incidence of small delays

Considerable problems were experienced in marshalling the appropriate data for the time-series analysis, and the final data set was substantially smaller than had been intended at the outset. Nevertheless, some results of interest were obtained. The report noted that the elasticities for the three reliability variables tested (average lateness minutes [ALM], average performance minutes [APM] and public performance measure [PPM]), while of the correct sign[4] and significant statistically, were all quite low. The estimated values, based on the "static model" reported in Table 7.1 of Deliverable 4, are as follows:

	Non-season tickets		Season tickets
	Full fare	reduced fare	
ALM	**-0.03**	**-0.04**	**-0.06**
APM	**-0.03**	**-0.03**	***-0.04***
PPM	**0.09**	**0.25**	**0.27**

It should be noted, however, that this in itself is not unreasonable, since elasticities are calculated in terms of **proportionate** change from an existing position, and the mean base value of average lateness – for example - is likely to be small: say, 1 minute. It is not clear whether or not these results are compatible with the current PDFH-based approach, which makes use of "generalised journey time" [GJT] elasticities.

Some controversy surrounds the SP analysis, which presented the following variables:

fare
scheduled journey time
a set of five possible journeys

The five journeys within the set had different amount of delay on a) boarding and b) arriving [NB while it was possible to arrive before the scheduled arrival time, it was not possible to depart before the scheduled departure time].

The recommended model included the variables Fare, Scheduled journey time, probability of arriving late, standard deviation of extra in-vehicle time, lateness on boarding, and mean positive delay on arrival.

It was considered of some interest that the disutility for lateness on boarding was **higher** than the disutility for mean positive delay on arrival. However, subsequent work by Bates demonstrated that although the first four variables given above produced robust results under a number of specifications, there were a number of alternative specifications for additional variables which gave more or less the same fit. This suggests that there is too much correlation between the candidate variables to allow reliable discrimination between alternative models.

[4] the elasticities for ALM and APM are negative and the elasticity for PPM is positive

The study concluded that there was a case for reviewing current PDFH guidance, which recommends a lateness multiplier of 3 across **all** segments. Note that while the PDFH guidance appears to be on a rather different basis, it can be reconciled with the highway results via the scheduling theory, as briefly described in Section 1.3.

2.5 Supply Work

The Department has accepted that the supply issues are even less tractable than the demand issues, and has accordingly devoted a greater part of the research resources to aspects of supply, almost entirely in the highway sector. In dealing with highway networks, an important question arises: whether route choice is affected by TTV at the link level? Almost all the research commissioned by the Department has assumed that route choice is **not** affected (this was also the assumption in the MVA work).

The exception is the Mott Macdonald 2000-2003 project (with assistance from Imperial College and David Watling of ITS Leeds) entitled "Travel Time Variability: Follow-On Research". While this was a thorough theoretical piece of work, it is difficult to conclude that it represent a practical way forward for a standard problem. Having a definition of TTV (the standard deviation) that is not additive across links causes major problems for network modelling. Three potential methods were tested on a simple network and a somewhat more complex CONTRAM network for York. Only one (path enumeration, requiring a pre-defined set of routes) could be relied on to give the correct results. The later (multi-modal) study by ITS *et al* also addressed this issue, and set out a clear theoretical account: nonetheless, it concluded that "there remain substantial challenges, [though] there are a number of current efforts which are promising."

Although further progress could probably be made, it would be wise, in the first place, to assemble evidence that TTV at the link level does indeed influence route choice. In the meantime, we are in the position of not knowing how important an effect TTV has on route choice, while recognising that, if it is important, modelling it will be very difficult. It therefore seems necessary, at least in the short term, to continue to **assume** that the impact on routeing is insignificant.

For the most part, the two aspects of incidents and DtDV have been separately researched, though the need to bring them together within a consistent framework has been acknowledged throughout. Since initial priority was given to incidents on motorways, we will discuss incidents first. It may be noted, however, that once we move from grade-separated roads with limited junctions, distinguishing between incidents and other effects on TTV becomes progressively more difficult. In particular, in urban areas it is difficult to take separate account of incidents (and of course motorists are not necessarily aware of whether variability is being caused by incidents or excess demand).

Incidents

Original work by John Fearon

At the end of 1997, John Fearon Consultancy (JFC) was commissioned to review MVA's working assumptions and to develop and implement an initial model. The ensuing model, named INCIBEN, followed the general MVA methodology on incidents, based on queuing theory, but re-worked the formulae, and made use of an actual data-base of motorway incidents.

Essentially, the model deals with the random loss of capacity due to incidents, and implements a deterministic queuing model for the period during which demand exceeds the reduced capacity. It is able to calculate the benefits associated with a policy which could affect both the level and type of incidents, as well as capacity: benefits arise both from changes in the delay caused by incidents and from changes in TTV (which includes some allowance for DtDV). The model assumes that, in the absence of incidents, demand will not exceed capacity: thus it does not allow for "transient excess demand" (as will be discussed later).

The model operates on a link basis. Note that with incidents the issue of link correlations is not really an issue (on the reasonable assumption that very few vehicles will encounter more than one incident): incident rates are expressed on a per Km basis, so that the overall probability of an incident occurring on a sequence of links is additive.

Subsequently, Mott Macdonald were commissioned to re-package the INCIBEN program in a more user-friendly form, and it was re-issued as INCA, though no substantive changes were made in the basic methodology, apart from an algorithm (DIVA) to take into account the assignment effects of diversion.

Further work on Incidents (TRL with John Fearon, Mott Macdonald)

In a later contract (TRL with John Fearon, 2004), the incidents database was updated, and a validation of the INCA queuing theory against actual incidents was carried out. This concluded that the theory should be retained, but with updated parameters. It expanded the list of classified incidents, and produced a new set of incident rates for various categories of road (though it is noted that they apply only to **grade-separated** roads). It also recommended new lane capacity factors and proposed that during the decline phase (after removal of the incident) a new lane capacity factor ("DLCF") should be implemented: in other words, it is assumed that full capacity is **not** restored until all associated queues have cleared. Based on an analysis of the sample of incidents and COBA-based estimates of lane capacity, an average value of 0.81 (with an associated standard deviation of 0.11) was recommended.

In addition, a new diversion procedure was developed, as a substitute for the DIVA routine developed by Mott Macdonald. As a result, DIVA has been replaced by the method set out in the 2004 TRL report on "Updating and validating parameters for incident appraisal model INCA". Appendix A.6 of the INCA 4.1 manual provides a summary.

More recently, Mott Macdonald were asked to produce a specification and develop algorithms and methods that would enable INCA to accommodate analysis for single carriageways, as well as exploring the availability of data for creating an incidents data-base for single carriageways.

In fact, a proposal for single carriageways had been made at the time of the development of INCIBEN, as a modification to the methodology for predicting the mean and variance of delays due to incidents on dual carriageways. The general approach to the calculation of delays was thought to be appropriate, but the definition of capacity needed to be changed (to allow for the impact of traffic in two directions. Motts re-assessed this, and proposed a minor modification, which was more consistent with COBA assumptions.

The conclusion therefore was that only minor modifications to the algorithm would be required for INCA to work with single-carriageway roads, but that further effort would be required to provide appropriate data, both for the additional items needed for the algorithm and for the incidents database. For the moment, given the lack of data, the program remains restricted to dual carriageways.

Motts were also asked to review and improve the assumptions within INCA relating to DtDV. We assess these separately in the following section.

It may be concluded that the principles under which INCA operates are clear and generally accepted (within the assumption about capacity referred to earlier). However, the INCA manual specifically notes “Typical schemes for which the INCA methodology is appropriate range from motorway traffic management measures to motorway widening schemes which change incident-related delays and travel time variability.” INCA is essentially a tool dealing with a limited sequence of grade-separated links, with a correspondingly limited set of demand flows. Only 63 links (including “feeder links”) and only 63 movements (which need to be pre-assigned to links) are catered for. Hence, it would not be trivial to generalise the methodology to more complex networks.

Day-to-Day Variability [DtDV]

Following the seminar on 2nd July 2001, a contract was awarded to a Consortium led by Arup, with John Bates, John Fearon and Ian Black. Although the remit of the study did cover work on Incidents and on demand issues, the focus was on DtDV, and for motorways in particular.

In relation to demand-related variability, it was observed that, as long as demand is **always** below capacity, the standard deviation of journey time rises more or less linearly with the mean. However, there are other cases where, even if **average** demand is below capacity, there are short periods when capacity is exceeded ("transient excess demand") on an essentially random basis, leading to queueing. Thus to deal with these cases a dynamic analysis is required, which allows for lagged effects of demand in previous "time-slices" whenever queues arise.

Current transport modelling is, in the main, time-aggregate. As long as demand does not approach capacity, this is unlikely to lead to significant problems. However, it is not clear whether the more complex situation of transient excess demand can be represented by a time-aggregate methodology, even as an approximation. To analyse TTV at a particular point in time it would seem essential to have information about the temporal profile of demand, at least from the first occurrence of transient excess demand, and to make allowance for serial correlation, which destroys the monotonic relationship between the mean journey time and the standard deviation, giving rise to the observed "looping" found in the analysis of ANPR data (ITS, 2001), (TRL, 2001).

In principle, the analysis of queues when demand temporarily exceeds capacity is similar to that of incidents. With incidents, queueing results from the **reduction** in capacity associated with the incident, demand staying constant, while with transient excess demand the capacity is not affected, but there is a temporary spike in demand. However, the impact of an incident is different from the excess demand case in that the incident occurs at a random time (and place).

On the basis of the initial theoretical work, Arup proposed some further research required to transform the simplified models into practical tools. Making progress in the more complex case of transient excess demand was hampered by the fact that most existing network tools do not adequately represent the profile of travel time, and, more importantly, by the general lack of understanding of traffic phenomena under conditions of flow breakdown. In addition there are merging and weaving issues particularly where motorways merge, and significant diverge problems, such as when capacity problems on the slip road lead to tailback on the main carriageway, or even back to an upstream merge.

Given the importance attached to a methodology for predicting TTV on motorways which experience transient excess demand, it was apparent that only very limited progress could be made with the analysis of **existing** datasets (as had been the intention of the Brief), so that there was an urgent need for new data collection. Noting that the aims of the study were limited to modelling a small number of successive motorway links, rather than the full complexity of a general network, Arup proposed collecting a year's data for an 11 Km section of the M6, comprising, *inter alia*:

- journey time (ANPR) data;
- volume and speed (MIDAS) data;
- incident data; and
- meteorological data.

If the project were successful in dealing with TTV in the context of sections of motorway, the methodology and data collection could be successively extended to other road types in order of increasing difficulty:

- inter-urban grade separated all purpose dual carriageways;
- inter-urban non-grade separated all purpose roads and junctions; and
- urban non-grade separated roads and junctions.

However, as a result of the timescale involved in the data collection, the priorities for the remainder of the study were altered. It was accepted that it would not be possible to analyse the full dataset within the project. In fact, no subsequent steps were taken to analyse it (and it was not until 2009 that, under the auspices of the Highways Agency, all the data was examined for the first time[5]). Prototype models were developed based on the analysis of the first month's data only.

As a change of focus, Arup were asked to develop a simple method of modelling TTV on urban roads, involving a re-analysis of the SDG data, together with the results of a comparable data collection process in Leeds. A relationship was developed which had generally appealing properties, whereby the ratio of the standard deviation of journey time to the mean ("coefficient of variation") increases with the "congestion index" (the ratio of actual journey time to minimum or "free-flow" time), and decreases with journey length. While the same functional form seemed to fit the data well, the estimated parameters revealed some differences between London and Leeds, which require interpretation. Nevertheless, as a short term measure[6], an appropriately averaged version of the model was proposed to give estimates of TTV in urban conditions.

Two important remarks should be made. Firstly, the SDG and Leeds data would have included TTV from **all** sources, ie including incidents (though in practice the measurements would have been aborted if **significant** incidents had occurred). Secondly, as had been the case with the 1993 SDG study, the proposed relationship was for the standard deviation of time **for the whole journey** (though the relationship was, in fact, estimated on "chains" of successive links).

Inter-urban work on TTV (Mott Macdonald)

More recently, as part of the work on INCA discussed above, Mott Macdonald were also asked to make improvement to the DtDV calculation in INCA, by estimating travel time variability as a function of traffic flows and journey time on motorways, inter-urban roads and rural or inter-urban single carriageways. The scope of the contract was later extended to include managed motorways, as an additional road type.

This work made use of the journey time data in HATRIS for a selection of 18 routes (in some cases consisting of only one link) covering a variety of road types. Because DtDV excludes the effect of incidents, it was necessary to remove records which could have

---

[5] Although this work was not commissioned by DfT, a brief summary is given at the end of this Chapter
[6] The results have subsequently been updated in the Hyder urban work – see below

been affected. In practice this was merely done ("given the time constraints") by using the HATRIS incident "flag" based on an outlier detector method.

In contrast to the urban supply analysis reported earlier, the Motts HATRIS analysis was done entirely at the **link** level. It noted: "When combining these to calculate the standard deviation of total OD travel times we should really take account of correlations between link travel times. However, we have not yet considered them as part of the work described in this note."

Previously used forms of urban relationship between CV and CI (eg as estimated by Arup and by Eliasson (2007) in Sweden) were tested, but they were rejected as giving poor fits. The generally preferred function form was one in which the link standard deviation per kilometer was modelled as a cubic polynomial function[7] of the link mean travel time per kilometer. Separate models were estimated for different road types. With the exception of the single carriageway results (where incident rates are not yet available), these relationships for DtDV have been incorporated in INCA.

Overall, the main result of the Motts work is an improvement in the DtDV calculations within INCA (together with the prospect of extending INCA to work with single carriageways). The **quality** of the DtDV relationships is more difficult to assess, and they are not generally compatible with the kind of **urban** models that had been previously developed. Furthermore, the use of such relationships to calculate DtDV for O-D movements remains unaddressed.

TTV in urban areas (Hyder)

More or less in parallel, a team managed by Hyder were commissioned to extend the supply models for TTV in Urban Areas. In contrast to the earlier urban work, this study, like the Motts work, made use of electronic data. The HATRIS data, while suggested in the Brief as a possible source, was rejected, as the majority of this dataset relates to **interurban** journeys. Instead, it was decided that the ITIS/CJAMS data was the most suitable, and a total of 34 routes were selected covering the 10 largest urban areas in England.

In line with previous work in the urban context it was accepted that it was not really feasible to distinguish between DtDV and incidents, though significant outliers were excluded. Because the ITIS database includes vehicle type, it was possible to confine the analysis to cars and LGVs (excluding breakdown vehicles).

The analysis began by fitting straightforward elasticity relationships between CV and CI for each **link**, after which various selections of "part routes" were made, and an overall average journey-based relationship obtained with the same form as the Arup relationships for Leeds and London, involving CI and distance.

[7] note: this is similar in form to the fuel consumption formula used in TAG unit 3.5.6

Given the substantial variation found in the link relationships, Hyder was commissioned, as an additional task, to collect link characteristic data for the selected routes, with a view to explaining the differences. However, despite significant effort using (this time) a linear relationship between standard deviation and mean delay, the results were generally negative, and it was acknowledged that the best model "leaves a large variation to be explained."

In fact, in terms of the predicted standard deviation the outcome of the "distance-based" model (O-D/route level) and the combination of link-based variances with an assumed adjacent link correlation is very similar. Since the "distance-based" model is easier to implement, this remains the recommended approach (as set out in para 3.3.2 of TAG Unit 3.5.7).

Recent M6 Work by Hyder *et al*

As noted earlier, the database collected in the Arup study remained unused until at the end of 2008, following discussions with the Highways Agency, a proposal was made by Hyder in collaboration with John Fearon, John Bates and Ian Black, with the aim of acting on the Arup recommendation for the analysis of the data (cited earlier), and ultimately developing a practical model, suitable for use in scheme evaluation and policy analysis, of TTV on motorways, that specifically allows for the phenomenon of transient excess demand, or flow break-down. This effectively re-united the original team for the Arup study.

Although this study does not strictly fall under the terms of reference of the current report, and was carried out for the Highways Agency rather than by DfT, its conclusions are of sufficient interest to justify a brief summary here. The following is based on the Executive Summary of the Final Report.

The study identified "flow breakdown" (the sudden change in the flow of traffic on motorways from smooth and fast to stop-go) and in particular in the form of "Standing waves" as the main cause of excess travel time and TTV on the congested section of the M6 that was studied. This is a stochastic phenomenon which is likely to occur in periods where demand is relatively high.

Delay and TTV caused by 'standing wave' flow breakdown in the case of M6 were found to be due more to supply than demand variability. The (random) time at which flow breakdown occurs is critical in determining the level of delay and TTV, since the impact depends strongly on when it occurs in relation to the profile of demand over the day.

The study was able to reflect the observed pattern of delay and TTV by means of an "Aggregate Model" which adopts a simulation approach at a relatively aggregate level, working in five-minute intervals. The model has four main components:

- the mean profile of demand for traffic entering the modelled network, together with its variation from day to day;

- a simple dynamic assignment model, in which the network is represented by a relatively small number of links;
- a “Breakdown Probability Function”, which forecasts the probability of flow breakdown occurring at a given time; and
- a “Queue Discharge Function”, which forecasts the rate at which the queue discharges once flow breakdown has occurred.

The model offers the potential for the prediction of both TTV and the delays caused by flow breakdown – something which conventional traffic models are not able to do. The Report makes a number of recommendations for further development and application.

A particularly important conclusion for modelling is the emphasis on the importance of the profile of demand. While the standard modelling for the appraisal of travel time benefits usually forecasts impacts for fixed periods, e.g. AM peak, PM peak, inter-peak, this approach is problematic for modelling the impacts of 'standing wave' flow breakdown: a) because travel time and TTV effects in any particular time period are often dependent on supply and demand in previous periods; and b) because the actual shape of the profile can be critical to the outcome.

3. ASSESSMENT

3.1 General Discussion

It will be clear from the foregoing that a significant body of work has been accumulated in relation to a consistent objective of introducing TTV into the modelling process. In this section we attempt to take stock of the current position, both for demand and supply components.

Though there is certainly room for improvement on the demand side, there is sufficient knowledge to proceed, at least on an interim basis. It is on the supply side that progress is weakest, in spite of the fact, it may be noted, that the Department has justifiably concentrated on this side of the work.

With this in mind, it is worth setting out clearly what the supply needs are. In the first place, since TTV is closely associated with travel time itself, it is necessary to consider what aspects of travel time are used in existing models.

For **rail**, it is invariably the scheduled journey time that is input into demand models, whether derived from networks or directly from published timetables. This means that the impact of TTV on delay time also needs to be taken into account. In other words, we require a model which, ideally, will predict the impact of any proposal/scheme on the distribution of delay relative to schedule time: as a minimum, we need the mean delay, and probably the standard deviation of delay as well. This should be available for all movements potentially affected by the scheme. There do not appear to be any tools which can currently provide such information.

On the **highway** side, existing supply models provide estimates of travel time between pairs of zones along predicted routes, using speed-flow (or volume-delay) relationships on a link basis and (usually) the assumption of Wardrop equilibrium, to deal with multi-routeing. It may be assumed, though it would be worth trying to confirm, that the link relationships, as originally estimated, are appropriate for the **average** volume of demand.

This means that, as well as quantifying TTV *per se*, it is necessary to consider the additional impact of TTV on mean travel time. As noted earlier, this needs to reflect both the effect of incidents and of DtDV. In the case of incidents, it may be assumed that their effects are **not** accounted for in existing link cost functions. In the case of DtDV, given the typically convex nature of the link cost relationships, the expectation is that the overall effect of DtDV will be to add to average journey time (assuming that the suggestion that existing link relationships are appropriate to mean demand is correct), though the effect will only be significant when average demand is close to capacity.

Note that this also raises some issues for the **validation** of supply relationships in particular circumstances. When network calibration is carried out (eg using journey time surveys), should it be assumed that both demand and travel time data represent average conditions?

3.2 Developing the Highway Supply Models

The requirements

Hence a full treatment of TTV on the highway side (implicitly assuming "macro" time periods, but see below) will involve:

additional mean travel time due to DtDV
additional mean travel time due to incidents
overall variability (eg, standard deviation) due to DtDV and incidents

This, again, will need to be available for each zone pair potentially affected by a given transport proposal/scheme, for the "with" and "without" cases.

A crucial question for the highway supply model is whether the appropriate "building blocks" of such information should be prepared at the link level (as is the basis of the existing network cost relationships) or at an O-D (zone pair) level. The current evidence is not at all conclusive. The attempt by Hyder to return to something of the link- and junction-based approach in the MVA work has not produced the hoped-for improvements.

Questions of Data

Partly this has to do with the available data, or perhaps more accurately, with the way the data is dealt with in the analysis which underlies the models. The earliest TTV data was based on a limited number of repeated journeys (eg the SDG 1993 study). By also recording the times at intermediate points on the route, it became possible to analyse the data on a link basis. This allowed intermediate "journeys" to be analysed as well, by designating appropriate points as the start and end points. If the route operated on a circular basis, as with SDG, it was possible to take any point-to-point combination as a journey.

If the data is only analysed on a link basis, then there will be a subsequent implementation issue, in terms of allowing for the combined TTV for the whole journey. By analysing the data on a journey basis, this problem is potentially avoided. It is reasonable to expect the **relative** level of TTV to fall as the journey length/number of links increases, because of the potential for making up time, essentially due to "flow metering" effects at congested junctions: indeed, such an effect can be demonstrated empirically. Nonetheless, the way in which links are combined into journeys for such analysis seems essentially *ad hoc*, and the "observations" used for model estimation are not independent.

With the more recent "electronic" data, the situation is somewhat different. Here in most cases the data does not relate to journeys at all: neither are there "link observations" as such, merely successive observations of the same vehicle along its path through a series of links. Considerable data manipulation is required to transform these observations into link-based data. Presumably these processes are described somewhere, but the studies which have used HATRIS or similar data sets do not provide a description or a reference to the methods. Whatever the value of such data sets for other purposes, their value for the analysis and modelling of TTV needs to be explicitly confirmed.

The diagram below illustrates some of the possibilities, using MIDAS, ANPR, and GPS observations (satellite). With sources such as MIDAS, where the "observation points" are fixed, the distances between them are known, but there is no ability to "connect" the data passing one point with the data passing the next point: in this case, a model is required to replicate the speed behaviour in the intervening section, and it will be necessary to carry out inter/extra-polation to allocate the estimated average travel times between points to the link. With ANPR, where the cameras are at a known fixed distance apart, the travel time, and hence speed, of individual matched vehicles is known (subject to recoding errors). However, only when cameras are mounted at the end of links will observations between successive cameras truly correspond with links: otherwise, once again, inter/extra-polation is necessary to allocate to the link. Finally, in the case of satellite-based observations (eg ITIS), the observations relate to moving vehicles, intercepted at regular intervals **of time**. In this case, the location itself is only known within definable limits of accuracy, and needs to be reconciled with electronic map information, in order to locate it on a particular link: successive "readings" will be at variable spacing along the links.

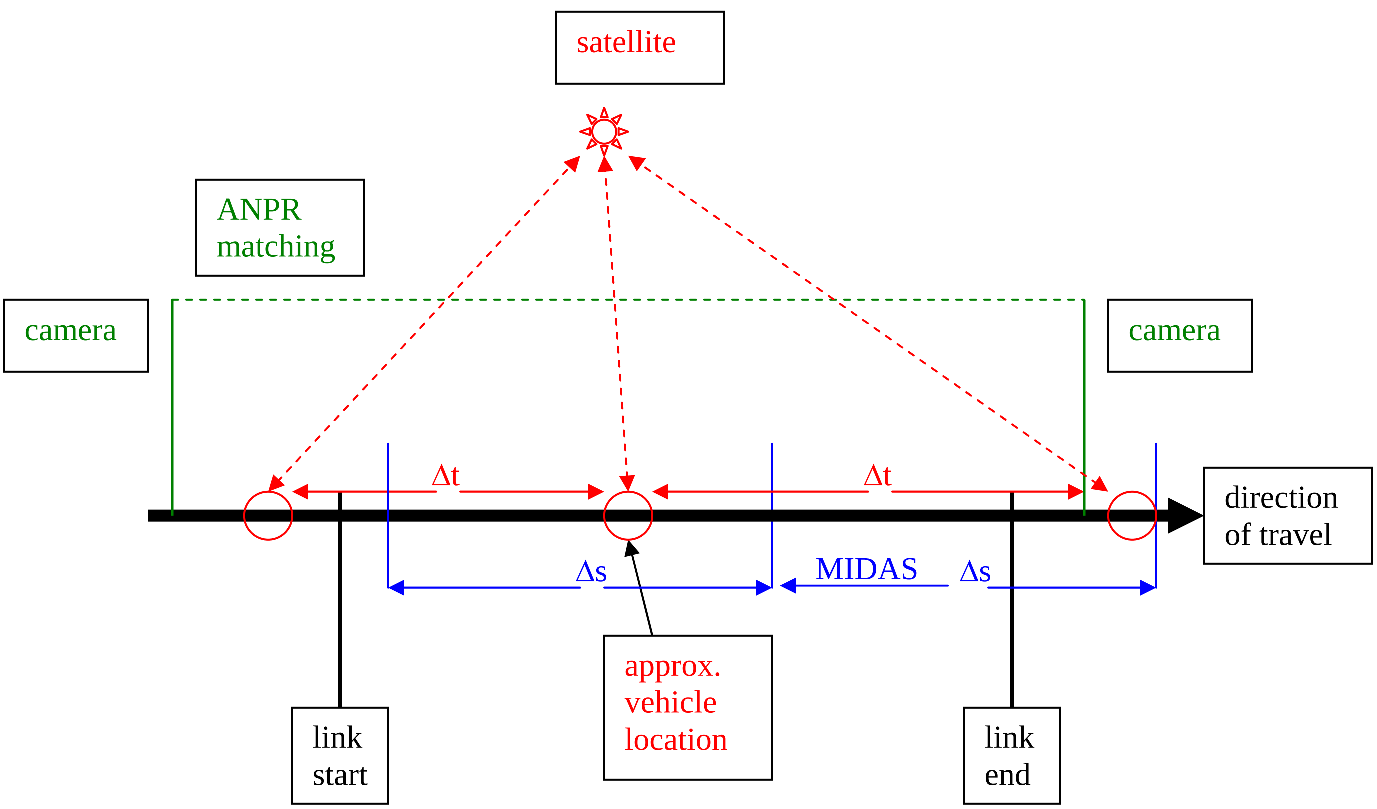

While the errors that attend each of these sources may not be important when estimating mean link speeds by time of day, it is much more likely that, when using them to estimate the **variability** of mean speeds, some of the measurement errors will be confounded with true variability: the way in which, and the extent to which, it occurs is likely to be source-specific. It would be re-assuring to think that these problems have been examined, but no evidence has been found of this.

Given an understanding of this, it would also be useful to define some rules for analysis so that results could be more easily compared between studies. Typically the link time data is available for short time periods (eg 15 minutes), and standard deviations and means are then calculated over a number of days. In constructing models for the standard deviation, it would seem preferable to carry out the analysis in a way which used all the data simultaneously. In addition, an agreement on the form of the dependent variable[8] would allow a more rigorous assessment of residuals. It is not clear how far issues of model form have been addressed in the existing work: there is some suggestion that a small number of relatively *ad hoc* formulations have been tried.

Clearly, the link travel time will depend on the link length, but whether the standard deviation will depend on the link length is less clear. Partly, this harks back to the original MVA suggestion that variability relates predominantly to junctions: on this basis, the standard deviation might be relatively independent of link length. Despite the conclusions of the Motts work, it seems that some index of capacity utilisation should have a role to play, and the "congestion index" (ratio of actual time to free-flow time) remains a convenient, and unit-free, index. Of course, once we try to allow for the observed correlation (at least at the link level) between **mean** travel time and standard deviation, we implicitly introduce the link length.

It may still be the case that simplified formulae which allow the TTV supply model to be applied directly on a route basis will be preferred. However, a proper understanding of the link-level relationship seems desirable. Of course, assuming a link-based analysis can be achieved, there remains the need to aggregate it to a route level, and this will require a further analysis of correlations. However, both the Hyder work and work by Eliasson (2007) suggests that relatively simple assumptions can be made without introducing major inaccuracy.

3.3 Summary of current position

While there is a need to make further progress on both the demand and the supply fronts, priority should continue to be given to the supply side. The demand side may be considered adequate for current requirements, but the supply side remains problematic. This is particularly the case for the highway mode, where both the level of delay and

[8] sometimes the dependent variable is link s.d., sometimes journey s.d., sometimes divided by distance, sometimes divided by mean travel time etc.

TTV itself are likely to be affected by the level of demand. The current guidance offers a rather disparate set of tools for different road types – INCA for motorways and dual carriageways, a general relationship for urban roads, and a “stress” indicator for other road types.

A first task is to consolidate **and understand** the empirical evidence for TTV. This could be achieved by a short piece of investigatory research into the basis of the various forms of electronic data (HATRIS etc), and their implications for modelling TTV. The older data sets – SDG, Leeds (and Edinburgh[9]?) – should also be carefully archived.

For those datasets which appear suitable, a more detailed investigation should be carried out to see a) whether there remains some scope for a distinction between “links and junctions” (as suggested in the MVA work), and b) what the appropriate treatment of link correlation should be. A successful outcome of this would imply that a link-level approach to TTV could be adopted for general work, and appropriate calculations could be commissioned within existing Assignment packages.

Alternatively, if this is unsuccessful, it will need to be concluded that the treatment of TTV will continue to be done by “post-processing” at the O-D level (as implied in existing Guidance TAG 3.5.7). At the moment such relationships only exist for urban journeys, so that there will be a need to re-visit interurban data at the journey, rather than the link, level.

Finally, the recent Hyder work for the HA, while confined to a particular section of motorway, does suggest that in order to have a realistic chance of capturing both delay and TTV in connection with flow breakdown, it is necessary to take a more “dynamic” approach to assignment. This involves making use of demand profiles over the (major part of the) day, rather than using separate models for different periods. The Hyder Report suggests a number of ways in which the development could be achieved.

[9] It is not known whether DfT has retained an archived copy of this data

4. PROPOSALS FOR FUTURE WORK

A short piece of investigatory research should be carried out into the basis of the various forms of electronic data (HATRIS etc), and their implications for modelling TTV. The older data sets – SDG, Leeds (and Edinburgh[10]?) – should also be carefully archived.

For those datasets which appear suitable, a more detailed investigation should be carried out to see a) whether there remains some scope for a distinction between "links and junctions" and b) what the appropriate treatment of link correlation should be.

The Department should give careful consideration to the recently completed Highways Agency project (managed by Hyder).

There should be a thorough review of bus reliability, particularly on the supply side. It may be noted that there was a FIT[11] project at ITS Leeds with the title "Model To Assess Public Transport Reliability" [see Liu, R. and Sinha S. (2007)]: this has not been investigated within the current research, however.

Separately, factors affecting rail reliability should be investigated, building on earlier work such as the November 2003 Report by AEA&T for the Department ("Predicting Rail Performance").

On the demand side, there is currently scope for investigation of the valuation of TTV under the scoping study "Updating Appraisal Values for Travel Time Savings Phase 1 Study". This should be co-ordinated so as to ensure that a compatible methodology is developed. The particular topics of bus travel and freight require special attention.

Finally, an attempt should be made to see how far the recent ITS rail work can be reconciled with the PDFH approach to reliability.

---

[10] It is not known whether DfT has retained an archived copy of this data

[11] **F**uture **I**ntegrated **T**ransport Programme – an earlier collaborative programme between the (then) DTLR and the research councils.

References

Studies commissioned by DfT

TSU Study on **Driver Stress**: see Bates *et al*, 1990
Cranfield study on **Demand Effects of Travel Time Reliability**: see Black & Towriss (1993)
SDG study on **Highway Reliability Supply Effects**: see Steer Davies & Gleave (1993)
MVA study on **Benefits of Reduced Travel Time Variability**: see MVA (1996)
John Fearon Consultancy study on **Inciben**: see "TAMTTV" DETR (1999)
Mott Macdonald study on **Travel Time Variability Follow-up Research –** 4 reports: see Mott Macdonald (2000)
Oscar Faber study on **Journey Time Variability**: see FaberMaunsell (2003)
Mott Macdonald study on **INCA**: see Mott Macdonald (2001) and Mott Macdonald (2009)
Arup Study on **Journey Time Variability**: see Arup (2002, 2003)
ITS Leeds **New Horizons Project**: see Watling *et al* (2004)
TRL/Fearon project **Updating and validating parameters for incident appraisal model INCA** : see Frith *et al* (2004)
ITS Study **Impact of Reliability on Passenger Rail Demand**: see ITS (2008)
Mott Macdonald study on **Estimation of Day To Day Variability Functions**: see Mott Macdonald (2008, 2009)
Hyder study **Forecasting Travel Time Variability in Urban Areas**: see Hyder (2007, 2008)
ITS/Polak/Bates study on **Multi-modal variability** : see Batley *et al* (2008)

(Highways Agency Project)
Hyder Study on **Motorway Travel Time Variability**: see Hyder (2010)

General references

Arup (2002) , Journey Time Variability, Deliverable D6.1, Modelling and Appraisal of Journey Time Variability – Review of Earlier Research, Research on this Contract, and Detailed Proposals for Further Research, 2002: available at http://www.dft.gov.uk/pgr/economics/rdg/ttv/backgroundstudies/

Arup (2003), Frameworks for Modelling the Variability of Journey Times on the Highway Network. London: available at http://www.dft.gov.uk/pgr/economics/rdg/ttv/backgroundstudies/

Association of Train Operating Companies (2005) Passenger Demand Forecasting Handbook. Association of Train Operating Companies, London.

Bates, J et al. Uncertainty and Driver Stress: A Review. Oxford University Transport Studies Unit, 1990.

Bates J.J., Polak J.W., Jones P. and Cook A.J., The Valuation of Reliability for Personal Travel, Transportation Research Vol 37E (nos 2-3) pp 191-230, 2001.

Batley R, Grant-Muller S, Nellthorp J, de Jong G, Watling D, Bates J, Hess S and Polak J (2008), Multimodal Travel Time Variability, Final Report

Black, I G & Towriss, J G (1993), Demand Effects of Travel Time Reliability, Department of Transport, London.

DETR, Guidance on the Methodology for Multi-Modal Studies (GOMMMS), 2000.

DETR, The Appraisal of Measures to reduce Travel Time Variability (TAMTTV), 1999.

DfT (2000). The Reliability Sub-Objective. Volume 2, in: Guidance on Methodology for Multi-Modal Studies.

DfT (2008). Draft WebTAG Unit 3.5.7: The Reliability Sub-objective. July 2008. http://www.webtag.org.uk/webdocuments/3_Expert/5_Economy_Objective/3.5.7-draft.htm Accessed 20/10/08.

Eliasson J. (2007). The Relationship between Travel Time Variability and Road Congestion, paper presented at World Conference on Transport Research, 2007

FaberMaunsell (2003). Journey time variability – outward-facing research. Prepared for the Highways Agency. Draft report, version 5.

Frith B A, Fearon J, Sutch T E and Lunt G (2004), Updating and validating parameters for incident appraisal model INCA, TRL Published Report PPR 030, available at http://webarchive.nationalarchives.gov.uk/+/http://www.dft.gov.uk/pgr/economics/rdg/ttv/incaresearch/inca/ [6 working papers also available]

Hyder (2007),:Forecasting Travel Time Variability in Urban Areas: Deliverable D1 Data Analysis and Model Development, (Nov 2007), available at http://www.dft.gov.uk/pgr/economics/rdg/ttv/highwaytraveltime/traveltimevariability/

Hyder (2008a), Forecasting Travel Time Variability in Urban Areas: Deliverable D2 Model Application, (Nov 2008), available at http://www.dft.gov.uk/pgr/economics/rdg/ttv/highwaytraveltime/traveltimevariability/

Hyder (2008b), Draft D3/D4 Final Report Forecasting TTV in Urban Areas (December 2008)

Hyder (2008c), Draft Link Based Report (November 2008)

Hyder (2010), Motorway Travel Time Variability, Final Report for Highways Agency

ITS Leeds (2008a), Revealed Preference Study to Assess Impact of Reliability on Passenger Rail Demand (Contract Number DfTRG/0027/2006), Deliverable 3: Report of Discrete Choice Analysis

ITS Leeds (2008b), Revealed Preference Study to Assess Impact of Reliability on Passenger Rail Demand (Contract Number DfTRG/0027/2006), Deliverable 4: Report of Econometric Analysis

Knight, Trevor E., 1974, "An Approach to the Evaluation of Changes in Travel Unreliability: A 'Safety Margin' Hypothesis", Transportation, 3: 393-408.

Mott MacDonald (2000). Travel Time Variability Follow-On Research 1.1. Report prepared as part of project for HETA division, DETR. available at http://www.dft.gov.uk/pgr/economics/rdg/ttv/backgroundstudies/

Mott MacDonald (2000), Travel time variability follow-on research 1.2 Report prepared as part of project for HETA division, DETR. available at http://www.dft.gov.uk/pgr/economics/rdg/ttv/backgroundstudies/

Mott MacDonald (2000), Travel time variability follow-on research: Conceptual and methodological issues Report prepared as part of project for HETA division, DETR. available at http://www.dft.gov.uk/pgr/economics/rdg/ttv/backgroundstudies/

Mott MacDonald (2000), Travel time variability follow-on research 1.3 Report prepared as part of project for HETA division, DETR. available at http://www.dft.gov.uk/pgr/economics/rdg/ttv/backgroundstudies/

Mott MacDonald (2001), INCA Version 1.1, Draft User Manual

Mott MacDonald (2008a). Estimation of variability functions for additional inter-urban road types. Report No. 233697/9/B. ., available at http://www.dft.gov.uk/pgr/economics/rdg/ttv/highwaytraveltime/daytodayfunctionsmotorways/

Mott MacDonald (2008b). Estimation of DTDV functions for motorways. Report No. 233697/5/B)., available at http://www.dft.gov.uk/pgr/economics/rdg/ttv/highwaytraveltime/daytodayfunctionsmotorways/

Mott Macdonald (2009), Development of INCA to incorporate single carriageways and managed motorways - PPRO 04/03/17, Summary Report

MVA (1996). Benefits of reduced travel time variability. MVA ref. 02/C/2923, June.

Steer Davies Gleave, Highway Reliability Supply Effects, Final Report prepared for London Assessment Division, UK Department of Transport, 1993.

Watling D, Sumalee A, Connors R and Balijepalli C (2004), Advancing Methods for Evaluating Network Reliability, Department for Transport "New Horizons" Project: Final Report (available at http://www.dft.gov.uk/pgr/economics/rdg/ttv/backgroundstudies/)