The Valuation of Travel Time Variability

04

Discussion Paper 2016 • 04

Mogens Fosgerau
Technical University of Denmark,
Kongens Lyngby,
Denmark

The valuation of travel time variability

Discussion Paper 2016-04

Prepared for the Roundtable:
Quantifying the Socio-Economic Benefits of Transport

9-10 November 2015,
Paris, France

Mogens Fosgerau
Technical University of Denmark
Kongens Lyngby
Denmark

February 2016

The International Transport Forum

The International Transport Forum is an intergovernmental organisation with 57 member countries. It acts as a think tank for transport policy and organises the Annual Summit of transport ministers. ITF is the only global body that covers all transport modes. The ITF is politically autonomous and administratively integrated with the OECD.

The ITF works for transport policies that improve peoples’ lives. Our mission is to foster a deeper understanding of the role of transport in economic growth, environmental sustainability and social inclusion and to raise the public profile of transport policy.

The ITF organises global dialogue for better transport. We act as a platform for discussion and pre-negotiation of policy issues across all transport modes. We analyse trends, share knowledge and promote exchange among transport decision-makers and civil society. The ITF’s Annual Summit is the world’s largest gathering of transport ministers and the leading global platform for dialogue on transport policy

The Members of the Forum are: Albania, Armenia, Argentina, Australia, Austria, Azerbaijan, Belarus, Belgium, Bosnia and Herzegovina, Bulgaria, Canada, Chile, China (People’s Republic of), Croatia, Czech Republic, Denmark, Estonia, Finland, France, Former Yugoslav Republic of Macedonia, Georgia, Germany, Greece, Hungary, Iceland, India, Ireland, Israel, Italy, Japan, Korea, Latvia, Liechtenstein, Lithuania, Luxembourg, Malta, Mexico, Republic of Moldova, Montenegro, Morocco, the Netherlands, New Zealand, Norway, Poland, Portugal, Romania, Russian Federation, Serbia, Slovak Republic, Slovenia, Spain, Sweden, Switzerland, Turkey, Ukraine, the United Kingdom and the United States.

International Transport Forum
2 rue André Pascal
F-75775 Paris Cedex 16
contact@itf-oecd.org
www.itf-oecd.org

ITF Discussion Papers

Table of contents

This paper provides an overview of some alternative conceptual definitions of travel time variability, discusses their implications about behaviour, and puts them into a broader context, including deviations from the underlying assumptions regarding rational behaviour. The paper then discusses the empirical basis for assigning a value to travel time variability. This discussion leads to the conclusion that a fair amount of scepticism is appropriate regarding stated preference data and that attention should turn to the possibilities that are emerging for using large revealed preference datasets. The bottom line is that travel time variability is quantitatively important and cost-benefit analysis should account for it, using the best values we can get, in order not to imply a bias towards project that do not reduce travel time variability. Omitting the cost of travel time variability is not the neutral option.

Introduction

Many commuters across the world find their travel time between home and work to be rather unpredictable. In addition to systematic variation by time of day and by day of the week, travel time has a considerable random component. Random travel time variability is a significant issue. The annual total time and money expenditures on transport in the US are worth more than USD 5 trillion (2007), which corresponds to more than 30% of the US GDP (Winston, 2013). For the US in 2011, road congestion for work trips alone caused an estimated 5.5 billion hours of travel delay and 2.9 billion gallons (11 billion litres) of extra fuel consumption with a total cost of USD 121 billion (Schrank et al., 2012). Nonrecurring traffic congestion (due to accidents, bad weather, special events, and other shocks) contributes 52-58% of total delay in US urban areas (Schrank et al., 2011). This indicates that accounting for travel time variability would add a very significant amount to the accounted cost of traffic congestion.

The variability of travel time is often large as a proportion of travel time on a given trip. Looking for example at a range of routes in central Stockholm, the standard deviation of travel time is on average 25% of mean travel time during commuting peaks with values ranging up to 75%, taking into account the systematic variation of travel time over the peak (Fosgerau et al., 2014). Accounting for travel time variability would then add very different costs to different routes, which would matter for traveller route choices, and also for project selection.

Cost-benefit analysis of transport projects builds on an assessment of the generalised travel costs to travellers with and without projects under consideration. The generalised travel costs include monetary costs as well as a range of time related costs accounting for the value of time spent on various parts of trips. Changes in the generalised travel costs induce changes in demand, which in turn are associated with changes in the consumer surplus that enters the cost-benefit analysis. To be able to include travel time variability into transport project cost-benefit analysis, we then need to include the cost of travel time variability into the generalised travel costs. This requires three things: we need to define a unit of measurement for travel time variability, we need to predict the quantity of travel time variability, and we need to determine the cost to travellers per unit of travel time variability.

We have in mind a setup where a traffic model predicts travel demand and computes travel time distributions, including means and variability of travel times. The output from the traffic model is subsequently used in a cost-benefit analysis. Values of time and of travel time variability are applied to the corresponding quantities in the traffic model output to compute changes in the generalised cost of travel. Travel time variability would also be accounted for in traffic model description of travel demand and thereby impact behaviour.

This paper first provides an overview and discussion of alternative conceptual definitions of travel time variability, emphasising models that have a foundation in micro-economic theory. Such a foundation is a strong advantage. It embodies rationality assumptions that constrain the behaviour implied by the models, thereby ensuring that the behaviour implied by the models makes sense. The information that we can obtain from empirical data will probably never be sufficient to pinpoint the best model. The restrictions implied by theory on the range of possible models are therefore useful to complement the empirical evidence that we can find.

Having established some conceptual models for travel time variability, we next discuss some broader issues. The first is the fact that travel time variability will often be costly to people other than the traveller, including other meeting participants. However, the standard models used to assign a cost to travel time variability take the perspective of a single agent and do not represent other people. The second issue is the role of information and expectations. What counts as random travel time variability depends on the information available to travellers and it is important that we take this into account. Very little is known, however, about how travellers actually form expectations of travel times. Third, the models that we use are based on neoclassical rationality assumptions, but these assumptions can be systematically violated, especially in stated preference experiments. The paper discusses what to do about this issue.

The paper continues to review some empirical evidence, indicating the main challenges we need to confront and arguing that we should move towards using revealed preference data to infer the cost of travel time variability to travellers.

This paper is not a literature review and does not mention all the relevant literature. Li et al. (2010) review both theoretical and empirical aspects of the value of reliability. Carrion and Levinson (2012) provide a review with much historical detail and a broader focus than presented here. Small (2012) reviews the broader literature on the valuation of travel time. Concerning freight transport, a recent review is provided by Feo-Valero et al. (2011). The reader may refer to these papers for a more comprehensive overview of the literature.

Conceptual models

The basic model

We consider the simplest possible models that meet the standards of classical micro-economics, while allowing us to talk about the cost of travel time variability in a meaningful way. This section outlines a general model structure with some specific instances presented in sections below. We discuss the assumptions involved in some detail such that we know how to interpret the numbers that arise. Some technical details are omitted in order to focus the exposition on intuition; these details may be found in the original papers. For concreteness, we talk about the case of a commute trip from home to work, but the models are applicable to any trip. The models describe passenger transport but may just as well be used to describe freight transport trips. Empirically, it is very important to account for the diversity of trip purposes as some trips are much more sensitive to delays than others: think for example about freight trips with perishable goods that may lose their value due to delays or about passenger trips accessing a flight connection at an airport or urgent trips to the hospital.

The first step in defining a model that is useful for valuing travel time variability is to consider what the outcomes of interest are. The minimum we have to consider is travel time itself; we shall consider models below where outcomes are described just in terms of travel time. Here we shall use a little more structure, assuming that what matters about a trip is *when* it takes place, using notation t_{dep} for the time of departure and t_{arr} for the time of arrival. This will allow us to distinguish trips of the same duration at different times of day. The travel time T is the difference between the departure time and the arrival time; from the point of view of the traveller, the travel time is random and we assume that the traveller knows the distribution of travel time. The traveller selects the departure time and then the arrival time is given by $t_{arr} = t_{dep} + T$.

The next step is to think about a utility function $U(t_{dep}, t_{arr})$ that ranks all possible outcomes. We call it a scheduling utility, since it concerns the scheduling of the trip. We shall consider several possibilities below with the thing in common that scheduling utility is money-metric, such that the utility difference between two potential outcomes is a monetary value. At this stage, we could equally well have formulated utility in terms of travel time and either departure time or arrival time, since travel time, departure time and arrival time are related by an identity. Formulating utility in terms of departure time and arrival time emphasises that it is these things that matter to travellers, rather than the travel time itself. We assume that travellers always prefer to depart later and to arrive earlier, *ceteris paribus*.

In order to arrive at tractable expressions for the cost of travel time variability, we will make the simplifying assumption that utility is separable, splitting into a part that depends only on arrival time and another part that depends only on arrival time, $U\left(t_{dep}, t_{arr}\right) = U_1\left(t_{dep}\right) + U_2(t_{arr})$. A natural way to think of the part of utility that depends on departure time is to say that utility is accumulated at home at some potentially time-varying rate $h(t)$ until the time of departure

$$U_1\left(t_{dep}\right) = \int^{t_{dep}} h(t)dt,$$

and similarly that utility is accumulated at work at another time-varying rate $w(t)$ after the time of arrival

$$U_2(t_{arr}) = \int_{t_{arr}} w(t)dt.$$

The utility rates should be understood as differences from the utility rate that is achieved while travelling (Oort, 1969). This generic form for scheduling utility was first formulated by Vickrey (1973). It is illustrated in Figure **1**.

Figure 1. **Utility rates**

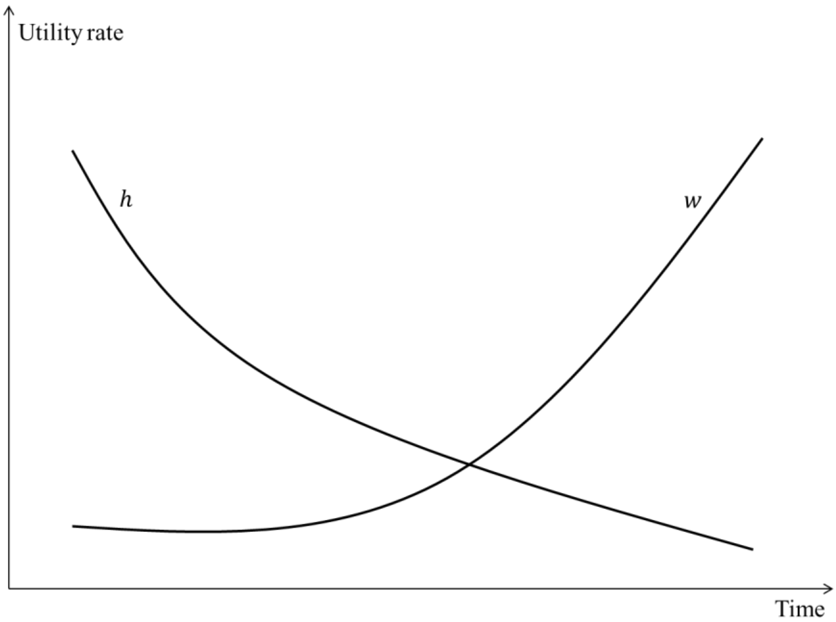

The third step is to say that travellers choose departure time, before observing the corresponding travel time. We assume that he or she chooses departure time to maximise the expected utility $E\left(U_1(t_{dep}) + U_2(t_{arr})\right)$.

The final element in the basic model is a technical assumption needed to ensure analytical tractability. We assume that the random travel time T has a distribution that is independent of the departure time. Strictly speaking, this is mostly inaccurate: in the morning peak, the travel time will be longer on average at the height of the peak compared to traveling at the shoulders of the peak. The model is in any case a useful approximation (Fosgerau and Karlstrom, 2010; Fosgerau and Fukuda, 2012).[1]

We can then rewrite travel time as $T = \mu + \sigma X$, where μ is the mean travel time, σ is the standard deviation of travel time and X is the standardised distribution of travel time, having mean zero and standard deviation 1 that is independent of departure time. This is useful to give some structure to the distribution of travel time. We denote by f the density of standardised travel time and by F the corresponding cumulative distribution of standardised travel time.

The expected utility is

$$E\left(U(t_{dep}, t_{dep} + T)\right) = \int^{t_{dep}} h(t)dt + E\left[\int_{t_{dep}+\mu+\sigma X} w(t)dt\right].$$

The traveller is assumed to choose departure time to maximise this expected utility. We have in mind trips where the departure time can be freely chosen, in which case the derivative of the expected utility with respect to departure time must be zero at the optimal departure time t^*_{dep}, i.e.

$$h(t^*_{dep}) = Ew(t^*_{dep} + \mu + \sigma X).$$

Given specifications of h and w, this equation could be solved to find t^*_{dep}. So in principle, this determines the optimal departure time t^*_{dep} as a function of the mean travel time, its standard deviation as well as the distribution of the standardised travel time, i.e. $t^*_{dep}(\mu, \sigma, f)$. Through this function, the scheduling model explicitly takes into account that travellers will schedule their trips in response to the distribution of travel time, and this is what makes the scheduling model different from models where utility depends only on travel time and not on the timing of trips. Plugging the optimal departure time into the expected utility shows, in principle, how the optimal expected utility depends on the distribution of travel time.

We can define the *value of travel time* as (minus) the derivative of optimal expected utility with respect to mean travel time and the *value of travel time variability* as (minus) the derivative of optimal expected utility with respect to some measure of dispersion. The term “the value of time” is a tradition in the transportation economics literature; a more precise term would be the “marginal expected disutility of travel time”. Similarly, a traditional term for the “value of travel time variability” is the “value of reliability” and it would more precisely be called the “marginal expected disutility of travel time variability”. We may define the value of travel time variability with respect to the standard deviation of travel time, in which case we will have the value of travel time standard deviation. Another possibility is to use the variance of travel time and consider the value of travel time variance. Which measure of dispersion is more appropriate depends on the model as we shall see below. In order to compare studies from different contexts we may consider the *reliability ratio*, which is the ratio between the value of travel time variability and the value of time. Usually, the value of travel time variability that enters the reliability ratio is expressed as the value of travel time standard deviation.

We have now arrived at a generic micro-economic model that comprises the minimal number of elements and assumptions, while accounting for the timing of trips. The generic model allows us to go some way in analysing the cost of travel time variability, but it is necessary to impose additional assumptions to make it operational. There are several attractive possibilities for these additional assumptions and they will lead to different measures of the cost of travel time variability that we may use in applications. They also imply different behaviour, which enables them to be distinguished empirically.

Step model

The most common specification of scheduling utility is due to Vickrey (1969), who included scheduling preferences in his now famous bottleneck model of commuting in a congested demand peak. Later, Small (1982), in a parallel development, formulated the same scheduling utility in an analysis of the timing of commuter work trips. It specifies the utility rate at home to be constant $h(t) = \alpha$, while the utility rate at work is a step function

$$w(t) = \begin{cases} \alpha - \beta, t \leq t^* \\ \alpha + \gamma, t > t^*, \end{cases}$$

where t^* is a preferred arrival time and α, β, γ are positive constants. The model based on this specification of scheduling utility is known as the step model. The utility rates are shown in Figure **2**.

Figure 2. **Utility rates in the step model**

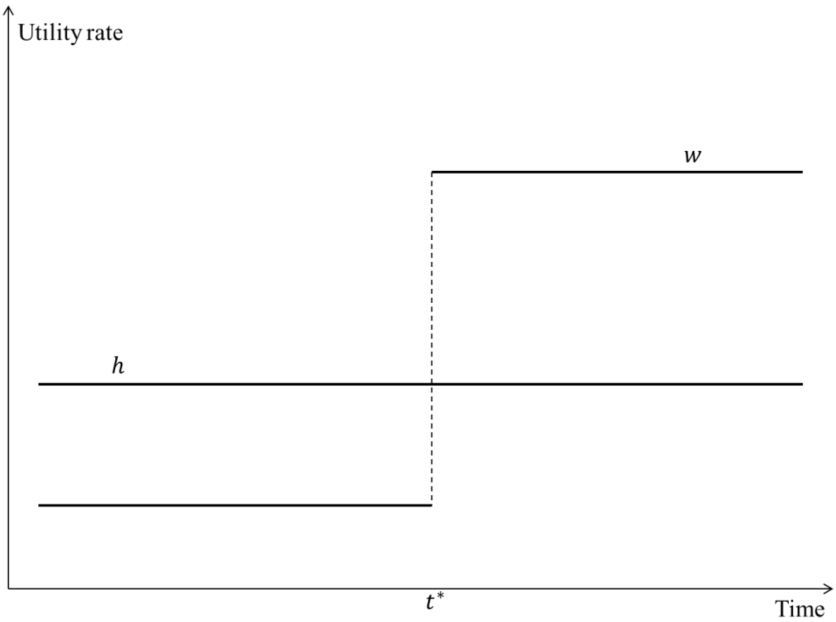

The optimal departure time can then be found to be

$$t_{dep}^* = t^* - \mu - \sigma F^{-1}\left(\frac{\gamma}{\beta + \gamma}\right).$$

This has an intuitively appealing interpretation: The traveller will depart in advance of the preferred arrival time, allowing for mean travel time μ as well as some additional head start or safety margin $\sigma F^{-1}\left(\frac{\gamma}{\beta+\gamma}\right)$ that is proportional to the standard deviation of travel time. The proportionally factor $F^{-1}\left(\frac{\gamma}{\beta+\gamma}\right)$ depends on the shape of the travel time distribution as well as on preference parameters β, γ that express the cost of arriving early or late. The traveller will arrive later than t^* with probability $\frac{\beta}{\beta+\gamma}$, which depends only on the preference parameters and not on the distribution of travel time.

Inserting the optimal departure time into the expected utility and differentiating, shows (after a bit of work) that the cost of mean travel time μ is a constant α per time unit, while the cost of travel time variability is $(\beta + \gamma)\int_{\frac{\gamma}{\beta+\gamma}}^{1} F^{-1}(s)ds$ per unit of standard deviation σ. The reliability ratio is

$$\frac{\beta + \gamma}{\alpha}\int_{\frac{\gamma}{\beta+\gamma}}^{1} F^{-1}(s)ds,$$

which depends on preference parameters and on the standardised travel time distribution but not on the mean μ and the standard deviation σ of travel time.

These convenient results were first established for some special travel time distributions (Noland and Small, 1995; Bates et al., 2001) and then later for a general travel time distribution (Fosgerau and Karlstrom, 2010).

To put a ballpark number on the reliability ratio, we may use the stylised values $\alpha = 2, \beta = 1, \gamma = 4$ based on Small (1982) and a value of 0.3 for the integral expression in the reliability ratio (Fosgerau and Karlstrom, 2010). Then the reliability ratio becomes 0.75.

The standard deviation may be replaced as a measure of travel time variability by any other statistic that is proportional to the standard deviation, provided that the shape of the travel time distribution may be considered to be constant. To see this, note that for any positive number ρ, we may rewrite the cost of travel time variability into the following, simply dividing and multiplying by ρ.

$$\left(\frac{\beta+\gamma}{\rho}\right)\left(\rho\sigma\int_{\frac{\gamma}{\beta+\gamma}}^{1} F^{-1}(s)ds\right).$$

Any measure of travel time variability that is proportional to the standard deviation has the form shown in the second parenthesis here. The first parenthesis is then the corresponding unit value of that measure of travel time variability.

Given a fixed shape of the distribution of travel time, the standard deviation is proportional to many other measures of the dispersion of travel time that have been used. This includes:

- the difference between two specific quantiles of the travel time distribution used, notably, by Small, Winston and Yan (2005)
- the difference between a quantile and the mean travel time
- the buffer time index (Texas Transportation Institute and Cambridge Systems, Inc., 2006), and
- the mean lateness, $\sigma\int_{\frac{\gamma}{\beta+\gamma}}^{1} F^{-1}(s)ds$.

All such measures are proportional when the shape of the standardised travel time distribution is constant. Hence, they all share the above micro-economic foundation in the step model.

A drawback of the step model is that the standard deviation is not additive across parts of a trip having independent travel times. If the travel time for the first part of a trip has standard deviation σ_1 and the second part has standard deviation σ_2, then the travel time for the combined trip has standard deviation $\sqrt{\sigma_1^2+\sigma_2^2}$, which is strictly smaller than $\sigma_1+\sigma_2$ with a difference that can be large. If, for example, $\sigma_1 = \sigma_2$ then $\sqrt{\sigma_1^2+\sigma_2^2} = \sqrt{2}\cdot\sigma_1$. This is inconvenient when working with network based traffic models, since the expected travel cost for a trip then cannot simply be added up from the level of links.

The step model may be used in several ways. First, it may be used as a structural model, where preference parameters α, β, γ and t^* are combined with an observed travel time distribution to compute cost measures. This is feasible in particular using stated preference data where the preferred arrival time t^* may be observed or inferred. It may also be used in a reduced form where the expected travel cost for a trip is simply described as being linear in the mean travel time and in the standard deviation (or some other measure proportional to the standard deviation).

Slope model

We now elaborate another version of the general model, but this time based on utility rates that are time-varying with constant slopes (Fosgerau and Engelson, 2011). The slope model specifies utility rates as

$$h(t) = \alpha - \beta \cdot (t - t^*)$$

$$w(t) = \alpha + \gamma \cdot (t - t^*),$$

where again α, β, γ are positive parameters while t^* plays the same role as the preferred arrival time in the step model of locating scheduling preferences in time. These utility rates are illustrated in Figure **3**.

Figure 3. **Utility rates in the slope model**

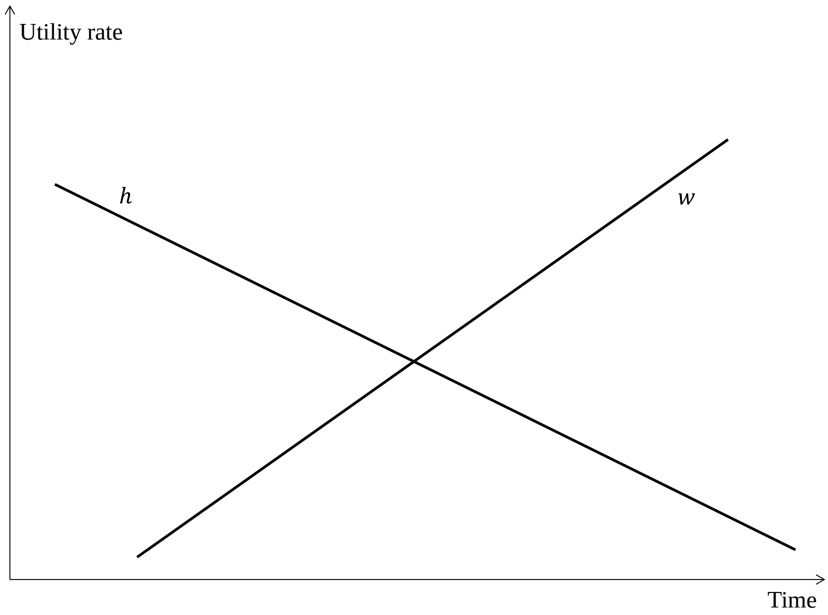

Source: Fosgerau and Engelson, 2011.

Proceeding in the same way as before leads to the optimal departure time

$$t_{dep}^* = t^* - \frac{\gamma}{\beta + \gamma}\mu.$$

The traveller will depart $\frac{\gamma}{\beta+\gamma}\mu$ before time t^*. On average he/she will arrive $\frac{\beta}{\beta+\gamma}\mu$ after time t^*. The amount of random travel time variability affects of course the distribution of arrival times, but has no effect on the departure time or on the average arrival time. The traveller in the slope model does not allow an additional safety margin when faced with random travel time. This distinguishes the slope model from the step model.

The value of mean travel time is $\alpha + \frac{\beta\gamma}{\beta+\gamma}\mu$ per unit, which depends itself on the mean travel time. Thus the value of time is higher for long trips than for short trips. This is because long trips begin at a time when the utility rate at home $h(t)$ is higher. When $\beta = 0$ such that the utility rate at home is constant, then also the value of mean travel time is constant and equal to α.

The cost of travel time variability can be expressed as the cost per unit of the variance of travel time and it is then a constant $\gamma/2$. Then the cost of travel time variability has the advantage, in contrast to the step model, of being additive across links if the random travel times on links are independent.

The value of the standard deviation of travel time is not constant but is proportional to the standard deviation. The reliability ratio in terms of the standard deviation is then not constant and is not immediately comparable to the reliability ratio from the step model.

Engelson and Fosgerau (2011) extend the slope model to the case where the utility rate at work is an exponential function and show that this exhausts all possibilities for slope models that are additive across links.

Non-scheduling models

The step and the slope models describe the timing of individual trips and the preferences concerning the times of departure and arrival, and combine this with expected utility maximisation. This structure has advantages; in particular it ensures that the behaviour predicted by the model makes sense. It leads to predictions regarding how departure times relate to the distribution of travel time, which may be used to evaluate models against empirical evidence. But the structure may also be a constraint if the models do not match actual behaviour. It is alternatively possible to use models that impose less structure.

Such models would ignore the timing of trips and just assume that travellers have preferences regarding travel time where less is better. The equivalent monetary cost of a trip would then be expressed as a convex function of travel time $C(T) = C(\mu + \sigma X)$, and the value of travel time and variability would be derived from the expected cost function, with the expected cost per unit of mean travel time becoming $EC'(T)$ and the expected cost per unit of standard deviation becoming $E[C'(T)X]$.

There are several convenient forms that can be used for the cost function. For example a quadratic cost function $C(T) = \mathrm{bT} + \mathrm{cT}^2$ leads to (Polak, 1987)

$$\frac{\partial EC(T)}{\partial \mu} = b + 2c\mu, \frac{\partial EC(T)}{\partial \sigma^2} = c,$$

which is a different model from the linear slope model presented above.

We are not free to specify any conceivable cost function: it may turn out that some forms are not consistent with underlying scheduling preferences. We would have less faith in such models as they would then not be consistent with any underlying rational scheduling behaviour.

Scheduled services

The scheduling models just presented assume that travellers are able to select their departure time optimally, as is the case for car drivers. Travellers who use scheduled services are constrained in their choice of departure time and this affects their value of reliability. In the case of the slope model it turns out that the value of travel time variability is unchanged relative to the case of car drivers for most of the models (Fosgerau and Engelson, 2011; Engelson and Fosgerau, 2011)

In the case of the step model, a similar result for the variability of total travel time is not available (Fosgerau and Karlstrom, 2010). Something can be said, though, for the case of a frequent scheduled service. When departures are sufficiently frequent, travellers do not aim for a specific departure, but arrive at the station to catch the next departure, not knowing specifically when that will be. The waiting time until the next departure is then random from the perspective of travellers. In this, case the general result from the step model applies with the arrival time at the station replacing the self-selected departure time (Benezech and Coulombel, 2013). Then travel time variability may be accounted for, due to random irregularity of departures as well as due to random variability of travel times.

Some broader perspectives

We now have some basic theory in place that allows values to be assigned to travel time variability. We shall now discuss this theory from a somewhat broader perspective and we begin by looking at the framing of the scheduling models in terms of a single individual.

External delay costs

The models that we have considered take the perspective of a single individual who travels from home to work or, more generally, between two activities. They consider the cost to the individual related to the time of departure from the first activity and the time of arrival to the second activity. In many cases, however, there will be someone waiting at the second activity and it may have consequences for them if the traveller experiences an unanticipated delay. Such costs are not included in the models we have discussed if the traveller does not include them fully in her scheduling preferences.

We may consider a situation where two persons facing random travel times schedule trips to a joint meeting and where the time one person spends waiting for the other is unproductive. Using the scheduling models above would lead us to consider each person independently, accounting just for the cost of travel time variability for each person, assuming that his/her scheduling preferences were constant and exogenous. Fosgerau, Engelson and Franklin (2014) analyse such a situation and find that it is not only the variability of travel times that matters but also their correlation: If travel times are positively correlated then random delays are less costly since the participants in the meeting will tend to experience similar delays and then not waste that much time waiting for the other. The setup also has consequences for how departure times depend on the distribution of travel times, since now the travel time distribution faced by one person, as well as the departure time decision made by that person, affects the outcomes for the other person. This reasoning applies to situations where more participants travel to a joint meeting, but not to situations where some participants do not have to travel to reach the meeting.

We can think of many situations where one person arriving late implies costs for others. Another central feature that seems to be commonly present is that some people are able to adapt to the travel time variability faced by others. So far, these issues seem not to have been investigated at all, and research to gauge the impact on the costs of travel time variability would be welcome. A realistic ambition may be to form an opinion about whether using the simple single agent scheduling models discussed above would lead us to an overestimate or an underestimate of the “true” costs of travel time variability.

Expectations and information

Travel times vary for many different reasons. There is systematic variation over the week and during the day due to systematic variations in demand. On top of that, there is demand variation that is more or less predictable in principle due to holidays, special events, weather and road works. Then there are incidents and accidents, which by nature are highly unpredictable.

To somebody having no information, just observing travel times, all travel time variation would be random. At the other extreme, a completely informed traveller might be able to predict travel time perfectly and for that traveller, travel time would be not random at all. Actual experienced travellers are somewhere in between. They might have a good feeling for systematic variation by day of the week and by time of day, they might know about special events and road works, and they could be more or less sophisticated in their ability to predict travel times.

In order to predict the cost of travel time variability, it is necessary to form an opinion about what information travellers use to form their expectations. A straightforward and practical solution is to assume that travellers know the traffic model used in a given application; they are just as well informed as the model. This solution has the attraction of being simple but it is not completely innocuous. Traffic models may predict travel times that depend on the time of day for different days of the week. In that case, travellers may be better informed than the traffic models since they may have information about special events and special days, that is not used in the traffic models. This may imply some overstatement of the cost of travel time variability.

Non-rational behaviour

The models presented above are based on classical economic rationality assumptions. These assumptions are useful in constraining models to deliver predictions that make basic sense. In brief, the assumptions amount to the following. Travellers have preferences over outcomes, meaning that they evaluate a potential trip in terms of the characteristics of that trip and nothing else. These preferences are such that they can be expressed in terms of a utility function. When faced with choices with uncertain outcomes, such as the arrival time when travel time is random, travellers make the choice with the maximum expected utility. It is clear that this is not an exact description of human behaviour, and certainly utility maximisation does not have an exact counterpart in the brain. But that is not really a problem: we are content if the description is about right on average.

A body of evidence is emerging, showing that human behaviour, especially in certain experimental settings, deviates systematically from rational behaviour. Two phenomena of particular relevance for valuing travel time variability are those described by prospect theory (Kahneman and Tversky, 1979). One is loss aversion: that preferences depend on the size and direction of change from a reference point. In the case of travel, the reference point may be a recent trip, or just some kind of “normal” trip. What may be considered as a reference point depends a lot on the context. In any case, the presence of loss aversion contradicts the assumption that preferences should depend on outcomes only. The other phenomenon is probability weighting of utility, which leads to small probability events having a larger influence on behaviour than they would under the plain mathematical expectation.

The stated preference experiments that are often used to measure the value of travel time or the value of travel time variability may induce non-rational behaviour. It is well documented that we can produce loss aversion in experiments designed to measure the value of travel time (De Borger and Fosgerau, 2008). Furthermore, there is much evidence of probability weighting in experiments involving gambles (e.g. Wu, 1996), and this would presumably extend to stated preference experiments presenting random travel time variability.[2]

It is clear that the issue of systematic deviations non-rational behaviour must be confronted. A view is emerging that actual behaviour may be seen as approximating rational behaviour, where the approximation relies on heuristics that work well on average but which may lead to clear deviations of behaviour from rationality in certain settings (Steiner and Stewart, 2015). Then the essential question from the point of view of this paper is whether the presence of travel time variability induces significant deviations from rational behaviour in the non-experimental settings that travellers encounter in their daily life. We will discuss this issue further below, after presenting some empirical findings specifically related to the value of travel time variability.

Estimating the parameters

Some empirical findings

Carrion and Levinson (2012) review a large number of studies of the cost of travel time variability. They face a number of difficult issues in comparing studies and are not able to provide firm conclusions, but their collection of estimates indicates that a reliability ratio (with respect to the standard deviation of travel time) of around 1 is plausible. The ballpark figure of 0.75 mentioned above is also plausible in the light of this evidence. Most of the empirical evidence is from stated preference studies. Revealed preference studies are emerging, primarily based on data from tolled lanes in the US (Small et al., 2005).

We shall discuss some recent stated preference studies in detail in order to assess the appropriateness of the scheduling model for such data. Hjorth et al. (n.d.) compare the step and slope models using stated preference data collected from commuters in Stockholm (Börjesson, 2008). Commuters with fixed work times were best described by the step model; this is a plausible result since that model incorporates a preferred arrival time. Commuters with flexible work times were better described by the slope model. The preferred slope model has constant utility rate at home and constant slope on the utility rate at work, which implies a constant value of travel time and a constant value of travel time variance.[3]

It is possible to estimate the step and slope scheduling models in their structural form, estimating scheduling preference parameters using observations of choices between trips with different departure times and either different travel times or different travel time distributions. It is alternatively possible to estimate reduced forms of the same models, using observations of choices between trips with different distributions of travel time. This provides an opportunity for testing the scheduling models. Börjesson, Eliasson and Franklin (2012) carry out such a test using stated preference data concerning public transport trips.

Their “structural” stated preference design lets respondents choose between trips that are specified in terms of departure time, a deterministic travel time and a cost. They estimate scheduling preference parameters from these data. Their “reduced form” stated preference design is somewhat different: respondents choose between trips that are specified in terms of an undelayed travel time, a probability of delay, a delayed travel time as well as a cost. The departure time is not specified in this design. Börjesson et al. compute the mean and the variance of travel time for each trip and use this to estimate reduced form models. They estimate both step and slope models in both structural and reduced form. The estimates show that the value of travel time is roughly the same in the structural and the reduced form models. The value of travel time variability, however, differs a lot: it is 5-10 times higher in the reduced form models than in the structural models.

A drawback of the Börjesson et al. study is that their two stated preference designs are somewhat different. The differences between the structural and reduced form estimates could be due to design differences. This is remedied by Abegaz, Hjorth and Rich (2015) who use a stated preference design with two very similar choice experiments. Both experiments comprise two alternative trips described by a cost and two potential travel times with corresponding probabilities. In addition, one of the experiments specifies the departure time for each alternative trip. The two experiments are then as similar as they can be; the only difference is whether the departure time is given to respondents.

Abegaz et al. estimate structural models on data from the experiment including the departure time attribute and reduced form models on data from the experiment not including the departure time attribute.[4] They reproduce the findings of Börjesson et al. that the value of travel time is about the same in the structural and reduced form models, but that the value of travel time variability is much higher in the reduced form models than in the structural models.

One objection that can be raised against these experiments is the following. The models that are estimated assume the scheduling preferences that underlie the stated preferences are stable and do not change during the experiments. It is however quite conceivable that respondents take into account that they are able to reschedule their activity at the end of the trip. When they are asked to choose between alternatives that specify departure times, then they could have in mind that they could act to change their schedule if they were somehow forced to depart later or earlier than their preferred departure time.

In any case, the finding of large discrepancies between the two kinds of experiments poses a serious problem, and something will have to give. There are essentially two possibilities: either we believe the theory or we believe the stated preference data at face value. It is not possible to do both.

We can think about this having Figure **4** in mind. We ultimately care about actual behaviour, but we require theory in order to develop models and to carry out cost-benefit analyses. Classical cost-benefit analysis is carried out under classical rationality assumptions as discussed above. For this analysis to be meaningful, we require that the rationality assumptions provide an approximation to reality that is not too bad.

Figure 4. **Theory and data**

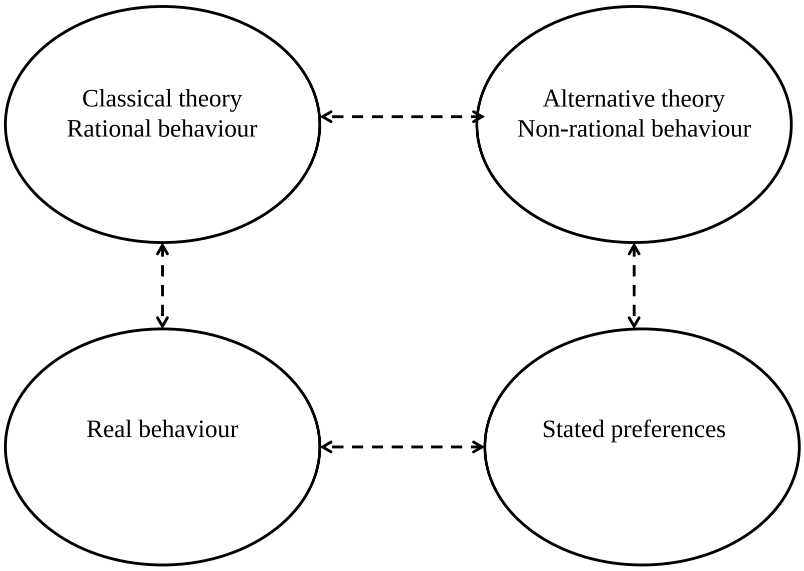

We may eventually be forced to adopt theory that incorporates non-rational behaviour. This will require fundamental changes to the way we think about and carry out cost-benefit analysis. An intermediate possibility is to introduce a distinction between classical (or hedonic) preferences and choice (or decision) preferences (e.g. Köszegi and Rabin, 2006; Steiner and Stewart, 2015). Classical preferences would capture underlying rational preferences and would enter cost-benefit analysis. Choice preferences would describe actual behaviour and would be explicitly linked to classical preferences through some account of systematic deviations from rationality. The link would be such that classical preferences could be backed out from observed choice preferences. Such a theory is conceivable but is not currently available for the case of travel time variability.

Finally, we have the behaviour that we observe in stated choice experiments. In that context we know that behaviour may deviate significantly from rational behaviour and it is by no means a given that the stated choice experiments provide valid information about real behaviour. Several conclusions are possible.

One possible conclusion is that the behaviour observed in stated preference experiments is simply not valid as a proxy for real behaviour. It is easy to argue in favour of this conclusion, especially in the context of travel time variability. The choice situations are hypothetical and unavoidably comprise a large amount of information to be digested and processed by respondents. We have little reason to believe they are actually able to do that. In that case, we should simply abandon stated preference experiments as a way of obtaining the value of travel time variability and look for ways to use revealed preference data. We may still think that people have heuristics that allow them to make reasonable choices in real context involving random outcomes, where they will have repeated experience and lots of time to learn and revise decisions.[5]

A second possible conclusion is that we should find a way to introduce the distinction between classical preferences and choice preferences that govern behaviour in stated preference experiments. Then we would use stated preference data to measure choice preferences, and then back out classical preferences to obtain values that can be used in cost-benefit analysis. The question of the validity of stated preferences regarding actual preferences would remain, however.

The third possible conclusion is to accept the behaviour we observe in stated preference experiments at face value. If evidence of non-rational behaviour is present in the data at hand, then we must account for that; we cannot just ignore this evidence. Then we will have an inconsistency between our estimates and the context in which we will use them. This is not a very attractive option.

Conclusion

Travel time variability is clearly quantitatively important. Including it in cost-benefit analysis will influence the ranking of projects and will therefore have significant real implications.

We have a simple and firm theoretical foundation in scheduling models for including travel time variability in traffic models and in cost-benefit analysis. Different models lead to different predictions that can be held against empirical evidence. We have research that provides some broader perspectives that allows us to form opinions regarding how well we might do in capturing the full costs of travel time variability.

Our ability to put numbers on the value of travel time variability can certainly be improved. We have evidence of serious problems with our stated preference experiments that are most often used to derive values. This does not mean, however, that zero is the best estimate of the value of travel time variability: not including travel time variability in analysis is not the neutral option, but will imply a bias towards projects that do not improve reliability.

Stated preference experiments, at least in the context of travel time variability, entail some fundamental problems that may be insurmountable. The reason why stated preference experiments are still popular might be that such data are comparatively cheap to collect and easy to analyse. It seems clear that practice should move towards using revealed preference data as far as that is feasible. The additional cost of acquiring data and estimating models should be held against the possibilities for improved selection of transport projects and policies. Given the very high stakes involved, the investment in data and analysis should easily pay off: after all, infrastructure is much more expensive than research and data.

The future

We are entering an era of big data. Research is now increasingly getting access to new large datasets that describe traffic conditions and actual travel behaviour. We have high-frequency measurements of speed and density at dense sets of locations in large road networks. We have datasets that track the location and time of every single train in national rail networks. We have datasets that track large numbers of GPS enabled cars, trucks and mobile devices through transport networks. We have travel card data that track travellers through large public transport networks. When the coverage of data is sufficient, we can infer travel times and travel time distributions across large networks. We can observe trip origins and destinations as well as route choices.

It is clear that it is not easy to use all this data. A lot of processing is required. But the gain, on the other hand, is clear: it is real data, not hypothetical choices. The revealed preference estimates of the value of travel time variability of e.g. Small, Winston and Yan (2005) relied on data describing a choice between just one tolled and one untolled route. To go to other countries and to other contexts, a number of challenges will have to be overcome. One challenge that will probably be common is that we prefer to have price variation in order to assign monetary values to travel time and travel time variation. Otherwise we have to infer monetary values in some indirect way, perhaps translating distance into an equivalent monetary cost, which introduces additional uncertainty. Another challenge is that we need to estimate models that describe route choice in large networks. A new generation of models is emerging (Fosgerau et al., 2013) that may be useful in this context. These new models resolve the issue that the number of possible routes in large networks is extremely large. Current practice route choice models use ad hoc devices to circumvent this problem but at the cost of systematically biasing estimates.

Making use of the emerging large datasets has many other benefits. They can be much more comprehensive than surveys and can therefore provide a much more accurate picture of traffic conditions and demand patterns than was feasible in the past. The opportunity for using these data for valuing travel time variability may come about just as a by-product of the many improvements to traffic modelling that these data make possible.

References

Abegaz, D., Hjorth, K., Rich, J., 2015. Testing the slope model of scheduling preferences on stated preference data. Working Paper.

Bates, J., Polak, J., Jones, P., Cook, A., 2001. The valuation of reliability for personal travel. Transportation Research Part E 37, 191–229.

Benezech, V., Coulombel, N., 2013. The value of service reliability. Transportation Research Part B: Methodological 58, 1–15. doi:10.1016/j.trb.2013.09.009

Börjesson, M., 2008. Joint RP–SP data in a mixed logit analysis of trip timing decisions. Transportation Research Part E: Logistics and Transportation Review 44, 1025–1038. doi:10.1016/j.tre.2007.11.001

Börjesson, M., Eliasson, J., Franklin, J.P., 2012. Valuations of travel time variability in scheduling versus mean–-variance models. Transportation Research Part B: Methodological 46, 855–873. doi:10.1016/j.trb.2012.02.004

Carrion, C., Levinson, D., 2012. Value of travel time reliability: A review of current evidence. Transportation Research Part A: Policy and Practice 46, 720–741. doi:10.1016/j.tra.2012.01.003

De Borger, B., Fosgerau, M., 2008. The trade-off between money and time: a test of the theory of reference-dependent preferences. Journal of Urban Economics 64, 101–115. doi:10.1016/j.jue.2007.09.001

Engelson, L., Fosgerau, M., 2011. Additive measures of travel time variability. Transportation Research Part B: Methodological 45, 1560–1571. doi:10.1016/j.trb.2011.07.002

Feo-Valero, M., Garcia-Menendez, L., Garrido-Hidalgo, R., 2011. Valuing freight transport time using transport demand modelling: a bibliographical review. Transp. Rev. 201, 1–27.

Fosgerau, M., Bierlaire, M., 2009. Discrete choice models with multiplicative error terms. Transportation Research Part B: Methodological 43, 494–505. doi:10.1016/j.trb.2008.10.004

Fosgerau, M., Engelson, L., 2011. The value of travel time variance. Transportation Research Part B: Methodological 45, 1–8. doi:DOI: 10.1016/j.trb.2010.06.001

Fosgerau, M., Engelson, L., Franklin, J.P., 2014. Commuting for meetings. Journal of Urban Economics 81, 104–113. doi:10.1016/j.jue.2014.03.002

Fosgerau, M., Frejinger, E., Karlstrom, A., 2013. A link based network route choice model with unrestricted choice set. Transportation Research Part B: Methodological 56, 70–80. doi:10.1016/j.trb.2013.07.012

Fosgerau, M., Fukuda, D., 2012. Valuing travel time variability: Characteristics of the travel time distribution on an urban road. Transportation Research Part C: Emerging Technologies 24, 83–101. doi:10.1016/j.trc.2012.02.008

Fosgerau, M., Karlstrom, A., 2010. The value of reliability. Transportation Research Part B 44, 38–49.

Hjorth, K., Börjesson, M., Engelson, L., Fosgerau, M., n.d. Estimating exponential scheduling preferences. Transportation Research Part B: Methodological. doi:10.1016/j.trb.2015.03.014

Hjorth, K., Ramjerdi, F., 2011. A prospect theory approach to travel time variability. Presented at the International Choice Modelling Conference.

Kahneman, D., Tversky, A., 1979. Prospect Theory: An Analysis of Decision under Risk. Econometrica 47, 263–292.

Köszegi, B., Rabin, M., 2006. A model of reference-dependent preferences. The Quarterly Journal of Economics 121, 1133–1166.

Li, Z., Hensher, D.A., Rose, J.M., 2010. Willingness to pay for travel time reliability in passenger transport: A review and some new empirical evidence. Transportation Research Part E: Logistics and Transportation Review 46, 384–403. doi:10.1016/j.tre.2009.12.005

Noland, R.B., Small, K.A., 1995. Travel-Time Uncertainty, Departure Time Choice, and the Cost of Morning Commutes. Transportation Research Record 1493, 150–158.

Oort, O., 1969. The evaluation of travelling time. Journal of Transport Economics and Policy 3, 279–286.

Polak, J., 1987. A more general model of individual departure time choice, in: Proceedings of Seminar C. Presented at the PTRC Summer Annual Meeting.

Schrank, D., Eisele, B., Lomax, T., 2012. TTI's 2012 Urban Mobility Report. Texas Transportation Institute, Texas A&M University, College Station.

Schrank, D., Lomax, T., Eisele, B., 2011. TTI's 2011 Urban Mobility Report. Texas Transportation Institute, Texas A&M University, College Station.

Small, K., 1982. The scheduling of Consumer Activities: Work Trips. American Economic Review 72, 467–479.

Small, K.A., 2012. Valuation of travel time. Economics of Transportation 1, 2–14. doi:10.1016/j.ecotra.2012.09.002

Small, K.A., Winston, C., Yan, J., 2005. Uncovering the Distribution of Motorists' Preferences for Travel Time and Reliability. Econometrica 73, 1367–1382.

Steiner, J., Stewart, C., 2015. Perceiving Prospects Properly. Working Paper.

Texas Transportation Institute, Cambridge Systems, Inc., 2006. Travel time reliability: Making It There On Time, All The Time (No. FHWA-HOP-06-070). Federal Highway Administration.

Tversky, A., Kahneman, D., 1992. Advances in prospect theory: Cumulative representation of uncertainty. J Risk Uncertainty 5, 297–323. doi:10.1007/BF00122574

Vickrey, W., 1973. Pricing, metering, and efficiently using urban transportation facilities. Highway Research Record 476, 36–48.

Vickrey, W.S., 1969. Congestion theory and transport investment. American Economic Review 59, 251–260.

Winston, C., 2013. On the Performance of the U.S. Transportation System: Caution Ahead. Journal of Economic Literature 51, 773–824. doi:10.1257/jel.51.3.773

Wu, G., 1996. Curvature of the probability weighting function. Management Science 42, 1676–1690. doi:10.1287/mnsc.42.12.1676

Note

1 It is feasible to check empirically whether the standardised travel time distribution can in fact be considered independent of departure time and how the mean and standard deviation evolves over congested peaks. For any given model, it is feasible to evaluate the numerical consequences of assuming that travel time is independent of the departure time in the case that it is not.

2 Hjorth and Ramjerdi (2011) estimate a cumulative prospect theory (Tversky and Kahneman, 1992) model on stated preference data presenting random travel times and find that extreme outcomes are over-weighted.

3 A more technical conclusion from their work is that models with multiplicative error terms were better than the standard models with additive error terms (Fosgerau and Bierlaire, 2009). There were empirical identification problems for the models with exponential utility rates, which points to a need for developing stated preference designs to strengthen identification of the slope model.

4 They use a multiplicative model (Fosgerau and Bierlaire, 2009) as that provides the best fit.

5 Stated preference experiments are routinely used to measure the value of travel time in cases where travel time is not random (Small, 2012). These choice settings do not involve random travel time outcomes and hence involve much less information, they do not require respondents to digest and process probabilities. Respondents should then be better able to make choices in these experiments than in experiments involving random travel times.

International Transport Forum
2 rue André Pascal
F-75775 Paris Cedex 16
contact@itf-oecd.org
www.itf-oecd.org