

Multimodal Travel Time Variability Final Report

11th November 2008

Imperial College
London

JOHN BATES SERVICES

Richard Batley, Susan Grant-Muller, John Nellthorp, Gerard de Jong, David Watling (ITS Leeds)

John Bates

Stephane Hess and John Polak (CTS Imperial)

A report for the Department of Transport

Multimodal travel time variability

VERSION Final
18th November 2008

Richard Batley
Susan Grant-Muller
John Nellthorp
Gerard de Jong
David Watling
Institute for Transport Studies,
University of Leeds,
Leeds, LS2 9JT
Tel: +44 (0)113 343 5325
Fax: +44 (0)113 343 5334

John Bates
John Bates Services
1 Picklers Hill, Abingdon OX14 2BA
Tel: +44 (0)1235 526702

Stephane Hess
John Polak
Centre for Transport Studies
Imperial College London
London SW7 2AZ
Tel: +44 (0)20-7594-6100
Fax: +44 (0)20-7594-6102

Contents

Executive Summary	4
1. Introduction.....	8
1.1 Outline of the study.....	8
1.2 Definitions of terms	11
1.3 Sources and Measurement of Travel Time Variability.....	12
1.4 Requirements for modelling and appraisal	13
2. Theoretical methods for demand modelling and valuation of variability	15
2.1 Choice under Uncertainty	15
2.2 The mean vs. variance approach.....	20
2.3 The scheduling approach	22
2.4 Equivalence between the mean vs. variance approach and scheduling approach ..	23
2.5 Applying the theory to the valuation of reliability.....	24
2.6 Challenges and opportunities for further theoretical research	25
2.7 Recommendations on theoretical approach	28
3. Evidence on Valuation of Travel Time Variability.....	30
3.1 Introduction.....	30
3.2 Presentational issues relating to variability (passengers).....	31
3.3 Derivation of PAT for scheduling models (passengers).....	37
3.4 Issues in valuation studies for freight transport	38
3.5 Empirical findings for passenger and freight.....	39
3.6 Recommendations on valuation.....	41
4. Supply issues.....	47
4.1 Highway assignment models: state-of-practice and issues for representing TTV..	48
4.2 Representing actual TTV in highway assignment models.....	51
4.3 Modelling highway network performance under TTV	54
4.4 Towards ‘Implementable’ Models.....	56
4.5 Widening the Range of Models	63
4.6 Public transport supply issues.....	65
4.7 Concluding remarks	66
5. Integration with Appraisal.....	68
5.1 Introduction.....	68
5.2 Implications for incorporating “Reliability” into the Appraisal Procedure	69
6 Conclusions.....	70
Annex I: Development of theory.....	72
Annex II: Valuation of travel time reliability.....	79
Annex III: Arup Transport planning Appendix B	88
Annex IV: References	100

Executive Summary

This report has been commissioned by the Department for Transport under contract number PPRO 4/3/5 as a scoping study with the aims of a) defining a way forward on the theory of (travel time) reliability across modes; and b) identifying areas of research needed to arrive at a practical solution to quantifying and valuing (travel time) reliability changes. A consideration of terminology concluded that the terms *reliability* and *variability* should not be used interchangeably because their meanings are opposites despite common practice. Hence the focus has been defined as ‘travel time variability’, one possible measure being the variance or standard deviation of travel time.

The underlying theory is outlined according to the basic principles and evolution of the theory in recent research. The two major candidates are the so-called “mean-variance” approach, and the “schedule utility” approach, each employing rather different utility functions. For the first it is assumed the traveller will consider travel time variability by choosing the departure time that maximises his or her expected utility, where expected utility could be approximated as a linearly additive function of the mean and variance of the travel time distribution, though in practice the standard deviation of travel time (having the same units as the mean) is often substituted for the variance: confusingly, both cases are referred to as “mean-variance” approaches. Empirical investigations have frequently simplified from the full functional form and with the assumption of constant coefficient of variation this has led to the so-called ‘reliability ratio’ term. This is given by the ratio of the coefficient of the standard deviation to the coefficient of the mean within the expected utility function. Variability can then be approximately allowed for simply by a “loading” for the value of time using this ratio.

The second basic approach is based upon the notion of scheduling and reformulates the utility function for a particular departure time choice as a function of four components; travel time, ‘schedule delay early’ (SDE), ‘schedule delay late’ (SDL), and a ‘lateness’ dummy variable (unity if schedule delay late is non-zero). The latter three components are conditioned by the notion of a ‘Preferred Arrival Time’ (PAT). Journeys arriving before the PAT are deemed ‘early’: the SDE is the difference between the PAT and the actual arrival time, SDL is zero. Journeys arriving after the PAT are deemed ‘late’: the SDL is the difference between the actual arrival time and the PAT, the lateness dummy variable is unity and SDE is zero. These four components of the utility function are typically specified as linearly additive. Variability in travel time is accommodated by taking expectations of each component of utility over the travel time distribution giving rise to a term representing the probability of lateness.

Work has been undertaken to investigate the equivalence between the two approaches at both a theoretical and empirical level. The evidence from both perspectives suggests this is highly dependent on the travel time distribution. Any equivalence between the two approaches also depends on the ability of travellers to continuously vary their departure time, which is generally not the case for public transport. In our view, it is not currently possible to choose between these two approaches without a clearer understanding of their

similarities and differences. While the mean vs. variance approach, together with the associated 'reliability ratio', has the advantage of greater simplicity, this does not detract from our belief that the scheduling approach may reveal important, and unique insights into behavioural responses to unreliability, particularly in the context of scheduled (but infrequent) public transport services. Specifically, further research is needed in the following areas: translating the theory to public transport, the interpretation of lateness and alignment with industry conventions, valuing travel time variability (as distinct from valuing time risk) and the extension of choice dimensions.

Empirical evidence on the valuation of travel time variability is then presented drawing on research evidence for passenger and freight journeys from the UK and abroad. Almost all the empirical work to obtain values for variability, using either the mean vs. variance or the scheduling approach or both, has been based on Stated Preference (SP) data. This is due to the difficulties in collecting Revealed Preference (RP) data that includes measures of variability, travel time and travel costs that are not heavily correlated. The study results are determined both by the way data has been collected and by the way it has been analysed, resulting in a wide range of studies with outcomes which are not necessarily presented in a comparable format. Where comparable results are given there is still diversity in the empirical values obtained due to the nature of the study, the sample taken and the level of disaggregation in the findings. While all reviewed studies agree that variability is a factor of substantial importance, there are no generally accepted monetary values for variability, or indeed a reliable estimate of the relative weight of travel time and travel time variability. A recommendation from this study is for substantive new research into the valuation of variability, investigating both how information on variability is presented to respondents and how it is later specified in econometric models estimated from the data. Before this can take place, agreement must be reached on the theoretical framework for journey time variability as outlined above.

For the purposes of this study, the review of supply effects concerning TTV has been largely confined to the role of the network model. Although the body of research has addressed the performance of both public and private transport networks, the majority of the work has focused on the highway context. Two broad areas are considered here: firstly, how to represent actual TTV, and secondly, how to model the effect of TTV on network performance. On the first issue the approaches may be divided into those (the majority) that look at the composite effect on TTV of all sources, and those that aim to decompose the variability into its component sources, ie day-to-day-variability of demand and incidents. Many studies have focused on the standard deviation as sufficient for representing TTV, but there is a question as to whether the complete distribution of travel times is required to obtain a true picture of TTV impacts. Further thought should be given to whether the objective of representing actual TTV needs to be widened to summary measures beyond the second moment, given the typically asymmetric impact of flow variations and incidents on the right-hand tail of the travel time distribution.

A number of theoretical and practical challenges remain in incorporating travel time reliability in network assignment models, echoing the debate on the relative merits of the mean vs. variance and scheduling approaches. A more fundamental question is also

raised, ie what the mechanism for transferring advances in network assignment theory into practice is. Looking to widening the scope of modelling tools considered for this purpose, the methods illustrated as examples are practicable and realistic. While they do not go so far as to criticise the theoretical foundation of current practice (e.g. they accept concepts such as 'equilibrium' even those these would not be accepted by all researchers when modelling a variable environment), they do need specialist software tools that can make the best use of current research and knowledge. The incorporation of travel time reliability within practical transport planning tools is an area that is evolving quickly and requires the continued support of the Department to keep apace with the research developments.

Against this background it is challenging to make recommendations on procedures for use in multi-modal appraisal. There is a practical imperative, however as failing to analyse variability implies omitting variability improvements from appraisal and decision-making, in direct contradiction both to policy (for example, the Eddington Study) and the majority of the research evidence. The evidence indicates that transport users are willing-to-pay something for improvements in travel time variability – the questions are really 'how much' and 'for which measure of variability'. A judgement is therefore needed on whether the theory and evidence is sufficient to propose an interim approach to evaluating variability. As a response to the practical imperative the DfT have recently produced a new consultation version of TAG Unit 3.5.7, recommending using the best available data/knowledge available for practitioners at this time. Essentially this implies use of mean-delay approach for rail and the reliability ratio for other modes.

We understand the need for a possible interim approach. Nonetheless, our general view is that proceeding along these lines is premature until a) we have agreed on a theoretical formulation, and b) resolved what practical models can be built. At this stage we have therefore gone no further than outlining a list of the elements of the appraisal that will need to be renovated once both the theoretical formulation and practical models are in place.

A key recommendation from this research is that there is a need for a major new study into the valuations of variability, investigating both the question of how information on variability is presented to respondents and how it is later specified in econometric models that are estimated on the data. This must, however, follow on from agreement on the theoretical framework for journey time variability. The aim should be to achieve a high level of consistency between the underlying theory, the data collected and the estimated model.

With respect to incorporating travel time reliability in network assignment models, a number of theoretical and practical challenges remain, mainly echoing the issues concerning the relative merits of the mean vs. variance and scheduling approaches. A more fundamental issue remains, however: namely, the mechanism for transferring advances in network assignment theory into practice. For further advances in the practice of TTV modelling to take place, the practice of issuing guidance for use in a small

number of accepted commercial packages may need to be reconsidered - the use of specialist software tools that can make the best use of current research and knowledge should also be supported by the Department.

1. Introduction

1.1 Outline of the study

This report has been commissioned by the Department for Transport under contract number PPRO 4/3/5 as a scoping study to address the issue of general topic of (travel time) reliability with the following aims:

- to define an agreed way forward on the theory of (travel time) reliability across modes;
- to suggest areas of research by which DfT could expect to arrive at a practical solution to quantifying and valuing (travel time) reliability changes.

Although the term “reliability” is widely used, we have been concerned as a team that it is too broad and ill-defined a term to be useful. In the next sub-section, we discuss general questions of terminology, as a pre-requisite to setting out the theoretical issues. However, in the light of that discussion, we have decided to use the term “travel time variability” in preference, and henceforth our discussion will be in those terms. This has also necessitated a change in the title of this report.

The range of transport interventions that the Department would wish to model or would wish scheme promoters to model is broad, and therefore the theoretical approach should be general enough to be potentially transferable to all of them. Whilst passenger trips are the primary focus, the role of freight in theoretical and empirical evidence is summarised, with key findings and references included. Moreover the objectives of the Scoping Study are twofold. Firstly, to review and summarise the current state of the art; covering the theory, evidence, appraisal, valuation and tools for travel time variability in a multimodal context. Secondly, to outline recommendations on research to take the state of the art forward towards a set of technically defensible but practical and operational recommendations on the measurement and inclusion of travel time variability impacts in new scheme appraisal.

A summary of the structure to the work is given in Figure 1.1 – this is a complex area of research and part of the challenge of the research is to provide a consistent and integrated approach to the treatment of variability based on a number of previous studies which have taken place on a more ad-hoc basis.

The first stage of this study has been to review and summarise the current state of knowledge on the modelling and valuation of travel time variability for surface transport. The underlying theory is outlined in section 2, describing the basic principles of dealing with travel time variability and evolution of the theory in recent research. Following on from this, empirical evidence on the valuation of travel time variability is presented in section 3, drawing on the research evidence for passenger and freight journeys from the UK and abroad. As the number of studies is relatively small, we summarise each one in terms of the key features of the method and the subsequent outcomes, and we identify gaps in knowledge in travel time variability valuation. In section 4 a discussion of how

travel time variability could be included in practical network modelling is given. Two broad areas are considered here: firstly, how to represent actual travel time variability and secondly how to model the effect of travel time variability on network performance. The focus does not extend to microsimulation models and remains largely on highway network performance due to the well developed nature of the research in that field. Finally, in section 5, the implications of the above for appraisal are outlined. Obviously, the ability to carry out appraisal is dependent upon being able to measure and forecast journey time variability, which in itself is dependent on the state of the art in theory and valuation.

Whilst the underlying objective in each section has been to aim towards a consistent way forward, areas of difference in academic perspective are also represented. During the process of gathering the theoretical and practical evidence a number of gaps in theory, evidence and methods have been identified. These form the basis for areas of further research which have been outlined within each section.

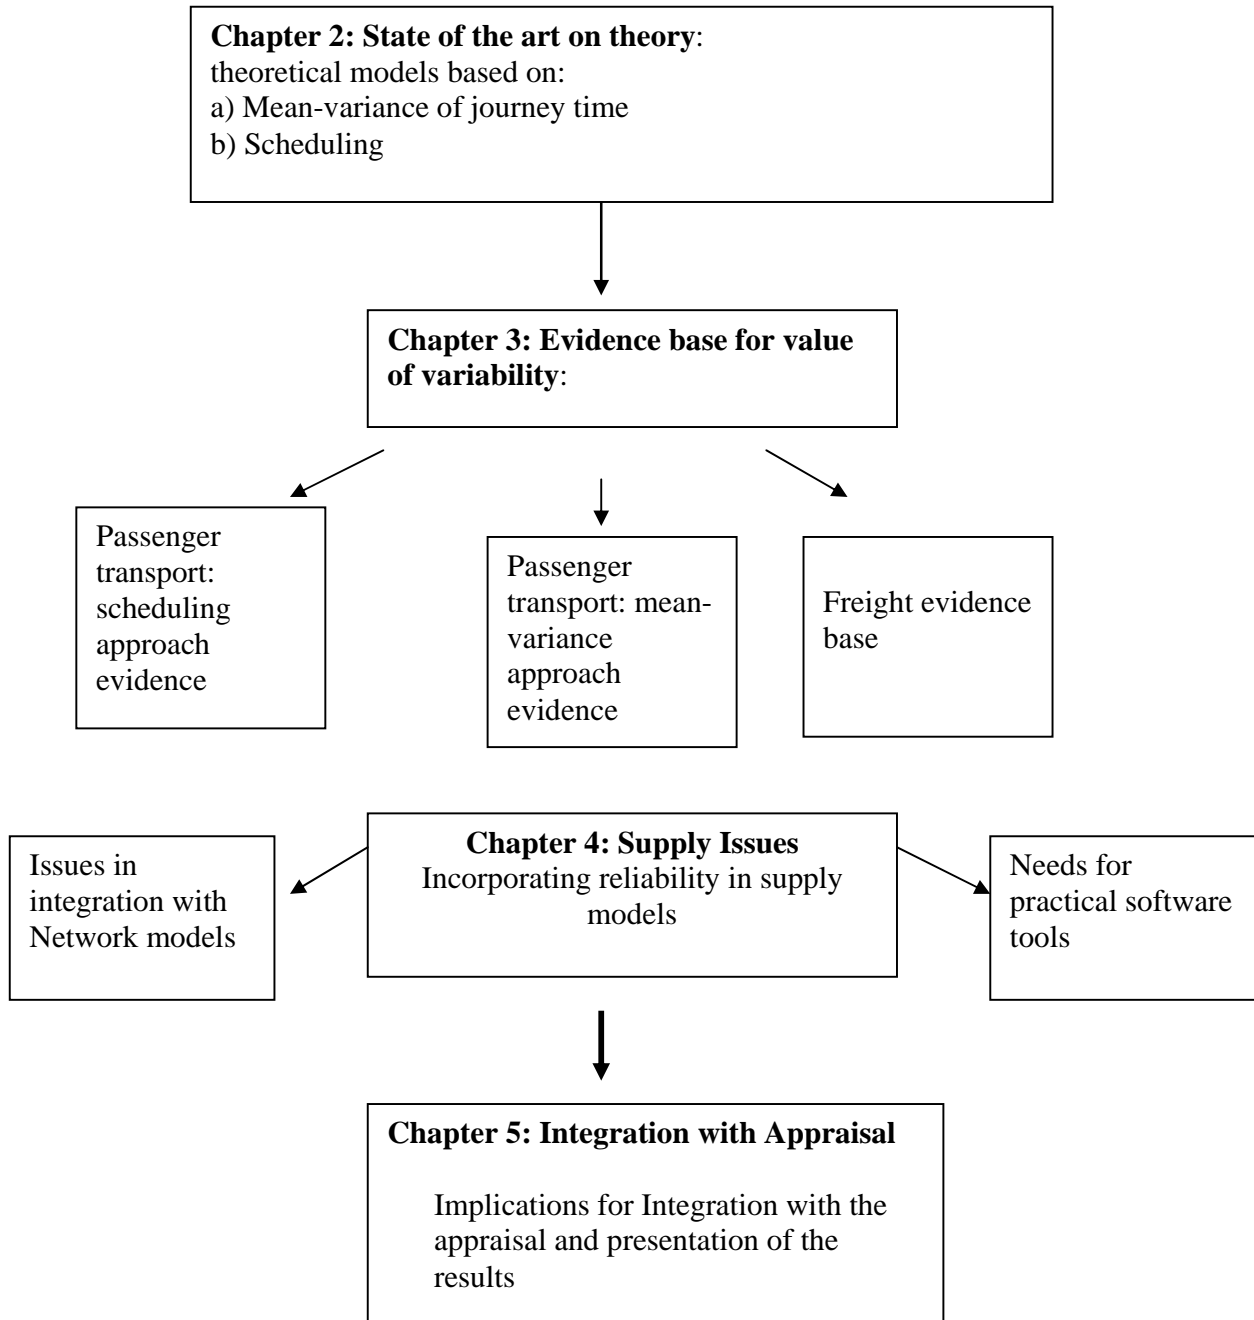


Figure 1.1: Structure of the research

1.2 Definitions of terms

It is essential to be clear about terminology from the start. We will use the commonly accepted scientific definitions for variance and standard deviation, as follows:

- *variance*, the mean of the squared deviations of a set of quantities, in this context journey times;
- *standard deviation*, the square root of the variance.

We will use a definition of *variability* from the Oxford English Dictionary:

- *variability*, the fact of, or capacity for, varying in amount, magnitude or value¹;
- hence *journey time variability*, the fact of, or capacity for, variation in journey time.

Therefore:

- *variance of travel time* or *standard deviation of travel time* are measures of the variability of journey time;
- other measures which could be used to characterise the variability of journey time include measures of skewness or kurtosis, although these measures are only meaningful if variance is non-zero;
- knowledge of the whole probability distribution of journey times would give the most comprehensive understanding of variability.

Reliability itself has been the most challenging to define; however, the OED helps us to see that there are multiple definitions:

- *reliability*, the quality of being reliable – *reliable*, that may be relied upon; in which reliance or confidence may be put; trustworthy, safe, sure;
- *reliability (Statistics)*, the extent to which a measurement made repeatedly in identical circumstances will yield concordant results.

The second, narrower definition of reliability seems to imply the reverse of *variability*, i.e. a journey time which does not vary when repeated. We have said that variability can be measured by the standard deviation or variance, and it seems that perfect reliability corresponds to zero variance. In this report we aim to simplify by using the term *variability* for this narrower concept.

The first, broader definition includes the ideas of trustworthiness and reliance, in this context whether a transport service can be trusted and relied upon by users, in terms of its journey time. The determinants of reliability in this broader sense are likely to include the *information* available to users, and the ways in which users form their *expectations* about journey time. For example, if a timetabled transport service is on average 10 minutes late

¹ Oxford English Dictionary, 2nd edition (1989), definition 2a. – definition 1. extremely similar.

compared with the timetable, whether or not this is a problem for users will depend their capacity to anticipate it. Users who rely on the timetable information – irregular users, for example – may find their expectations are wrong and they are late at their destination. Regular users of the service may form different expectations, learning by experience, and in effect modifying or rejecting the timetable information.

These issues are not limited to public transport, since irregular users will also face challenges predicting car journey times on the road network, even with the help of journey planning tools. These are under-researched areas of the topic of reliability, in which it turns out to be harder to make recommendations for modelling and appraisal practice, and where we are working on ideas for relevant research.

In the light of these definitions, it would be preferable not to use the terms *reliability* and *variability* interchangeably, both because that risks confusing a broader with a narrower interpretation of reliability, and because their meanings are opposites, the former suggesting low variance, the latter suggesting high variance.

On this basis, as already noted, we prefer **not** to use the term reliability, which we will leave to be used for wider concepts. We re-define our subject matter as relating to travel time variability, and we note that one possible measure is the variance or standard deviation of travel time. Nevertheless, the most appropriate measure is not being defined at this stage, for reasons which we now go on to discuss.

1.3 Sources and Measurement of Travel Time Variability

Having clarified the general definition, there remain problems of application and measurement, and these are largely related to the source and context of travel time variability. For example, the journey time between two stations on the rail network may show appreciable variation due to a different pattern of services (eg fast trains on the hour, slow trains on the half-hour). These are **scheduled** differences, and are in principle completely predictable (ie timetabled). Hence, even though some travellers may not be aware of them, they can be ascertained, and should therefore **not** be included in a measure of variability.

In a similar vein, on the highway side there are observable regularities in the variation of travel time by time of day (eg peak vs off-peak), and day of week. Much of this “regular” variation will be due to the impact of **demand** variation, through its influence on congestion. These variations are not **scheduled**, as in the rail service pattern case discussed above, but they are, at least to some extent, predictable. As Bates (2000) has put it:

*Since the essence of any measure of variability (such as the variance) relates the variations to the **expected** value, alternative definitions of the expected value will clearly have an impact. A failure to clarify this point in the past has led to much confusion of*

measurement. In general, it is sensible to remove as far as possible any non-random effects.

For example, a person who regularly travelled by car into the city centre in the early hours of the morning **might** be surprised by the time that the same journey took when he made it in the morning rush hour, but it would be perverse to view this as an example of unreliability. it is reasonable to expect travellers to foresee the impact of **predictable** variations in demand. Such impact relates predominantly to the highway mode, via the usual supply (speed-flow) relationships, but there are also potential impacts on rail modes.

From a modeller's point of view, there are, in addition to variations in the (average) demand profile during the course of a day, seasonal effects, day-of-week effects, and of course specific period effects (eg school holidays). After accounting for all this, the residual day-to-day variations in demand will typically be classified as essentially random. We will also need to take account of random effects on the **supply** side, which will be predominantly due to "incidents" (including vehicle breakdowns, signal failures, burst mains etc.).

On the basis of this kind of reasoning, the Arup study (Arup 2003) adopted the definition of "JTV" [journey (or travel) time variability] as "unpredictable variation in journey times", and went on to clarify:

One of the components of JTV is due to 'incidents', while what remains is referred to as 'day-to-day variability' (DTDV). and this in turn can be divided into two components: one due to unpredictable variations in demand, and the other due to random fluctuations in capacity, as represented in the following 'equations':

$$JTV = DTDV + \text{Incident-related variability}$$

$$DTDV = \text{Demand-related effects} + \text{Capacity-related effects}$$

Hence, in making any measurements of travel time variability, we need to define clearly what observations are in scope, and with respect to what the deviations will be measured.

Note that although this last sentence is cast in terms of what would typically be needed for the calculation of **variance**, measurements such as variance (and mean) are simply descriptive statistics for the distribution of travel time $f(T)$ and don't on their own define the full theoretical distribution.

1.4 Requirements for modelling and appraisal

In 1999, DfT's ITEA Division reviewed the scope and aims of their research programme in respect of travel time variability, envisaging 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 networks' TTV performance (both highway and public transport) - need to be handled together, to ensure the two are in equilibrium, as illustrated in figure

1.2 below. It also 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.

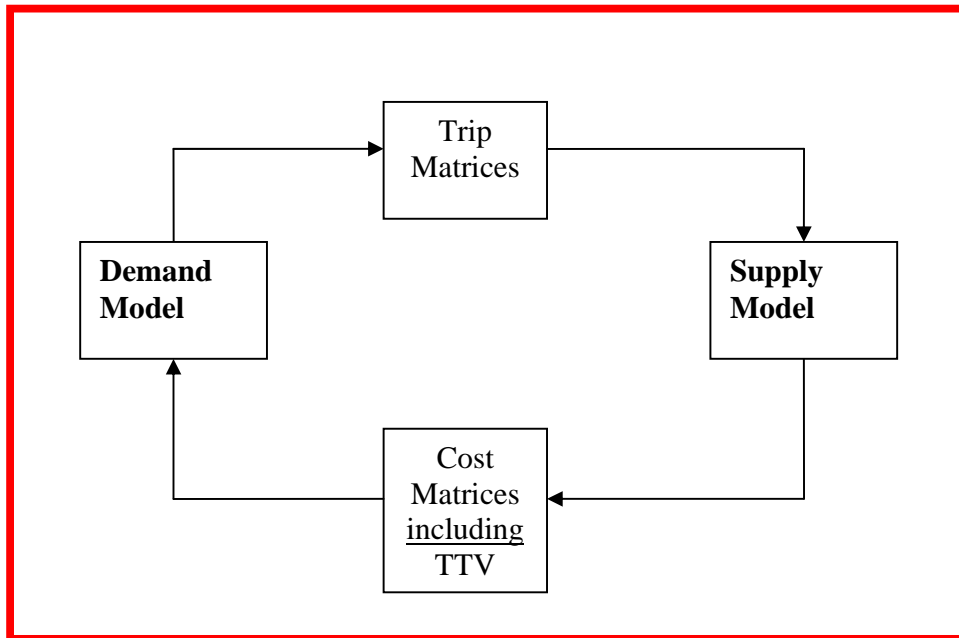


Figure 1.2: Equilibrium for generalised cost including TTV

Hence, this requires a) an agreed measure of travel time variability, b) a forecasting procedure (“supply model”) which can output the agreed measure in both a reference and “scheme” context, and c) a demand model which can respond to the difference in this output measure. Given all this, the supply and demand models need to be embedded in an equilibrium system (at least insofar as there are likely to be supply effects due to variations in demand), along the lines of current WebTAG recommendations for “variable demand modelling”.

As well as being potentially responsive to different levels of demand, the supply model is also the source of most policy input. Thus, in TTV terms, it will need to provide forecasts of TTV as a result of specific policies relating not only to standard network improvements and pricing but also (for example) the impact of improved signalling of rail junctions, provision of passing places for rail and buses, improved response to incident detection on the highway, etc.

It is essential to be clear about the different remits of the demand and supply models. As already indicated, while there is a reasonable amount of evidence relating to demand (and, related to that, **valuation**), the supply side modelling of travel time variability remains rudimentary. Of course, both sides must be in place if there is to be a workable overall model.

2. Theoretical methods for demand modelling and valuation of travel time variability

(Note: this chapter aspires to offer a reasonably succinct summary of the theory, some of which is discussed in more detail in Annex I).

2.1 Choice under Uncertainty²

The fundamental proposition is that travel time is randomly distributed. The travel time distribution is, importantly, conditioned by the time of departure. For a given time of departure we may assume that in addition to the ‘free flow’ travel time there is also potential ‘recurrent delay’ due to congestion, such that the summation of the two is considered ‘predictable’. The first moment of this distribution accommodates both these notions, as well as the average of the purely random effects. The second moment of this distribution (the variance) provides one measure of the notion of ‘travel time variability’, distinct from recurrent delay in the sense that it is ‘uncertain’.

Given the general practice in transport modelling of making use of discrete choice theory, based on an underlying microeconomic “utility” theory, it will be useful to maintain this approach. However, the random nature of the travel time component means that the utility itself becomes random, so that the discrete choice has to be exercised under conditions of uncertainty: we can refer to this as “Risky choice”³.

As noted by Liu and Polak (2007) this notion of a stochastic utility is essentially different from the notion of “random utility models” in discrete choice theory, where the “error terms” reflect the **modeller’s** uncertainty about aspects of the choice process: here we are dealing with the case where, from the **traveller’s** point of view, some key aspects relating to the choice cannot be treated as deterministic. In random utility theory (RUT) we deal with situations in which there is a one-to-one correspondence between an act (e.g., implementing a travel decision) and the consequence of that act for the decision agent – i.e., it is a theory of riskless choice. In particular, the error term(s) in RUT accommodate *inter alia* unobserved heterogeneity and inter-alternative correlation but *not* agents’ uncertainty. Thus RUT does *not* provide an inherent treatment of risky choice – we need to extend it to do so. We can characterise this in Figure 2.1

In risky decision problems an act may give rise to several different consequences and we don’t know, *ex ante*, which will occur. We can therefore speak of choice between different *prospects*, with each prospect r being defined as $(x_1, p_1; x_2, p_2; \dots, x_n, p_n)$, where

² The presentation and characterisation of risky choice problems and their relationship to existing travel demand models presented in this section draws extensively on the work of Liu and Polak (2007) and Michea and Polak (2006).

³ Note that while in general we refer to choice under uncertainty, in practice (from a microeconomic standpoint) we mean “risky choice”, which assumes that we know the probabilities of different possible outcomes. This is in a sense an approximation to true choice under uncertainty, which occurs when we **don’t** know the probabilities (which is of course the position).

each x represents the outcome vector of a possible consequence and each p represents the corresponding probability.

		Agent's Uncertainty?	
		No	Yes
Modeller's Uncertainty?	No	Deterministic utility theory	Expected utility theory
	Yes	Random utility theory	?

Figure 2.1: Risky choice and Random Utility Theory

This provides an entry point to the vast literature on economic choice under uncertainty. The dominant theoretical paradigm within this literature is expected utility maximisation (the classic approach is by von Neumann & Morgenstern (1947) - for a useful discussion of choice under uncertainty, see Deaton & Muellbauer (1980: §14)), though other variants have also been considered, as we discuss below.

According to this paradigm, choice outcomes follow a probability distribution, and expected utility is the expectation of utility across these choice outcomes. Given a risky prospect $r = (x_1, p_1; x_2, p_2; \dots, x_n, p_n)$ agents are assumed to seek to maximise

$$V(r) = \sum_j p_j v(x_j)$$

where $v(x_i)$ is termed a (von Neumann and Morgenstern [vNM]) utility function, mapping the outcome vector x_i into utility space


In other words, travellers choose the course of action which, bearing in mind the probabilities of different outcomes, has the highest expected utility. The approach implies that the traveller needs to assess all the eventualities resulting from different possible outcomes, though in practice, of course, he is likely to adopt a simpler version of this strategy.

A common variant on the expected utility model is to dispense with the use of the objective probabilities for different outcomes, i.e. those presented to respondents, instead

using ‘subjective’ probabilities. This allows for the possibility of respondents giving different weights to different outcomes. In this approach, the “true” probabilities p would be replaced by $\pi(p)$, where this potentially also has a number of other parameters, and possibly also interacts with the attributes of the different outcomes. An in-depth discussion of this topic is given in Michea & Polak (2006), who test a number of different specifications, some of which outperform the standard expected utility model in terms of model fit. The outcomes of this study suggest that respondents tend to distort the presented probabilities, underweighting the probability of negative outcomes relative to positive ones, hence reflecting optimism.


Whereas the economic literature devotes its interest almost exclusively to uncertainty in monetary outcomes, this contrasts with our present interest in reliability, which embodies uncertainty in travel time rather than uncertainty in cost (acknowledging that the former may in itself incur some incidental cost). From the point of view of travel time variability, the most interesting choice to consider is that of departure time.

When the outcome vector contains a single variable (eg the amount of money gained or lost), then if the (vNM) utility function is linear in this variable, the expected utility will correspond identically with the utility in which the variable is set to its mean value. In such circumstances, the theory itself does not imply any loss of utility due to uncertainty. From this it can be seen that the agent’s attitudes to risk are embodied in the curvature (non-linearity) of $v(\cdot)$, as illustrated in the diagrams below. The so-called “risk premium” is defined as the difference between the mean value of the variable \bar{X} and the value \hat{X} at which the mean utility is obtained. In the first diagram, the risk premium is positive, demonstrating an aversion to risk (the standard case). By contrast, the second diagram depicts a “risk-prone” situation, where the agent is willing to pay to face a risky situation.



$$\text{Risk premium} = \bar{x} - \hat{x} > 0$$

Figure 2.2 Risk averse agent



$$\text{Risk premium} = \bar{x} - \hat{x} < 0$$

Figure 2.3 Risk prone agent

However, in cases where the outcome vector consists of more than one variable, there is potential for confusion, as Liu & Polak (2007) have recently pointed out. Essentially, the concept of utility can be used in two subtly different ways. In line with RUT, utility can be thought of as a way of combining a set of variables into a single “metric”, usually, but not exclusively, using a linear form. An important example of this is the standard “generalised cost” metric combining cost and time variables. However, as we have seen, the vNM utility function carries with it some attitude to risk, and this is not present in the RUT formulation.

When the (RUT) formulation of utility is linear, so its’ expected value can be obtained by setting all random elements to their mean values, there is no reason for confusion: any “premium” associated with uncertainty must be dependent on the curvature of the vNM utility function. However, when the expected value of the (RUT) utility is **not** obtainable by merely substituting the mean value of the random elements, it is possible to generate an apparent uncertainty premium which does **not** require the introduction of a vNM utility function. We will see below that this is the case with the so-called “scheduling” model when it is used in conjunction with a travel time distribution.

In this latter case, Liu & Polak show that it is still possible to propose a von Neumann and Morgenstern utility function, and they demonstrate, using an appropriate dataset, that

an improvement in explanation can be obtained by doing so. On this basis, they propose that the term “expected utility” should be reserved for models where a vNM utility function is explicitly present, and the term “expected **value**” should be used for the case where we are merely taking the expectation of an RUT utility formulation, given randomness in some of the components, although it must be acknowledged that there is some disagreement within the team about the value of this terminology.

In practice, as noted by Liu & Polak, “the existing transport literature has combined the EUT and RUT perspectives in a variety of entirely plausible but nevertheless essentially *ad hoc* manners to draw conclusions regarding attitudes towards risk, principally, but not exclusively, those associated with variability in travel time.” The two major candidates in this respect are the so-called “mean-variance” approach, and the “schedule utility” approach, which we now go on to discuss.

2.2 The mean vs. variance approach

Suppose we write $U = U(T)$ where U is a **RUT** utility function which we may assume is linear in T (ie, in essence, of the generalised cost form), and T is a random variable with distribution $f(T)$. In addition, we postulate a non-linear vNM function $v(U)$. It follows that the expected vNM utility is given by:

$$E[v(U)] = \int_0^{\infty} f(T) \cdot v(U(T)) dT$$

Write $\mu = \int_0^{\infty} f(T) \cdot T dT$ for the mean travel time. We can develop a general formula for $v(U(T))$ as a Taylor series about the value $v(U(\mu))$, ie:

$$v(U(T)) \approx v(U(\mu)) + (U(T) - U(\mu)) \cdot v'(U(\mu)) + \frac{1}{2}(U(T) - U(\mu))^2 \cdot v''(U(\mu)) + \dots$$

Taking the expected value of both sides, this gives:

$$E[v(U(T))] \approx v(U(\mu)) + (E[U(T)] - U(\mu)) \cdot v'(U(\mu)) + \frac{1}{2}E[(U(T) - U(\mu))^2] \cdot v''(U(\mu)) + \dots$$

Now **if** U is linear in T , then $E[U(T)] = U(\mu)$ so that the 2nd term on the RHS falls out, while $E[(U(T) - U(\mu))^2] = \lambda^2 \cdot \sigma^2$, where λ is the coefficient on T in the RUT utility function. Hence for linear U , we obtain:

$$E[v(U)] \approx v(U(\mu)) + \frac{1}{2} \lambda^2 \sigma^2 \cdot v''(U(\mu)) + \dots$$

This shows that the expected utility can be approximated by a) inserting the mean value for T in the utility function, and b) adding an additional term in the travel time variance

(σ^2), which is multiplied by half the second derivative of the Utility function at the mean. On this basis, we can write expected utility:

$$E[v(U)] \approx \phi(\bar{T}) + k.\sigma^2 \quad (2.1)$$

Hence in these terms, the traveller will, in the face of travel time variability, choose the departure time that maximises his or her expected utility, where expected utility is a linearly additive function of the mean and variance of the travel time distribution. Empirical evidence suggests that both the mean and variance are “bads”, and it is implied that the term in σ^2 can be seen as a reflection of the “cost of unreliability”.

In practice, there is some ambivalence between this “pure” derivation using the travel time variance σ^2 (as used, for example by Jackson & Jucker (1981)) and approaches using the standard deviation of travel time σ , which has the advantage of being the same units as the mean. In a somewhat unrigorous way, both approaches are referred to as “mean-variance” approaches: most cases in the literature have, in fact, used the standard deviation. In addition, for empirical investigations, the implications of the functional form $\phi(\bar{T})$ have been ignored, and the most commonly applied function has the simplified form:

$$E[v(U)] = \alpha\bar{T} + \lambda\sigma \quad (2.2)$$

Practitioners commonly refer to the so-called ‘reliability ratio’ $\tilde{\rho} = \lambda/\alpha$. This allows the relevant part of (2.4) to be written: $\alpha\bar{T} \cdot (1 + \tilde{\rho} \cdot \sigma/\bar{T})$, so that if the coefficient of variation (σ/\bar{T}) is constant, the allowance for reliability can be seen to be simply a “loading” for the value of time.

In practical terms within the transport field, models of risky choice based on mean and standard deviation have primarily been applied to route choice. Suppose a traveller has the choice of J routes, each of which has a travel time T_j which is a random variable with distribution $f_j(T_j)$. A number of authors have suggested models in which the traveller chooses the route j according to the principle:

$$\min_j [E(T_j) + \lambda SD(T_j)]$$

Essentially, this collapses the “prospect” for each route to an equivalent single alternative.

2.3 The scheduling approach

This derives from the pioneering work of Vickrey (1969) and Small (1982) to departure time choice, originally applied in a deterministic context. Later work by Noland and Small (1995) developed these ideas in the context of travel time variability, building on the earlier theoretical contributions of Gaver (1968) and Polak (1987).

One might see the scheduling approach as an attempt to instil the previous discussion with greater intuition in terms of travel behaviour. To this end, the scheduling approach reformulates the (RUT) utility function for a particular departure time choice as a function of four components; travel time, ‘schedule delay early’ (SDE), ‘schedule delay late’ (SDL), and a ‘lateness’ dummy variable that is set to unity if schedule delay late is non-zero. The latter three components are conditioned by the notion of a ‘Preferred Arrival Time’ (PAT), as follows.

- Journeys arriving before the PAT are deemed ‘early’. In this case, the SDE is derived as the difference between the PAT and the actual arrival time (i.e. the number of minutes of earliness), and SDL is zero.
- Journeys arriving after the PAT are deemed ‘late’. In this case, the SDL is derived as the difference between the actual arrival time and the PAT (i.e. the number of minutes of lateness), the lateness dummy variable is unity, and SDE is zero.

The four components of the utility function are typically specified as linearly additive, thus:

$$U = \alpha T + \beta SDE + \gamma SDL + \delta L \quad (2.3)$$

where:

SDE is schedule delay early

SDL is schedule delay late

L is a dummy variable set to unity if *SDL* > 0, otherwise zero

Note, however, that because of the truncated form of the last three components, the overall function is not linear in travel time.

As in the case of the mean vs. variance approach, empirical evidence suggests that all components of utility are ‘bad’. If we now admit variability in travel time, then we can once again use the ideas of expected utility maximisation (but NB not assuming a vNM utility function), taking expectations of each component of utility over the travel time distribution.

$$E[U] = \alpha E[T] + \beta E[SDE] + \gamma E[SDL] + \delta E[L] \quad (2.4)$$

where $E[L]$ may be re-interpreted as the probability of lateness.

Because of the discontinuous nature of the terms SDE, SDL and L, this expectation is **not** the same as would be obtained from inserting the mean value \bar{T} in the variable definitions: U is not linear in T.

2.4 Equivalence between the mean vs. variance approach and the scheduling approach

Since the terms $E[SDE]$ and $E[SDL]$ embody variability in travel time, albeit conditioned by the notion of the PAT, it has long been suggested that there exists an (approximate) linear relationship between the summation $\beta E[SDE] + \gamma E[SDL]$ and the standard deviation of the travel time distribution σ . Strictly speaking, however, this approximation relies on departure time being continuously variable (as with the car mode, or with high-frequency public transport perhaps), and is closer for some distributions (e.g. exponential, uniform) than for others.

Until recently, a complete account of the conditions relating to this equivalence had not been demonstrated. Using the approach of Noland & Small, Polak (1995) had demonstrated it theoretically for the exponential distribution, when the optimum travel time is chosen. Appendix B of Arup Deliverable 6.1 (see Annex) takes this a little further, in also considering the uniform and logistic distributions, as well as one based on two “mass points”. It is noteworthy that the equivalence appears to be dependent on the form of the distribution.

In an important new paper, Fosgerau (2007) has demonstrated both the generality of the equivalence, and the dependence on the distribution of travel time. In what follows, we present his general argument.

For simplicity, we drop the “ δ ” term and assume that the travel time distribution does not depend on the departure time⁴. **Whatever** the distribution $f(T)$ is, we can represent the random variable T by a combination of the mean and variance, using Z as a standardized random variable with mean 0 and variance 1: in other words we write $T = \mu + Z \cdot \sigma$ and write $\phi(Z)$ as the distribution of the standardized variable. It can then be shown that the optimum departure time is given by

$$PAT = \left(\mu + \sigma \Phi^{-1} \left(1 - \frac{\beta}{\beta + \gamma} \right) \right)$$

where Φ^{-1} is the inverse cumulative distribution of Z.

Substituting this into the utility function, it can be shown that we obtain an optimum utility equivalent to:

⁴ In Annex I it is shown generally how these restrictions can be relaxed

$$\alpha \mu + \sigma \cdot \left\{ (\beta + \gamma) \cdot \int_{\left(1 - \frac{\beta}{\beta + \gamma}\right)}^1 \Phi^{-1}(y) \cdot dy \right\}$$

Hence, at least under the restrictions stated, the reliability ratio $\tilde{\rho}$ can be written as:

$$\tilde{\rho} = \left\{ \frac{\beta + \gamma}{\alpha} \cdot \int_{\left(1 - \frac{\beta}{\beta + \gamma}\right)}^1 \Phi^{-1}(y) \cdot dy \right\}$$

This makes it clear that, for a given distribution ϕ , holding the mean constant, the impact on the optimum value of the schedule delay terms in (2.3) is linear with changes in the standard deviation. For fixed values of α , β , and γ , the reliability ratio is fixed, **but** it is dependent on the distribution function ϕ ⁵.

Fosgerau calculates the integral in the above formula, both for the normal distribution and for an empirical distribution which appears to fit the data for a particular road in Central Copenhagen. Using the commonly quoted Small (1982) coefficients for α , β , and γ gives a value for the reliability ratio of 0.84 for the normal and somewhat higher (0.93) for the empirical distribution. As we will see, these are close to values which have been reported in the literature (eg 0.8).

It is noteworthy that the Arup work attempted some straightforward simulation work which changed the nature of the travel time distribution, and demonstrated that in some circumstances the implied benefits could be very different according to whether they were computed using the scheduling formulation or the mean vs. standard deviation formulation. This bears out Fosgerau's conclusions about the dependence on the distribution.

It will be noted that the equivalence between the two approaches discussed here depends on the ability of travellers to continuously vary their departure time. This will in general not apply to public transport.

2.5 Applying the theory to the valuation of reliability

Typically we wish to apply the above theory to the (monetary) *valuation* of reliability. In order to do this, the convention is to supplement the expected utility functions (2.1) and

⁵ It should also be noted that while for some simple distributions, and for the normal distribution, the cumulative distribution of the standardized distribution is independent of the mean and the standard deviation, so that the reliability ratio is fixed, this is not generally the case. For example, with both the gamma and lognormal distributions, Φ^{-1} is a function of the coefficient of variation (σ/μ). Hence, for a given μ , variations in σ will lead to different values of Φ^{-1} . Hence the reliability ratio will not be constant. What is required is an understanding of how serious this is over the likely range of variation in σ . This is not a criticism of Fosgerau's mathematical treatment, but it means that the interpretation that can be placed on it is more limited than it at first appears.

(2.4) with travel costs, and to derive valuations of reliability as the marginal rates of substitution between cost and reliability variables. As already noted, in practice the mean vs. variance approach in fact usually considers standard deviation rather than variance (perhaps because standard deviation is measured in time units). If we expand (2.4) thus:

$$E[U] = \alpha \bar{T} + \lambda \sigma + \theta C \quad (2.5)$$

where C represents travel costs. (One might also wish to supplement with other variables depending on the practical context and policy interests, for example interchange and crowding.)

Then the marginal valuation of mean travel time (“value of time”) is given by α/θ , and the marginal valuation of the standard deviation in travel time is given by λ/θ . The ratio of these valuations is, of course, the ‘reliability ratio’ $\tilde{\rho} = \lambda/\alpha$.

Turning to the scheduling approach, we again supplement expected utility with travel costs as follows:

$$E[U] = \alpha E[T] + \beta E[SDE] + \gamma E[SDL] + \delta E[L] + \theta C \quad (2.6)$$

Then we can derive marginal valuations for each of the time and scheduling components, thus the marginal valuation of expected travel time is given by α/θ , the marginal valuation of expected schedule delay early by β/θ , the marginal valuation of expected schedule delay late by γ/θ , and the marginal valuation of the probability of lateness by δ/θ .

2.6 Challenges and opportunities for further theoretical research

Translating the theory to public transport

The theory as outlined above appeals particularly to the car mode, where it may be assumed that departure times are continuously variable. By contrast, public transport departures are often governed by fixed timetables, and this induces a discrete set of possible departure time choices. Bates et al. (2001) devote attention to this challenge, as does Batley (2007) from an alternative standpoint, but further work is probably needed. Such work would be usefully guided by an aspiration to harmonize the theory across modes, thereby permitting application to multi-modal analysis. Indeed there would be obvious attraction in proposing a single theoretical approach that is applicable to all modes.

There is evidence (eg Bates et al) to suggest that in addition to scheduling costs with respect to PAT, , public transport users may also take into account:

- Punctuality (i.e., compliance with the published schedule) which will in general be correlated with but not the same as the classical notion of schedule delay
- Pure variability (i.e., even if TTV does not make them late, they still may not like it)

Thus both mean-std and scheduling models may in fact be special cases of a general model.

The interpretation of lateness, and alignment with industry conventions

Following from the above point, and in particular the aspiration to harmonize theory, it would seem fundamental to seek greater consistency in how our notion of ‘lateness’ applies to different modes. In the empirical application of Bates et al. (2001), which relates to the rail sector, the scheduling model is supplemented with an additional variable representing mean delay (i.e. the *operators’* notion of lateness, with reference to timetable). In many cases it appeared that the estimated coefficients of the lateness [mean delay] variable and $E[SDL]$ were not significantly different. This perhaps provokes the question of how public transport travellers interpret the PAT, and in particular whether they align the PAT with timetabled arrival times.

If the PAT and timetabled arrival time are one and the same, then this would seem to bring some convenience to the theory, as well negate the need to undertake the notoriously difficult activity of eliciting PATs from travellers. More generally, can the scheduling approach be re-couched without the need to explicitly specify the PAT? It might be noted that, in terms of modelling choice of time to travel, both PRISM and the Dutch national model operate in this vein, representing the time dimension in terms of large discrete periods (such as a two-hour peak period), and modelling shifts between periods as a result of changes in travel time and cost. Nonetheless, more thought would be required before it could be proposed for the investigation of reliability.

In order to illustrate the impact of unreliability in the public transport (in point of fact, rail) case, Bates *et al* made certain assumptions about the pattern of delays, and estimated the expected schedule disutility, conditional on the PAT. They concluded that, at least for low levels of variability, the impact of variability is dominated by the mean delay, especially for less frequent services (say, fewer than three trains per hour). This is in reasonable agreement with the existing practice of valuing unreliability for rail. Nonetheless, more generally the impact is sensitive to the distribution of delays, not merely the mean. This accords with intuition. In addition, they conclude that the actual distribution of PAT with respect to the timetabled arrival is potentially of great importance. Clearly, the impact on schedule delay of unreliability will be much greater when PATs are concentrated around the timetabled arrival. This is another area where more research is needed.

The work by Bates *et al* makes it clear that, with some effort, it is possible to apply a scheduling approach to the investigation of travel time variability on rail, but that the outcome is affected by a number of factors, not all of which are well understood. In spite

of the difficulties, it appears that the scheduling approach offers a greater chance of developing a harmonised approach across modes than does the mean-variance approach.

Valuing travel time variability, as distinct from valuing time risk

The distinction made earlier between the two types of utility function (RUT and vNM) and the possible confusion in some cases means that there is an unresolved debate as to whether the marginal valuations of reliability derived in the following section 3 encompass the full range of costs that unreliability imposes on the traveller. It should be noted that these valuations - whether derived from the mean vs. variance approach or from the scheduling approach - are strongly dependent on the form of the expected utility function. Acknowledging the earlier remark in section 2.1 concerning the analogy between uncertainties in money and travel time, it would seem instructive to compare marginal valuations of reliability with the conventions of the mainstream economic literature of choice under uncertainty.

Such a comparison provokes the suggestion that these marginal valuations, whilst adequately representing the impact of travel time variability on utility and choice, omit the cost of risk-bearing imposed on the traveller as a result of this variability, which may be referred to as the *risk premium*. If travellers are inclined to 'insure' against travel time risk, in the same manner that they routinely insure against money risk, then there is the potential for a risk premium.

Economic appraisal should, in principle, include any risk premia relevant to economic behaviour, and omission could potentially introduce bias. Further theoretical work is required to formalise the proposition of a risk premium for travel time, as well as empirical work to investigate its prevalence and magnitude (whilst this will likely require fresh data collection, it would be worth considering the potential for re-analyzing existing data).

This point is discussed at some length in Liu and Polak (2007) and also in Richard Batley's recent work (Batley, 2007). The fundamental issue is the following. When we consider choice under uncertainty in the context of scheduling costs we are dealing with a non-linear system which, as demonstrated, induces a value of travel time variability. However, we can parameterise this system in different ways. In particular we can load explanatory responsibility either on the taste parameters that we are familiar with from choice contexts with **no** uncertainty, or we can add additional parameters that seek to characterise tastes with respect to risk, or we can do both. Each approach has different implication regarding what we mean by the value of variability. As noted, Liu and Polak have shown that models with an explicit characterisation of tastes with respect to risk out-perform models without such a characterisation.

Extension of choice dimensions

According to the above theory, the only available response to unreliability is to adjust departure time. In the context of the present commission, it would seem pertinent to

consider whether a broadening of scope to include (at least) mode choice as well as departure time choice issues any challenge to the theory. There are examples from practice of joint mode choice-departure time choice models, but do they stand up to theoretical dissection? One might argue that mode choice can be straightforwardly included through the time, cost and scheduling variables. As we have already noted, however, some modes are constrained by timetabling whilst others are not, and this could reasonably feasibly introduce complications. Indeed, for any given traveller with a particular PAT, improvement in the reliability of one mode might yield benefit whilst improvement in the reliability of another mode may incur cost. Further complexities perhaps arise from multi-stage journeys involving various modes; in this regard Bates et al's (2001) discussion of interchange offers a useful starting point.

Ultimately, of course, if we are successful in constructing a method for including travel time variability within some kind of "generalised cost" metric, then it will be necessary to address the same set of issues about choice hierarchy as are currently faced in constructing demand models. The discussion of mode choice was only singled out on the basis that a) some analysis of joint choice processes with departure time choice has already been carried out in this case, and b) there are different factors attending the departure time choice process when we are dealing with scheduled services, rather than continuous flexibility.

Heterogeneity in valuations and PATs

Finally, we should note that while the theory is typically expressed in terms of an individual traveller with a particular PAT, in practice it is necessary to take account of variations in PAT among the travelling public, as well as "tastes" relating to scheduling disutilities. We can expect that some segments will place a higher value on reliability than other segments, hence the call for methods that support a segmented analysis.

2.7 Recommendations on theoretical approach

Two versions of the theoretical paradigm to representing travel time variability have been presented here, the mean vs. variance approach, and the scheduling approach. In both cases, a strong recommendation is that further work is needed to progress the theoretical basis and to test the hypothesized model forms.

In our view, it is not currently possible to choose between these two versions without a clearer understanding of their similarities and differences. Certainly we concede that the mean vs. variance approach, together with the associated metric of the 'reliability ratio', has the advantage of greater simplicity, which gives it some appeal. This should not however detract from our belief that the scheduling approach may reveal important, and unique, insights into behavioural responses to unreliability, particularly in the context of scheduled (but infrequent) public transport services. Moreover, if theoretical methods are to permit meaningful comparison across modes, then any remaining challenges to the

correspondence between the mean vs. variance and scheduling approaches must be resolved.

3. Evidence on Valuation of Travel Time Variability

3.1 Introduction

This chapter deals with the question of finding monetary values for variability, discussing both passenger transport and freight, with more emphasis on the passenger side. Empirical outcomes are given in tables at the end of this chapter.

The key distinctions, which comprise the two main strands discussed in the theory chapter, used in this chapter are listed in Table 3.1.

Table 3.1 Key distinctions in empirical studies on valuing travel time variability

	PAT obtained	PAT not obtained
variation presented	A	B
variation not presented	C	D

If in empirical studies variation in travel time is not presented and PAT (preferred arrival time) not obtained (D) then the data cannot be used for “reliability” purposes, so this case can be dropped.

However, scheduling parameters can be obtained whether or not variability is presented, as long as we have the PAT (or some reflection of it). We can then derive the response to variability by confronting the schedule delay formula with a distribution of travel time. These are cases A and C, which follow the schedule delay approach discussed before.

If PAT is not obtained, then clearly we have to resort to mean vs. variance approach (case B). We can also estimate mean/variance for case A.

As background to the work reported here and recommendations, however, it should be noted that practically all the empirical work that has been done to obtain values for variability, using either the mean vs. variance or the scheduling approach or both, has been based on Stated Preference (SP) data. It is generally very hard to collect Revealed Preference (RP) data that includes measures of variability, travel time and travel costs that will not be heavily correlated. Also, with RP data, there is the perennial difficulty of getting information on the attributes of the non-chosen alternatives, e.g. on the travel times at different moments (periods) in time. Nevertheless, researchers have tried to find situations with variability variation (e.g. the choice between two car routes where one is less reliable because of more congestion or bridges that might be closed) or between two train services. Examples of practical RP studies known to us are the ones carried out in California (SR91) comparing a route with a variable toll to an untolled (uncongested) route, and an ongoing study by ITS Leeds on rail time variability.

A few studies (e.g. Bates et al. (2001); Copley et al. (2002); Noland et al. (1998); Hollander (2005); Significance et al. (2007)...) have presented both a measure of variation of travel time and obtained a PAT (case A in Table 3.1). Most empirical studies

have either used the mean vs. variance approach (case B) or the schedule delay approach (case C). The two dimensions noted in Table 3.1 both relate to significant difficulties for data collection. For cases A and B, there is the issue of presenting variability to respondents in a way which gives them sufficient assistance to provide reliable and useful data in terms of the choices offered. For cases A and C, a major difficulty is how to obtain the PATs. We discuss the difficulties relating to these two dimensions respectively in sections 3.2 and 3.3.

3.2 Presentational issues relating to variability (passengers)

In the context of travel time variability research, models often look jointly at valuations for the mean journey duration along with the valuation for the variance (or the standard deviation) of this journey duration. While this poses no issues as such from a modelling perspective, it is generally recognised that a non-trivial part of the population have difficulties in understanding the concept of the variance of journey duration in an SP survey. Aside from general design questions such as the number of alternatives and attributes, along with the attribute levels, the main issue that needs to be addressed in the survey design phase is the decision as to what approach should be used in the presentation of travel time variability in the survey questionnaires. A number of different possibilities arise.

The first possibility is to actually present respondents directly with a measure of travel time variability, such as for example the standard deviation in the travel time across a number of trips. The closest example of such an approach is the inclusion of travel time variability as an actual attribute in the surveys used by Hensher and colleagues in Australia. As an illustration, Figure 3.1 shows a screenshot from such a survey undertaken in Sydney, where this is taken from Hess et al. (2006).

As mentioned above, it is often argued that such direct measures of variability are difficult to understand for a non-trivial portion of the population. This is partly reflected in the low levels of significance obtained for such attributes (cf. Hess et al. (2006)), and also in the low relative valuations when compared to travel time (cf. Hensher, (2007)).

A very basic approach consists of presenting respondents with the probability of a certain journey being affected by issues with variability. An example of this approach is given in Figure 3.2, for the choice between two car routes in the context of the mobility pricing study undertaken by Vrtic et al. (2006). Here respondents face a choice between a tolled and an untolled alternative, where for the latter, in this case, one out of every 20 journeys takes at least 10 minutes longer than scheduled.

Sydney Road System

-Practice Game-

Make your choice given the route features presented in this table, thank you.

	Details of Your Recent Trip	Road A	Road B
Time in free-flow traffic (mins)	50	25	40
Time slowed down by other traffic (mins)	10	12	12
Travel time variability (mins)	+/- 10	+/- 12	+/- 9
Running costs	\$ 3.00	\$ 4.20	\$ 1.50
Toll costs	\$ 0.00	\$ 4.80	\$ 5.60

If you make the same trip again, which road would you choose? Current Road Road A Road B



If you could only choose between the 2 new roads, which road would you choose? Road A Road B

For the chosen A or B road, HOW MUCH EARLIER OR LATER WOULD YOU BEGIN YOUR TRIP to arrive at your destination at the same time as for the recent trip: (note 0 means leave at same time) min(s) earlier later

How would you PRIMARILY spend the time that you have saved travelling?

Stay at home Shopping Social-recreational Visiting friends/relatives
 Got to work earlier Education Personal business Other

Figure 3.1: Example of survey including travel time variability as an attribute (taken from Hess et al. (2006))

	 Route 1	 Route 2
Abfahrtszeit	16.45	17.00
Fahrtzeit (min)	60 min	40 min
Treibstoffkosten für die Fahrt	6,40 CHF	3,20 CHF
Art der Strassengebühr		Für eine zusätzliche Mautspur
Höhe der Strassengebühr		4,00 CHF
Verlässlichkeit	Verspätung grösser 10 Minuten: etwa jede zwanzigste Reise	keine
Ankunftszeit	17.45	17.40

← Ihre Wahl →

Figure 3.2: Example of a survey design presenting respondents with the probability of a delayed journey (taken from Vrtic et al. (2006))

In the example shown, respondents are not given any information on the likely delay, although an indication of the lower limit on any delays is given (i.e. 10 minutes). In fact, studies often do not even give this information, and are based solely on the probability of a delay, or indeed the probability of arriving on time, such as in the example in Figure 3.3, taken from König (2004).

Ich fahre mit	der Bahn		<input type="checkbox"/>
	Dauer:	90 Minuten	
	Pünktlichkeit:	100%	
Ich fahre mit	dem PW die normale Strecke		<input type="checkbox"/>
	Dauer:	45 Minuten	
	Pünktlichkeit:	50%	
Ich fahre mit	dem PW den grossen Umweg		<input type="checkbox"/>
	Dauer:	90 Minuten	
	Pünktlichkeit:	100%	

Figure 3.3: Example of survey providing respondents with a probability of arriving on time (taken from König (2004))

In this example, the choice is between travelling with the train (“mit der Bahn”) with 100% certainty of punctuality (“Pünktlichkeit”), or by car (“PW”) either along the normal route with a 50% expectation of arriving by the stated time, or by a major deviation (“Umweg”) with a much longer time but with 100% certainty.

The approach of presenting respondents with a probability of a delayed journey can be extended straightforwardly by additionally providing them with a measure of the likely delay (when not zero), where this then again allows for a calculation of some form of travel time variability. An example of such an approach is given in Figure 3.4, again taken from König (2004). Here, respondents are presented with a choice between a slower and a faster alternative, where the latter (route B – a bus-rail connection) gets delayed by 30 minutes on 2 days per week (owing to the delays on the bus, which mean a missed connection). In this case, there is no further uncertainty as to the length of the delay, where such uncertainty can be incorporated by providing respondents with an anticipated mean delay in addition to the probability thereof, rather than a fixed delay if and when such a delay occurs.

Fahrzeit Route A: 50 Minuten

Fahrzeit Route B: 45 Minuten

Aufgrund Ihrer Erfahrung **verpassen Sie an 2 Tagen der Woche den Anschluss**, weil der Bus unpünktlich ist.

Dann benötigen Sie für die Route B 75 Minuten.

Aber das wissen Sie ja erst, wenn Sie bereits unterwegs sind. Welche Route wählen Sie für diesen täglichen Weg?

Antwort:

Ich wähle Route A.

Ich wähle Route B.

Figure 3.4: Example of survey presenting respondents with probability of a delay and the extent of this delay

A different approach is to present respondents, for each alternative, with a set of possible travel times. This goes some way towards presenting the complete travel time distribution to the respondents, which from a research perspective would be the “ideal” representation of travel time variability. Often, this takes the form of presenting respondents with a set of possible journey durations (say 5 to 15) within each choice alternative, where this information is sometimes presented graphically. The researcher can then calculate the variance that is consistent with each set (or, more likely the other way around, generate a set of journey durations that matches a target variance). Average journey time and the variation in travel time presented in the SP survey can be chosen with a minimum level of correlation. Alternatively, respondents may be presented not with the actual travel times but with the early or late arrivals (compared to schedule) for a number of services. This is

for example the approach used in the work of Bates et al. (2001), where, as in many other studies of this type, a graphical approach ("clock-face") is used to present the possible delays (cf. Figure 3.5). Here, respondents are told that each one of the possible outcomes has the same probability.




Figure 3.5: Example of a survey presenting respondents with a number of possible delays for each alternative (taken from Bates et al. (2001))

The question invariably arises whether this method of presentation is not in fact more demanding for respondents than an approach giving a direct measure of travel time variability. As such, with the example in Figure 3.5, there is clearly a possibility that respondents do not in fact evaluate the alternative on the basis of mean delay and variation in the delay. Rather, they might simply work on the basis of the mean delay across the ten scenarios (e.g. 21.7 minutes late for operator A), or the probabilities of early and late arrival (e.g. 20% vs 80% for operator A).

In the Netherlands, a major empirical study commissioned by the Dutch Ministry of Transport is underway to measure the value to society of travel time benefits and travel time variability benefits in passenger and freight transport (see: Significance, VU University Amsterdam and John Bates, (2007)). The values will be estimated using Stated Preference experiments. In the past years several researchers have designed formats to present unreliability of travel times to respondents in SP experiments. These formats use concepts from statistics like "average travel time", "travel time variance", "probability of arriving late", etc. To test whether these fairly advanced concepts are understandable for laymen travellers without a degree in statistics or even higher education, in-depth interviews among 30 respondents were carried out. The objectives of these (face-to-face) interviews were:

- Test the respondents' understanding of different reliability presentation formats;

- Investigate the respondents' assessments of these presentation formats with respect to clarity, ease of handling, and visual attractiveness;
- Collect the respondents' preferences of the presentation formats.

In the analyses eight formats of presenting travel time variability were tested. All eight relied on giving five possible travel times within a single choice alternative, but differed in the wording and the use of graphics (e.g. bar charts, clock-face presentation as in Bates et al. 2001). Respondents were stratified according to their education level. There were three groups of questions. First, questions about how respondents conceptualise unreliability themselves. Do they think in terms of average, minimum, maximum travel time or probability? And how complicated do they find these concepts? Secondly, respondents were prompted with questions to test whether they gave the “right” answer for the different presentation formats. These test questions were designed to check to what extent the respondents have the “correct” perception of reliability, i.e., the same as expected by researchers, for each presentation format. Thirdly, there were questions about the respondents' assessments of the eight presentation formats regarding clarity, ease of handling, and visual attractiveness, and which ones were preferred.

The interviews supplied a clear “winner” (see Figure 3.6) among the eight formats, which was the format without any graphical presentation, but just a list of five equi-probable travel times. This format is not only preferred by a majority of respondents, but also equally by people with low and high levels of education. What matters, however, is not what the respondents prefer but how well they understand the information that is presented. This format (with the corresponding arrival times added) will be used in the forthcoming main study.

In deze versie tonen we u de 5 mogelijke reistijden onder elkaar.
Stelt u zich voor dat u met de auto naar een winkelcentrum wilt reizen. U kunt kiezen uit twee ritten: A en B. Welke zou u kiezen?

<i>Rit A</i>	<i>Rit B</i>
Gemiddelde reistijd: 40 min	Gemiddelde reistijd: 41 min
U heeft een even grote kans op elk van deze 5 reistijden:	U heeft een even grote kans op elk van deze 5 reistijden:
35 min	30 min
40 min	35 min
40 min	45 min
40 min	45 min
45 min	50 min
Kosten: €3,80	Kosten: €2,80

A

Figure 3.6: SP presentation format that worked best in Dutch trial survey (taken from Significance et al. (2007))

3.3 Derivation of PAT for scheduling models (passengers)

There is a substantial body of work on the use of scheduling models (see the theory chapter), some of which yields empirical evidence on the value of travel time variability. The data used are often of an SP nature. A practical difficulty is that the model specification requires that the respondent elicits his or her preferred arrival time (PAT). It is by no means straightforward to ask this in an SP survey in a way that will correctly define what the researcher needs and will also be understood as such by the majority of respondents.

Some of the early scheduling studies (e.g. Small (1982)) have used the (fixed) start of the factory/office working hours as the PAT for commuting. Agreed starting or delivery times can also be used for (large sections of) business travel and freight transport, but usually do not exist for other travel purposes. Moreover, for commuting this approach has become less appropriate because of the increasing share of workers with flexible working hours.

An example of how questions into the PAT are phrased in a questionnaire survey (and how complicated things become) can be found in Significance et al. (2007):

1. The travel time of your trip that day was **TRAVEL TIME** minutes. How long would this travel time have been without any traffic jams or other delays?

_____ minutes

2. Please keep thinking back to the trip you made that day. What would have been your preferred departure time if you had known **with absolute certainty** that there would be no traffic jams or other delays (so that the travel time would have been **ANSWER QUESTION 1** minutes)?

_____ : _____ (use 24 hour notation, for example 15:45)

CALCULATE PREFERRED ARRIVAL TIME FROM QUESTION 2 AND QUESTION 1

3. This implies that your preferred arrival time would be **PREFERRED ARRIVAL TIME**. Is this correct?

- yes
- no, I would have liked to arrive at _____ : _____

Furthermore, the monetary values obtained for being early or late are very difficult to implement in a standard cost-benefit framework using the standard transport models, because the link to travel time period choice is not made in the such analysis (there is no reference to clock time, only to journey durations), and the preferred arrival times are unknown. This however is a modelling and implementation issue (see next chapter), not a problem in valuation itself.

3.4 Issues in valuation studies for freight transport

The presentational issues are also relevant for freight transport. Just as in passenger transport, the concepts of variance and standard deviation are also often considered as too difficult for the respondents in a freight context, i.e. the shippers and carriers. Most studies use the probability of delay or the percentage of not on time arrivals instead.

In freight transport, the probability of delay is often measured as the probability of not arriving at the specified (by the shipper/receiver) time or within the specified time interval. This is then used as an equivalent of the PAT or preferred arrival time interval of passenger transport. The schedule delay could also include being too early, which leads to extra costs at the destination. Therefore, the outcomes are related to those of the scheduling approach. An explicit application of the scheduling approach to freight transport has also been undertaken by Small et al. (1999), and by Fowkes et al. (2001).

The values for variability in freight transport from the different studies are difficult to compare, because of the differences in the measurement units.

A specific difficulty in valuation studies for freight transport is who should be interviewed. Several agents are involved in decision-making on the same shipment (shippers, carriers, third party logistics service providers, truck drivers, ...). Massiani (2005) gives a theoretical argument that shippers will only give the VoT and travel time variability value of the cargo and carriers will include all elements (including costs for drivers, vehicles, cargo). We are not sure whether this argument will hold in a practical SP (which could also be seen as a game between shippers and carriers, with agents acting on how they would expect the others to act). Significance et al. (2007) makes the following assumptions (a priori hypotheses) on the freight VoT. Shippers with own account transport can give information on both the cargo-related and the vehicle/staff-related elements, shippers that contract out provide values for the first element and carriers for the second element. Of course there may be exceptions to the above general pattern, but in the freight questionnaires the shippers that contract out are steered (by very explicit instructions) to only answer on the components they generally know most about, and likewise for carriers.

There has also been some work in Australia on the valuation of variability in the area of freight, by Hensher et al. (2007), where they develop a framework that captures the interactive element of choice by incorporating ideas of concession and power for shippers and carriers, making use of mixture models (also see Puckett and Hensher, (2006); Puckett et al. (2007); Paglione et al. (2007)). In the survey, respondents are presented with the probability of an on-time arrival for the shipment, along with a host of other attributes.

3.5 Empirical findings for passenger and freight

It must be accepted the outset that the results which have been presented in individual studies are determined to a considerable extent not only by the way in which the data has been collected but also by the way in which it has been analysed. Therefore considerable care is needed in comparing the results of different studies. In almost all cases, a RUT approach has been taken to the analysis. With the exception of the recent work by Michea & Polak (2006) and Liu & Polak (2007), we are not aware of much work in this area which has explicitly specified a vNM utility function. A recent paper by de Palma and Picard (2005) did this, though not quite in the same way or in the same context. Other recent relevant work is by Recker et al (2005):

A key issue is whether the specification used in the models (in terms of the variables included in the utility function) is the same as that used in the presentation of the choice situations to respondents. It is of course important, in model specification, to stay as close as possible to the scenarios that respondents were actually faced with. In the case of data where respondents were presented, for example, with the probability of a delay, it is straightforward to include this directly in the same way in the models. This also still

applies when respondents are additionally faced with the average duration of any delay. As such, notwithstanding issues of respondent understanding, for the scenarios presented in Figures 3.1 to 3.4 it is possible to use the same specification in the models as in the survey data. The attributes relating to variability would be treated in the same way as the attributes relating for example to travel cost and travel time, typically using a linear-in-parameters approach.

The situation however becomes more difficult in the case of a design such as that used in Figure 3.5. Here, respondents are presented with a number of different outcomes, and it is clearly inappropriate to simply include the various outcomes in the utility function, each with their own estimated coefficient. Rather, we need to attempt to find a specification of the utility function that replicates the approach used by respondents in evaluating the various possible outcomes.

The number of possible approaches in this case is very large indeed, and the choice of an approach is difficult. There is a need to find an approach that attempts to do justice to the possible approaches used by respondents, but also one that is practical for use in modelling. The other issue is that by choosing a specific modelling approach this acts as a very strong assumption, imposing a certain evaluation tool on the data, where this may in fact differ significantly from the approach used by respondents. The other problem is that already at the stage of survey design, researchers might have an idea of the factors they are interested in, say for example the mean and standard deviation of travel time. They then produce a proxy for this information that they deem suitable for presentation to respondents and base their analysis of the observed choices on the original concept, say an evaluation of mean and standard deviation in travel time. When the assumed method however differs from that actually used by respondents, misleading results can be produced. This risk is potentially higher in the case of more complex specifications, such as percentiles or highly non-linear methods that may be difficult to translate into a simple set of equally likely outcomes. In short, the issue is to reconcile a mathematical specification of variability at the modelling end with the presentation of such information at the data end.

A summary of the main findings from the studies reviewed is presented in Tables 3.2 and 3.3. This includes quantitative outcomes, the method used in the study and a brief description of other relevant issues. Monetary outcomes from various empirical studies have been converted to Euros of 2003 (using international exchange rates and consumer price index numbers); some outcomes were in the original reports and papers only given in terms of minutes of travel time, and these were kept as such. It is clear from tables 3.2 and 3.3 that a very wide range of studies have been undertaken with outcomes which are not necessarily presented in a comparable format. Where comparable results are given there is still some diversity in the empirical values obtained. There are some apparent reasons for this, including the nature of the study, the sample taken and the level of disaggregation in the findings. This diversity in the evidence body is reflected in the overall recommendations outlined in section 3.5 below.

3.6 Recommendations on valuation

The review of empirical evidence on values for variability has highlighted the existence of a growing body of work on the modelling of the valuation of travel time variability. It should be acknowledged that this review is by no means complete. This is not helped by the fact that some of the existing work is not available in English, or is indeed not publicly available at all. As such, it can for example be assumed that there is a substantial body of work in the private sector, for clients such as railway companies, a point alluded to in König (2004).

While all reviewed studies agree that variability is a factor of substantial importance, there are no generally accepted monetary values for variability, or indeed a reliable estimate of the relative weight of travel time and travel time variability. It should be acknowledged that the valuations of variability in the present literature come from very specific investigations and are not even used in cost-benefit analyses in the respective countries of origin. Often, the studies are relatively small scale, and in some cases, variability is not the main topic of investigation.

The primary recommendation from this piece of work would thus be that there is a need for a major new study into the valuations of variability. Such a study would have to investigate both the question of how information on variability is presented to respondents and how it is later specified in econometric models that are estimated on the data. However, prior to this, agreement needs to be reached on the theoretical framework for journey time variability (as discussed in the previous chapter). The aim should be to achieve as high a level of consistency as possible between the underlying theory, the data collected and the estimated model, without unduly affecting understanding by respondents or the value of model outcomes. The a priori assumptions as to how respondents evaluate the information in the surveys should be kept to a minimum. Such work should also look further into the benefits of advanced modelling techniques such as discussed by Michea & Polak (2006) and Liu & Polak (2007).

Table 3.2. Value of variability (in travel time or Euros of 2003) in passenger transport: quantitative outcomes, methods used and other lessons.

Study	Quantitative outcomes (+definition)	Method	Other lessons
Accent and HCG, (1995)	Doubling the chance of delay is equivalent to 13 min. travel time (commuting) or 20 min. (business and other travel); halving the chance of delay: 3 min. (commuting) or 5 min. (business and other travel).	Stated preference (SP) in road transport in the UK, with the following attributes: travel time, provision of information and chance of delay.	For some segments (e.g. business travel, time gains) the value of travel time per minute is higher in congested than in uncongested conditions.
ATOC, (2002)	Reliability ratio for train is in the range 0.6-1.5; most values in the range 0.8-1.3	Literature review as part of Passenger Demand Forecasting Handbook, UK	
AVV, (2003)	Reliability (operating on time) is the most important aspect (importance of 3.58 on a scale from 1 to 5, with 5 being best) for bus, tram and metro, and the actual reliability performance is regarded as mediocre (5.94 on a scale from 1-10, with 10 being best).	SP among 3,387 users of bus, tram and metro in The Netherlands.	
Bates et al. (2001)	Found significant valuations placed both on the inherent variability in travel time and (for scheduled services) on schedule compliance. Value of expected late schedule delay twice that of expected early schedule delay.	SP amongst 200 rail travellers	Underlying theory and recommendations for empirical research: scheduling model and SP data. Dealt with both unscheduled (car) and scheduled (public transport) contexts.
Brownstone and Small, (2002)	Value for 90 th minus 50 th percentile of the transport time distribution: 11-14 Euro/hr (males) and 28-30 Euro/hr (females).	RP: travel time measurements on State Route 91 in California, with variable tolls.	Method: can be done with RP data (in special cases such as this), use of percentiles.
Brownstone and Small, (2002)	Value for 80 th minus 50 th percentile of the transport time distribution: 26 Euro/hr.	RP (see above) and SP.	Travel time accounts for two-thirds of the service quality differential between the tolled and the alternative route; reliability one-third.
Copley et al. (2002)	The value of the standard deviation of travel time is 1.3 times the value of travel time (both per minute).	SP among 167 car drivers commuting in Manchester; mean versus variance method.	
Copley et al, (2002)	1 minute late or early are valued less than 1 minute travel time.	SP (see above); scheduling model.	Method for valuation: mean versus variance approach or scheduling model.
Eliasson, (2004)	Reliability ratio of 0.95 for commuting 0.3 for business travel and 0.59 for other purposes (all car)	SP among 600 car drivers, Sweden	
Hensher, (2007)	Value of travel time savings much higher than value of travel time variability	Travel time variability presented directly in SP	
Hollander, (2005)	Bus: mean travel time and early arrival valued at 8 Eurocent/minute; late arrival at 22 Eurocent/minute	SP among 244 bus users in York; scheduling model	mean-variance method gave no significant variance coefficient

De Jong et al. (2003)	Commuting, business and leisure travel: 1 minute late or early is 1-1.5 times as bad as 1 minute travel time; Education: 1 minute late or early is less important than 1 minute travel time; All purposes: 1 minute longer or shorter participation in activity at destination is less important than 1 minute travel time.	SP among around 1,000 car drivers and train users in the peak periods in The Netherlands; Scheduling model.	
König, (2004)	Valuation of 34CHF for delay of 60 minutes, slightly lower for PT, 20% higher than value of time.	Various SP designs presenting direct measure of variability.	Important to look separately at probability of delay and length thereof.
Liu and Polak, (2007)	Value of risk aversion parameter is approximately -0.2, indicating a moderate level of risk aversion in the sample. Addition of risk aversion specification does not materially affect other model parameter valuations compared to Bates et al. results.	SP amongst 200 rail travellers (data from Bates et al.)	Proposes a (genuine) expected utility scheduling model with explicit characterisation of travellers' attitudes towards risk. These model out-perform existing (expected value) models
Michea and Polak, (2006)		SP amongst 200 rail travellers (data from Bates et al.)	Proposes a number of non-expected utility models including various forms of reference point model. These models offer only moderate improvements compared to conventional models.
MVA, (1996)	Ratio of value of standard deviation of travel time to value of in vehicle time for car: business travel: 0.36, commuting and other: 0.78	Literature review, UK	
MVA, (2000)	The value of the standard deviation of time in the bus is 24% of the value of travel time in the bus (when seated; less for travel time standing; both measured in minutes). The value of the standard deviation of waiting time is 48% of the value of waiting time.	SP among 309 bus users in France; Mean versus variance approach.	
Noland and Polak, (2002)		Literature review of empirical and theoretical work up to late 1990s.	Method: scheduling model or mean versus variance approach (both in combination with SP).
Polak et al. (2008)	Similar value of risk aversion parameter to those found in Liu and Polak (2007). Evidence of significant observed and unobserved heterogeneity in risk aversion parameter. Observed heterogeneity related to traveller demographics.	SP amongst 200 car and rail travellers.	Refinement of the work of Liu and Polak (2007) to accommodate heterogeneity in attitudes to risk.
Rietveld et al. (2001)	A decrease in the probability of a 15 min. delay from 50% to 0% is worth 2.35 Euro (30% of the value of an hour travel time). A reduction in the probability of a 2 min.	SP among 781 public transport users in The Netherlands, with the	

	delay from 50% to 0% is worth 0.32 Euro (therefore 1 min. delay is 2.4 times as bad as 1 min. travel time: risk-averse). Can be converted into a RR of 1.4	following attributes: travel time, probability of a delay, probability of a seat.	
SACTRA, (1999)	By ignoring travel time variability the economic benefits of trunk road schemes are underestimated by 5-50% (UK).		
Senna, (1991)	The disutility of the standard deviation of travel time is around 2.5 times as high as for travel time.	SP survey among 301 respondents in Porto Alegre (Brasil), with a range of travel times, mean travel time and travel costs as attributes.	Underlying theory: utility theory and attitude towards risk; Methods: SP and mean versus variance approach.
Stockholm public transport study, (2001)	Value of a minute delay is 3 times the VoT for metro and 4 for bus	unknown	Higher VoR for bus than for metro because higher share have connections that they may miss; indications of non-linearities: value of a 10 minute delay is less than twice the value of a 5 minute delay.
Swedish railway study (2004)	Value of a minute delay is 6 times the VoT or 100 euro/hour for private trips and 135 euro/hr for busbies trips	unknown	Indications of non-linearity: small delays have higher unit values than large delays.
Vrtic et al, (2006)	Reduction in proportion of late trips by 1% valued at 0.5CHF for car travel, and 0.2CHF for public transport.	Probability of late arrival included as attribute in SP.	

Table 3.3. Value of variability in freight transport (in travel time or Euros of 2003): quantitative outcomes, methods used and other lessons.

Study	Quantitative outcomes (+definition)	Method	Other lessons
Accent and HCG, (1995)	A 1% increase in the probability of delay of 30 or more min. is equivalent to 0.45 – 1.8 Euro of 2003 per transport.	Stated preference (SP) in road transport in the UK with the following attributes: travel time, travel costs, provision of information and probability of delay.	Method: SP
Bogers and van Zuylen, (2005)	Truck drivers value the unfavourable travel time twice as high as its objective (risk-neutral) worth. Managers of shippers and carriers did not have this relatively higher value for unfavourable travel times. This measure of unreliability cannot easily be transformed into a reliability ratio for the standard deviation of travel time.	SP among truck drivers and managers of shippers and carriers, used a visual presentation with one favourable travel time once in 10 days, one unfavourable and 8 normal travel times	
Bruzelius, (2001), based on Transek, (1990), (1992)	Sweden: for rail transport, a 1% increase in the frequency of delays is equivalent to 4.7-7.0 Euro per wagon; For road transport: 3.5-32.6 Euro per transport.	SP survey among shippers in Sweden in 1989/1990, including the following attributes: costs, transport time and probability of delay.	
Bruzelius, (2001), based on INREGIA, (2001)	Sweden: the value of the risk of delay is 6.1 Euro per pro mille per transport for road, 111.3 for rail and 25.7 for air transport.	SP survey among shippers in Sweden in 1999, including the following attributes: costs, transport time and probability of delay.	
Fowkes et al. (2001)	UK, road transport: the value of the difference between the earliest arrival time and the departure time is on average 1.18 Euro per min. per transport (more or less the free-flow time); for the time within which 98% of the deliveries takes place minus the earliest arrival time, the value is 1.44 Euro ('spread'); for deviations from the departure time (schedule delay) the value is 1.12 Euro.	SP survey among 40 shippers and carriers in the UK in 1999 with the following attributes: time, costs, latest departure time, earliest arrival time, arrival time for 90, 95 and 98%.	Method: SP
Hensher et al. (2007)	Valuation of reliability gains of 2.20 Euro per percentage point for transporters, 6.50 Euro for shippers. This is obtained when looking solely at the freight rate; when further incorporating all costs in the calculation, the VRG rises to 7.90 Euro. Giving an actual meaning to these values, the results would imply that, if a toll free route had a 91% probability of on-time delivery, with 97% for the tolled route, the value of trip time variability for transporters would be 13.30 Euro per trip.	Mixture models incorporating ideas of concession and power, probability of on-time arrival presented in SP.	
HCG, (1992a)	The Netherlands: an increase in the percentage not on time by 10% (e.g. from 10% to 11%) is just as bad as 5-8% higher transport costs.	SP survey in 1991/1992 among 119 shippers and carriers in goods transport by road, rail and inland waterways with the following attributes: time, costs,	Method: SP.

		percentage not on time, probability of damage and frequency.	
HCG, (1992b)	A decrease in the probability of delay by 10 index points (e.g. from 15% to 5%) is worth 0.5 – 2 Eurocent per tonne-km.	SP surveys in 1992 in The Netherlands, Germany and France with around 50 interviews per country with the following attributes: time, costs, probability of delay, frequency and flexibility.	Method: SP.
MVA (1996)	Reliability ratio for transport: 1.2	Literature review	
RAND Europe et al. (2004)	The Netherlands: a change of 10% in the percentage not on time (e.g. from 10% to 11%) is equivalent to 1.77 Euro per transport for goods transport by road. Also values for rail, inland waterways, sea and air transport.	SP/RP survey among 194 shippers and carriers in road transport with the following attributes: time, costs, percentage not on time, probability of damage and frequency.	Method: SP model.
Small et al. (1999)	USA: A reduction in the deviation from the agreed delivery time (schedule delay) by 1 hour is worth \$ 393 Euro per transport.	SP survey among hauliers in the USA; Scheduling model.	

4. Supply issues

As was made clear in Section 1.4, there has been a good deal of work in relation to modelling TTV in the demand area. However, if TTV is to be brought into the general transport modelling domain, it is essential to also represent the supply effects⁶, and in this area there has also been a great deal of research. In this Chapter we discuss the issues, with examples of the current state of theory and practice. Although the work has addressed the performance of both public and private transport networks, it is probably fair to say that the majority of the work has focused on the highway context. This balance between highway and public transport network modelling is also reflected in UK planning practice, where the tools applied for highway modelling are much further advanced than those typically applied for public transport modelling, and goes beyond the question of whether and how to model TTV. Because of these two considerations, our review therefore focuses on the highway context, before considering briefly the position for public transport. A further point to make about the scope of this chapter is that we will not consider micro-simulation methods. Although these are becoming increasingly used in the UK, the route-choice/demand side of micro-simulation tools is relatively poorly developed, at least in practice, meaning that their application has so far almost exclusively focused on supply-side issues. The starting point of this Chapter, by contrast, is that a combined approach is needed, which deals consistently with demand, route choice and supply-side factors.

The Chapter begins by briefly describing the basic principles of state-of-practice highway assignment modelling, without TTV. We then move on to consider the issues for introducing TTV within such approaches. In order to bring some structure to the discussion, we divide the modelling discussion into two broad issues, namely the largely empirical issue of how the actual variation in travel times is represented (section 4.2), and the largely theoretical/algorithmic issue of how network performance is modelled in the light of TTV (section 4.3). From this review, we see that many developments are taking place in the first of these areas that give cause for optimism regarding practical implementation, whereas the second area has some practical constraints (identified in detail in section 4.4) that do not appear to be the focus of any present study. Based on the constraints identified, we move on to consider how to take forward the modelling of network performance under TTV (section 4.5). With the focus of the chapter on highway network assignment, we briefly consider in section 4.6 the distinct issues that arise in public transport networks. In section 4.7 we conclude with a summary and recommendations for how to take this work forward.

⁶A point of terminology is that in this report we shall refer to the role of the network model as a 'supply-side process', although this is not a universal usage of this term. In particular, the route-choice process could be considered to be a component of demand, and in that case it is the implementation of capacity restraint on the links of the network that represents the true supply function; this is the terminology that would be more commonly understood in network and/or traffic modelling. However, as the current report has been written at the outset from the perspective of demand modelling, we have found it more natural to adopt the terminology that is common to that field, where the understanding of the role of an assignment model is as a method for managing the interface between demand (essentially between pairs of zones) and supply (essentially at the network/link level).

4.1 Highway assignment models: state-of-practice and issues for representing TTV

Network models are one of the commonly used tools to forecast changes in flows, demands and travel times under the conditions of a newly applied policy instrument. Typically such models assume driver decisions are motivated by the generalised cost, composed of a combination of distance (as a proxy for vehicle operating cost), travel time and any tolls payable. In this context, we might ask: what exactly do we mean by ‘travel time’? In normal circumstances, it is presumed to represent an average, but if we also wish to reflect that travel times are variable, then how might we achieve this with such models?

Since it will be important for the subsequent argument, it is worth spending a little time discussing this standard procedure, and introducing some notation. The aim is not to give a complete account of highway assignment, but to set down the main points.

Let us consider the simplest form of highway assignment model that represent well the current state-of-practice: fixed demand, one user class, static, deterministic (Wardrop) user equilibrium. The historical origin of the name ‘assignment’ in transport derives from the first all-or-nothing style models, in which the purpose was to take a matrix of vehicle travel (as movements from origins to destinations) and to assign it onto an appropriate network. Subsequently equilibrium models extended this notion to the representation of a kind of ‘game’ in which the aim is to find a consistent solution to several processes, which may be described as:

- 1) drivers on each origin-destination movement choosing minimum cost paths through the network, with cost including travel time as an attribute;
- 2) the flows on the links of the network being an aggregation of the flows on all paths for all origin-destination movements;
- 3) dealing with supply-side effects (*capacity restraint*) as a result of the volume of link flows relative to capacity, implying that the (mean) link travel times are a function of the link flows (possibly including junction interactions).

It has long been known that such fixed demand approaches may be extended to handle at least some forms of demand response; alternatively it may be desirable to model the demand responses by some external model, in which case one may also consider the output of such a model to be the travel costs on each origin-destination movement. Note that such a concept is well-defined, at least for the deterministic user equilibrium model, since at equilibrium all the used routes for a given OD movement have the same travel cost: so there is a unique notion of an OD travel cost. (The same is not true for stochastic user equilibrium models, or models where the outputs themselves are stochastic, where different measures need to be considered).

In line with the terminology of Clark & Watling (2005), we will define:

v_a	flow on link a ($a = 1, 2, \dots, A$),
\mathbf{v}	the vector of flows across all links
q_w	mean demand on O–D movement w ($w = 1, 2, \dots, W$)
\mathbf{q}	W -vector of mean demands
R_w	index set of acyclic paths serving O–D movement w
f_r	flow on path r in R_w serving O–D movement w
\mathbf{f}	vector of path flows across all paths and O–D movements
δ_{ar}	indicator variable, equal to 1 if path r contains link a , 0 otherwise
$t_a(v_a)$	(mean) travel time on link a as a function of v_a ($a = 1, 2, \dots, A$)
$\mathbf{t}(\mathbf{v})$	vector of functions $t_a(v_a)$ ($a = 1, 2, \dots, A$)

Then a fixed demand Wardrop User Equilibrium (UE) model with travel time the only component of travel cost, is defined as an assignment of flows and costs to the network that simultaneously satisfies the following conditions (with an asterisk denoting equilibrium values of the variables concerned):

- 1) Equilibrium route costs are consistent with link costs that would arise at the equilibrium link flow levels:

$$i. \quad c_r^* = \sum_{a=1}^A \delta_{ar} t_a(v_a^*) \quad (r \in R_w; w = 1, 2, \dots, W).$$

- 2) Equilibrium link flows are consistent with the equilibrium path flows:

$$i. \quad v_a^* = \sum_{w=1}^W \sum_{r \in R_w} \delta_{ar} f_r^* \quad (a = 1, 2, \dots, A).$$

- 3) Equilibrium paths flows are consistent with the OD demands:

$$i. \quad \sum_{r \in R_w} f_r^* = q_w \quad (w = 1, 2, \dots, W).$$

- 4) Wardrop condition is satisfied for each OD movement: used routes having equal cost (which we can therefore term the OD cost), and routes with higher cost carrying no flow:

$$a. \quad \text{For each } w = 1, 2, \dots, W, \text{ then } f_r^* > 0 \Rightarrow c_r^* = \pi_w^* \text{ and } c_r^* > \pi_w^* \Rightarrow f_r^* = 0 \\ (\forall r \in R_w)$$

$$b. \quad \text{where } \pi_w^* \text{ is therefore the equilibrium OD cost for movement } w.$$

Some care needs to be taken in interpreting this definition: for example, it is well known that the equilibrium path flows are non-unique, and so typically we envisage the output of such a model being in terms of entities that can be guaranteed to be unique (under certain conditions), namely the link flows, and link/path/OD costs. For a Stochastic User Equilibrium model a somewhat different form of definition is possible, since in such cases there is unique notion of the route flows (and the route choice proportions). To maintain maximum similarity with the definition above for UE, it is probably most convenient to imagine the SUE model requiring the same conditions 1-3, but replacing condition 4 with the SUE requirement:

$$\text{For each } w = 1, 2, \dots, W, \text{ then } f_r^* = q_w p_r^w(\mathbf{c}^*)$$

where for any vector of route costs \mathbf{c} , the function $p_r^w(\mathbf{c})$ denotes the proportion of demand that would use route $r \in R_w$ for movement w . Note that a similar form of definition is not possible for the UE model as such a well-defined function $p_r^w(\mathbf{c})$ does not exist in the UE case.

Before proceeding we should note that a further aspect which needs to be discussed is the consequential impact of the supply model output on demand. In terms of developing an understanding of new techniques, it makes sense to focus initially on the simpler inelastic demand case, where the only response represented in the model is one of route choice. However, it will also be appreciated that standard practice in network assessment is now to perform a variable demand evaluation, in the spirit of the VADMA advice and in conformity with DfT guidance. Thus, whatever approach is developed will ultimately require to be embedded within a supply-demand equilibrium system. While this may be viewed as primarily an algorithmic issue, in terms of whether the convergence of such a system can be guaranteed/attained at all and the run-times that would arise, it is also likely to have practical implications in terms of the generality of models that may be accommodated, and the consistency of the theoretical assumptions made about perceptions of TTV on the demand-side and the route-choice side. All such issues will need to be borne in mind at some stage; we begin to address some of them in subsequent sections, even though our primary focus is on the inelastic demand case.

Having set out the preliminaries, therefore, for conventional, fixed demand assignment modelling, the question is then: how does TVV fit into this framework? This question has several facets to it. In order to bring some structure to the subsequent discussion, we divide the modelling discussion into two broad issues that any approach to modelling TTV must consider:

- 1) How is the actual variation in travel times represented and modelled?
- 2) How is network performance affected by TTV, incorporating both the behavioural (route choice) response to TTV and the impacts this has on the equilibration process.

The first issue, of representing actual TTV, is considered in section 4.2, and is seen to divide once more into two further classifications, namely those approaches that aim to model the composite effect of all sources of TTV (without regard to the causes), and those approaches that aim to represent the distinct impacts of the various causal factors, notably demand variability and incidents.

The second issue, of modelling network performance under TTV, is a much wider issue, meaning that it is not feasible to review all the theoretical approaches that have been proposed for dealing with it. Instead in section 4.3, we choose to illustrate some of the key considerations that arise in modelling network performance by reference to a sample of approaches in the literature. Throughout we pay particular attention to the relevance of previous DfT-supported research in these areas. Although this broad, two-way classification is not ideal, as there are clearly overlaps between the decision as to how to model actual TTV and the decision as to how to model network performance under TTV,

it at least gives some way of structuring a complex research area, and some way of comparing the disparate approaches followed.

4.2 Representing actual TTV in highway assignment models

As discussed in section 4.1, an issue that cuts across the question of the overall framework to be used for network performance under TTV is: how do we actually represent the TTV itself (rather than the drivers' or system's response to it)? Clearly TTV is not constant across the links and routes of a network, and is certainly likely to be policy-sensitive. This suggests that as a minimum (and for ease of connection to observable data on variability), we would consider specifying variability on a link level. For example, in this context we could mention the work by Arup (discussed previously in section 1.4), who showed that with a linear speed-flow relationship of the form $\alpha_a - \beta_a v_a$, where v_a is link flow and α_a is the free-flow speed, it is possible to derive a formula for the coefficient of variation of link travel time on link a , CV_a :

$$CV_a = \sqrt{[\text{var } \alpha_a + \beta_a^2 \text{ var } v_a]} / (\alpha_a - \beta_a v_a)$$

and that this can be generalised, without substantial effort, to more complex speed-flow curves.

Taking a somewhat different approach, Arup (2003) have also developed original work by SDG (1993), based on the observation that TTV is likely to be greater as flows reach capacity. This observation motivated the use of the so-called *congestion index* CI as a key explanatory variable in representing TTV, the CI being defined as the ratio of the mean travel time to the free flow travel time for a *journey* (note: not for a link). In particular, it has been proposed to relate the coefficient of variation ($CV = \text{ratio of standard deviation to mean}$) of travel times to the CI, which gives model-users a simple way of generating travel time variances from standard data that already exists on mean travel times. Using data collected from routes in London and Leeds, relationships between CV and CI were estimated, and are reported in Arup, 2003 (chapter 15). In carrying out the estimation, it was noticed that there was a tendency for the journey CV to decline with journey length d . The results for the two datasets were not identical, but showed some degree of consistency, and therefore in order to achieve some kind of transferability to other locations, the consultants proposed a "compromise equation" giving a relative weight of 75:25 in favour of the Leeds results, which was of the form:

$$CV = 0.148 CI^{0.781} d^{-0.285}$$

and this has been proposed in the latest WebTAG guidance (Unit 3.5.7) for use in urban studies. An important point to note about this work is that it is based on relationships between CV and CI for a whole journey, not for individual links of a network: the relationship proposed does not decompose in a way that we can consistently assume it to hold on both a link level and a journey level.

There are several other, on-going works of relevance in this area, whose results when reported should further enrich such modelling. On-going DfT-funded research by Hyder has the aim of “develop[ing] sound theoretical models supported by robust data to give firm underpinning to future estimates of the relationship between travel time variability traffic flows and journey time.” In a similar spirit, on-going work by Mott MacDonald for the DfT (yet to be reported) has used extensive inter-urban data from the HATRIS database to derive various forms of model, relating travel time variability to *link* (rather than whole trip) attributes, for different link types and day-types, thus providing potential for link-based variability inputs to be readily generated for network models. Alternative model forms for inter-urban studies are likely to arise from this research.

However, a somewhat different approach to this problem is to start with a belief that we need to explicitly model the *causes* of TTV, rather than seeing TTV as either a fixed, inevitable element or one that is only linked to the mean link travel time. In this respect, we might refer to the Arup work noted in Section 1.4, which adopted an approach in which highway TTV derives from two general sources, what they referred to as “incident-induced” variability and “day-to-day variability” [DTDV]. The latter term refers to the variability in travel time brought about by changes in **demand** from one day to the next, considered at the same time of day, and allowing for seasonal and other, essentially predictable, effects. This consideration implies that the earlier, CV-related work described above describes the composite effect of variability from both DTDV and incidents. An argument in favour of the (composite) CV approach is that a) no evidence has so far been adduced that the *demand* response is different between the two sources, and b) at least in an urban context, the impact of most incidents is likely to be difficult to distinguish from that of the random demand fluctuations. The counter argument is that by representing the underlying causal factors, we may gain a deeper understanding of TTV and introduce more potential policy levers to affect TTV in our models. Adopting this kind of philosophy, allowing for DTDV would in principle evaluate the impact on link travel time variability by allowing explicitly for the variation in the OD demand matrix \mathbf{q} . We discuss this further in the next section.

Continuing with the line of modelling the underlying causes of TTV, we should mention that the treatment of *incidents* is conceptually separate and rather more problematic than that of representing DTDV. These are essentially capacity-reducing random events whose frequency is nonetheless sensitive to the level of demand, and whose *consequences* are also sensitive to the level of demand, in terms of the length of queues that may be generated. One approach to conceptualising such incidents is as follows: Define an incident as a discrete event E with random characteristics \mathbf{h} (which may include the type of incident, the amount of the carriageway blocked etc.). On any given link, the random probability of E occurring is conditioned by the flow v : $p[E|v]$. *Given* E , the consequent delays are a function of \mathbf{h} and the flow v . Once again, therefore, if we could input the associated probabilities (essentially those of an incident of given characteristics occurring), then we could generate the outcome distribution in terms of travel time. It will be appreciated, however, that this is an even more demanding task than the case of DTDV.

The impact of incidents is likely to be greatest when alternative routes are limited, as in the case of Motorways. It is precisely in such cases that explicit attention may be given to projects for reducing the effects of TTV due to incidents. The interest in this topic has led to the work underlying INCA (which in turn derives from INCIBEN), which has been funded by the Department. This was carefully reviewed within the Arup work, together with a comparable approach developed in the US by Cohen & Southworth (1999). Chapters 5 and 6 of D6.1 (Arup, (2002)) provide a clear theoretical account. Key factors noted by Arup (Chapter 2) are as follows:

- The impact of an incident is different from the DTDV case in that the incident occurs at a random time (and place), and this random time of occurrence itself induces some variability in terms of impact.
- If the motorist arrives before the incident occurs he will experience no delay. If he arrives after the queue has dispersed he will again experience no delay. Otherwise, the delay that he experiences depends on his time of arrival relative to the time of the incident. If the total duration of the incident is T minutes, then the probability of a motorist experiencing some delay is equivalent to the probability of the incident occurring up to T minutes earlier than the time of “arrival”, while they argue that the delay experienced by an individual motorist who encounters a queue can be considered to be uniformly distributed.
- This allows us to calculate the overall mean and variance of delay from a randomly occurring incident, taking into account the probability of occurrence. An important result is that the mean is proportional to the square of the incident duration, while the variance is related to the third and fourth powers.

Arup considered how to allow for random characteristics of an incident, in terms of its duration and severity: by assuming either a theoretical distribution, or an empirical distribution drawn from a sample of observations based on recorded data from actual incidents, as is done in the INCA methodology. The conclusions of their analysis was that it is legitimate, as an approximation, to add the delay variances over incident types and links, as is done in INCA, provided: (a) the square of expected delay is small relative to expected value of the squared delay; and (b) not more than one incident (causing delay) is encountered during the course of the journey. The underlying basis is currently being extended to single carriageways, it being originally for motorways only and later extended to grade-separated dual carriageways. The program is recommended in WebTAG Unit 3.5.7. Although the previous version of INCA did not incorporate an explicit calculation of TTV due to DTDV, the new INCA 4.0 now models DTDV explicitly.

In summary, there is a great deal of work in the area of representing actual TTV, whether in composite form or by its component causes, much of this work DfT-funded, and several studies on-going at present. The understanding emerging from these studies is sufficient to believe that specifying models of TTV is feasible in realistic networks, and such models will only improve with further investigation of empirical sources.

4.3 Modelling highway network performance under TTV

As explained in section 4.1, we have divided the modelling of highway networks under TTV into two broad themes; in the previous section we considered how we might represent the actual variation in travel times, without regard to how the drivers or the network may be impacted by this variation. In the present section, on the other hand, we presume that suitable models of actual TTV exist (whether composite or as separate cause-related models), and the key question is how to model network performance in such an environment. In other words, if travel times are variable, how do we take forward the conventional UE/SUE approaches described in section 4.1?

Taking a software-oriented viewpoint, a common first response to this question runs as follows. Let us consider, purely for illustration, a cause-based approach to representing TTV based on representing DTDV effects only; that is to say, stochastic travel times arise from the stochastic demand. It seems reasonable, then, that we could proceed along the following lines, using Monte-Carlo (MC) methods:

Assume a probability distribution for the day-to-day distribution of demand \mathbf{q}

Simulate day n , initially with $n=1$:

Randomly select a demand $\mathbf{q}^{(n)}$ from the assumed demand distribution.

Assign $\mathbf{q}^{(n)}$ using a standard UE/SUE approach.

Record the link travel times $\mathbf{t}^{(n)}$ that arise from assigning this demand.

Increase n by 1 and return to simulate the next day, until $n=N$.

Collect together the results from all N days to give an empirical distribution.

Conceptually the ideas seem straightforward, and we are left only with a numerical challenge of replicating many equilibrium assignments. In fact, we may circumvent this numerical speed issue by using an approximation to the equilibrium model, known as sensitivity analysis, which means that we can produce the same (in fact better) results as the process above in a fraction of the time, using a single equilibrium assignment run. Clark & Watling (2005) critique this kind of approach in some detail, which has been suggested in several previous studies, and highlight the key flaw in its reasoning. Namely, since equilibrium is reached on each day, it implies that drivers have perfect predictive knowledge of the travel times they will experience on that day, before they make their journey. This seems a reasonable assumption for slowly-changing or systematic trends where drivers may have an opportunity for repeated experience: for example, if the different samples were to represent the different mean demands arising in school-holidays and term-times. For events that can change on a daily basis, such as demand and incidents, it seems more difficult to justify however, and so even though the

resulting problem is tractable it is highly questionable as to whether it is the appropriate approach for representing the kind of TTV impacts we have in mind here.⁷

Although the process described above does not, then, seem the appropriate solution, it is useful as a reference point to consider how we might go about incorporating TTV in a suitable way. In particular, it is clear that a representation of TTV in the model is an acknowledgement of some kind of uncertainty – but whose uncertainty? Once we are clear on this question, we are much closer to finding an appropriate method. Two alternative answers to this question that will motivate the work presented below are:

- a) The drivers have uncertainty in making their route choice decision.
- b) The planner has uncertainty, even if the input data to the model are error-free, because the actual day-to-day performance of the network is variable daily.

A rather fundamental question to consider first is: does TTV bring into question our whole equilibrium framework? Since the equilibrium models, be they UE or SUE, all implicitly assume that travellers have perfect predictive knowledge of travel costs/times, does TTV then bring with it a contradiction, since drivers could not possibly anticipate the times they will experience when they are making their route choice decisions? While several authors have argued for such radical departures from accepted methods, it seems that even in such situations it is still possible to extend and utilise our existing tools: this is an issue to which we shall return in subsequent sections.

A much simpler way of interpreting the role of TTV, and the one that has been adopted in previous DfT-funded work, is that essentially TTV is just another attribute to include in the list of components of drivers' generalised costs. That is to say, there is some measure of TTV (say, the travel time standard deviation on the alternative routes/links) that drivers are assumed to perceive from their repeated use of the transport system, and that the issue then is only how to 'value' this additional attribute. This sort of thinking allows us to stay within our conventional equilibrium modelling framework, although (as we discuss below) there are still some significant algorithmic challenges that then arise. Clearly there are also empirical questions as to what extent drivers value and respond to such variability when making their route choice decision. Certainly, though, it may be supposed that particular routes with a high level of TTV (for example, one including a non-priority junction at the intersection with a major road) might be avoided even if their average performance is acceptable, making the case that allowing for TTV should improve our modelling of drivers' route choice decisions.

⁷ Just to emphasise the fact that the reason for rejecting the Monte Carlo procedure proposed at the start of the section is one of appropriateness, rather than its computational difficulty, it does seem that this kind of approach is appropriate for other areas of the Department's interests, such as representing the impact of uncertainty in traffic forecasts due to limitations in the input data provided to a model. In a different source paper to that cited above, Clark & Watling (2006) describe the use of such an approach (comparing Monte Carlo and more efficient analytic methods), in which the probability distribution of demand now represents the sampling error of the observer/planner in estimating the current mean demands, and the objective is to derive a confidence interval for the resulting model outputs that reflects this uncertainty in the input data.

B1.1 ArupTransportPlanr

APPENDIX C

c1.1.1 Demand Response to Uncertainty and the Weight Given to Major Incidents

