International Transport Forum

Reconciling Accessibility Benefits with User Benefits

Discussion Paper

182

Roundtable

Jonas Eliasson

Linköping University,
Sweden

Reconciling Accessibility Benefits with User Benefits

Discussion Paper

182

Roundtable

Jonas Eliasson

Linköping University,
Sweden

The International Transport Forum

The International Transport Forum is an intergovernmental organisation with 62 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, Georgia, Germany, Greece, Hungary, Iceland, India, Ireland, Israel, Italy, Japan, Kazakhstan, Korea, Latvia, Liechtenstein, Lithuania, Luxembourg, Malta, Mexico, Republic of Moldova, Mongolia, Montenegro, Morocco, the Netherlands, New Zealand, North Macedonia, Norway, Poland, Portugal, Romania, Russian Federation, Serbia, Slovak Republic, Slovenia, Spain, Sweden, Switzerland, Tunisia, Turkey, Ukraine, the United Arab Emirates, the United Kingdom, the United States and Uzbekistan.

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

ITF Discussion Papers

Cite this work as: Eliasson, J. (2020), "Reconciling Accessibility Benefits with User Benefits", *International Transport Forum Discussion Papers*, No. 2020/21, OECD Publishing, Paris.

Table of contents

Introduction

It is increasingly argued that transport policies and projects should be evaluated in terms of the "accessibility benefits" they bring, rather than via the traditional approach of valuing "travel time savings" or, more generally speaking, "user benefits".

This paper attempts to reconcile the two approaches and demonstrate that they are essentially identical, provided that the parameters of the accessibility measure are derived from the implicit preferences of travellers[1], as revealed by their travel choices, and that the same procedure to aggregate benefits accruing to different individuals is used.

This argument is formulated using a combination of reasoning and mathematical notation. The intention is to describe the reasoning in a way that is accessible with a minimum of mathematics, while at the same time indicating through mathematics how the arguments can be made exact.

Before considering the quasi-mathematical description, it is important to summarise the general line of argument. Assume that each time someone considers where, how and whether to travel, they are faced with a number of travel options (modes, destinations, etc.) as well as choice constraints. Assume that the traveller chooses their preferred travel option, subject to the choice constraints.[2] A more complicated but very useful way to formulate how the traveller chooses the preferred option is to assume that each travel option yields a *net benefit*, specific to that individual and choice situation, and that the traveller chooses the travel option that gives them the highest net benefit, subject to the choice constraints. Presumably, the net benefit of a travel option will depend on aspects such as travel time, travel cost, reliability, convenience, the quality of opportunities located at the destination, as well as characteristics linked to the traveller such as income, education, preferences, age and so on. But the concept of the benefit of an option is completely general; it is just a stylised way to talk about the traveller's implicit choice process.[3]

A traveller's *accessibility* can then be defined, very generally, as the net benefit the traveller achieves by choosing the travel option they prefer. To turn this general concept into an operational accessibility measure, the analyst must construct a way to calculate the net benefit of each option – some quantitative function of the measurable characteristics of the travel option, the choice situation and the traveller. We can think of this as determining the "parameters" of the accessibility measure. The crucial assumption is that these parameters are determined based on traveller behaviour; in the intuitive sense that if a traveller chooses one option over another, the calculated benefit of the former option should be higher than the latter. If the analyst manages to do this, they have succeeded in constructing an accessibility measure that is consistent with travellers' behaviour. Another way to formulate this is that the "parameters" of the accessibility measure reflect travellers' implicit preferences – what relative weights they put on, for example, travel times, travel costs, convenience and destination qualities when deciding which travel option to choose.

Given these definitions, the central proposition of this paper can be formulated as follows: Almost any conceivable accessibility measure, with parameters calibrated to be consistent with travellers' behaviour, will give results equivalent to user benefits when used to assess a change in variables that affect accessibility. In particular, if the change in accessibility is a reduction in travel time to a destination, the change in accessibility will equal the famous "travel time savings". But this is completely general: *any* change in accessibility will be reflected in user benefits in an analogous way. This includes, for example, changes in destinations' attractiveness, including location changes but also any measurable change in the quality of destinations, or changes in aspects of travel impedance such as convenience, security, comfort and so on – provided that the effects of these aspects on behaviour can somehow be measured (which is more of a practical problem).

Two caveats should be pointed out from the outset. First, any method for evaluating policy measures will face the problem of how to aggregate benefits (and costs) accruing to different individuals. In other words, any measure of accessibility (or user benefits) needs to specify what weights should be given to benefits accruing to different individuals. A conventional answer in economics is to aggregate different individual's benefits by summing the individual's willingness to pay for those benefits. But this is not the only answer; benefits can be weighted in an infinite number of ways. Specifying such weights is an ethical question without any clear answer, but it cannot be avoided. It is possible to discuss pros and cons with different weighting systems, such as their internal and external consistency.

Second, a change in a variable that affects accessibility often has other consequences as well, apart from the change in accessibility. For example, emissions, health, air quality, tax revenues, producer surplus, and many other societal costs and benefits may change as a result of an accessibility change. Clearly, these effects must be taken into account when evaluating a policy, but they are not part of the *accessibility* change, merely caused by it, so they are not the focus of this paper.

Equivalence between accessibility benefits and user benefits

Assume that there is a number of individuals, indexed by n, considering where and how to travel. For simplicity, let's say that we are considering a specific trip purpose: where to go shopping, for example, or where to work (which is obviously a long-term choice). One option can be not to travel at all. Individuals choose between alternative travel options j, each consisting of a combination of destination, mode, route, departure time and so on. An individual's choices are restricted by certain constraints, for example budget and time, car availability, physical ability, etc. With each travel option is associated a *net benefit*, specific to that individual and choice situation. The net benefits capture, in principle, all aspects of travel impedances and destination qualities. Presumably, it therefore depends on travel impedance variables such as travel times, travel costs, convenience and security, on destination-related variables such as the variety and quality of shops, jobs and services, and on individual-specific variables such as residential location, income, age, abilities and so on. If we consider job choices, the wage the individual would get for different jobs is also included. Let V_j^n be the net benefit of travel option j for individual n.

Assume that individuals choose the travel option with highest net benefit for them. This assumption simply captures the idea that individuals choose the option that they perceive is the best, which provides the most attractive combination of low travel impedance and high destination quality. Given this, we can define a very general accessibility measure A^n as follows:

$$A^n = \max_j \left(V_j^n \right) \quad (1.)$$

This simply says that the accessibility for individual n is the net benefit they derive when choosing the best destination/mode option (j) – the one yielding the highest net benefit, or the highest value for the travel effort, as it were. Expressed like this, we are thinking in terms of choosing *one* option for *one* trip, but the reasoning can obviously be generalised to several trips and trip purposes.

So far, this is perfectly general. The crucial assumption is that travellers choose the option they perceive as the best for them, i.e. they act according to their preferences. This assumption can of course be questioned. People sometimes act against their own best interests for various reasons, for example lack of information, lack of self-control and cognitive limitations. But the assumption of consistency between preferences and behaviour is, arguably, the only reasonable and democratically and ethically justifiable starting point. I will return to possible limitations later.

Quantifying the accessibility measure

Now, assume that we want to quantify this measure. After all, if we don't, we can never hope to measure accessibility at all. The idea is to quantify the net benefits of different travel options based on observable variables and in a way that is as consistent as possible with individuals' choices. That is, the "parameters" of the measure are "calibrated" in a way that makes them as consistent as possible with how people actually choose between various travel options.

What does "correspond to" or "be consistent with" mean in this context? Arguably, the most natural way is this: look at travel choices made by a number of individuals, and get as much as data as possible about them and their respective choice situations, travel options and choice constraints. Given this data, construct a measure of the net benefit of a travel option. Obviously, we can never hope to estimate measures that will predict people's choices perfectly. Instead, we put the calculated benefit measures into a function that calculates the probabilities that each individual will choose each particular option. The goal is to construct such a "predictor function" which "guesses right" as often as possible. This idea can be formalised into the so-called maximum likelihood method for estimating parameters of such "predictor functions" – which are, essentially, transport demand models.

Expressing the same idea formally, let v_j^n be the analyst's approximate measure of the true net benefit V_j^n. Assume that v_j^n captures some part of an individual's true net benefit of a travel option, while some part of the true net benefit is "unobservable" by the analyst. This "unobservable" part of the true benefit can be called the "idiosyncratic" benefit, or "random" benefit, since it varies idiosyncratically, or "randomly", between seemingly identical (from the point of view of the analyst) trips and individuals. The better we manage to predict individual behaviour, the smaller this "random benefit" will be, and vice versa. Expressing this idea in formulae, the true net benefit V_j^n is the sum of the measurable part v_j^n and an unobservable (to the analyst) part ε_j^n:

$$V_j^n=v_j^n+\varepsilon_j^n \qquad (2.)$$

The individual chooses the option that gives the highest true net benefit, but since the analyst can only observe the measurable part of this, only the probability P_j^n that individual n will choose option j can be calculated. Choosing different probability distributions of the "random" terms ε_j^n will yield different mathematical expressions for the choice probabilities, so one part of an analyst's task in practice is to choose these probability distributions so they reflect behaviour as well as possible.

Expressed in this way, the analyst's best measure of individual n's accessibility a^n can be written as the expected value of the highest net benefit we defined in (1.):

$$a^n = E(A^n) = E \max_j \left(v_j^n + \varepsilon_j^n\right) \tag{3.}$$

That is, individuals makes the best possible choice for themselves, given the random variables (which are known to them); but the analyst, having imperfect information, can only evaluate the benefit of this by taking the expected value over the random terms, knowing that they are known to the individual who is making the choice.

This formula can be interpreted in a slightly different way, which is equivalent for our current purposes. The random terms can also be interpreted as representing the slightly different circumstances and purpose of each trip one individual makes. With this interpretation, the formula can be interpreted as the accessibility the individual perceives when looking into the future, knowing that there will be a long sequence of trips, each with different "draws" of the random terms. When making each trip, the random terms will be known, and the choice made accordingly – but looking into the future, the individual can only evaluate their "future accessibility" by taking the expected value over all of these future choices, each with its own "draws" of the random terms.

Again, it is important to stress that despite the perhaps intimidating mathematical notation, the formulation is essentially perfectly general. We have only expressed in general mathematical notation the idea that the analyst tries to formulate an accessibility measure, capturing the net benefits of different travel options for different travellers that correspond as closely as possible to observed choices.

Equivalence between accessibility benefits and user benefits

So far, we have not really accomplished anything: we have just set up a framework to define an accessibility measure that is as general as possible. It is not at all clear how this can ever be operational or useful. To get further, we need a general and rather astonishing mathematical result, namely this:

$$\frac{\partial a^n}{\partial v_j^n} = P_j^n \tag{4.}$$

This says that if we change the net benefit of option j slightly, the corresponding change in the overall accessibility a^n will be equal to the likelihood of that option being chosen. Intuitively, this is natural: if the option is really good, with a high likelihood of being chosen, then the overall accessibility changes a lot when the net benefit of that option is improved, and vice versa. What is astonishing is that the result is completely general; it holds for virtually *any* kind of accessibility measures of the general type above, which are only built on the simple assumption that the traveller chooses the option with the highest net benefit. Fosgerau, McFadden and Bierlaire (2013) provide a proof.

With this in hand, we can obtain the central and rather remarkable result that equates a change in the completely general accessibility measure defined above with user benefits. The idea will be easier to explain if we take a very simple example, just to reduce the complexity of the mathematical notation. Consider a specific travel option consisting of taking the bus to a certain shopping centre, and assume that the net benefit of that travel option can be measured by a weighted sum of the number of shops at the destination S_j, the travel time with the bus t_j and the waiting time at the bus station w_j. In addition, the net benefit depends on a large number of other variables which we do not (or cannot) specify but denote by x_j.

That is, we have the following measure of the (observable) net benefit of this travel option:

$$v_j^n = \alpha_j^n S_j - \beta_j^n t_j - \gamma_j^n w_j + x_j^n \tag{5.}$$

Note that the parameters α, β and γ are different for each individual, since everyone is different. The parameters α, β and γ measure how the benefit of the travel option changes when the corresponding variable changes, hence they are referred to as the *marginal* net benefit of the corresponding variable. Again it must be stressed that this particularly simplistic formula is used to make the exposition simpler.

Consider an improvement of the bus travel time from t_j^0 to t_j^1. Divide the improvement into tiny incremental steps Δt_j. Use the definition of a derivative to see that the corresponding incremental increase in overall accessibility Δa^n will be

$$\Delta a^n = \frac{\partial a^n}{\partial t_j} \Delta t_j \tag{6.}$$

Using (4.) and (5.), we can rewrite this as[4]

$$\Delta a^n = \frac{\partial a^n}{\partial v_j^n} \frac{\partial v_j^n}{\partial t_j} \Delta t_j = P_j^n \beta_j^n \Delta t_j \tag{7.}$$

The total increase in overall accessibility B^n of the entire improvement from t_j^0 to t_j^1 will be the sum of all the incremental accessibility increases Δa^n, and by the definition of an integral we get

$$B^n = \int_{t_j^0}^{t_j^1} P_j^n \beta_j^n dt \tag{8.}$$

This describes the increase in accessibility for individual n, or in other words the accessibility benefit of the improvement. We then need to specify some procedure to aggregate the benefits accruing to different individuals. The next section considers different procedures, but for now, assume that the net benefits of different individuals have been scaled in some way so that we can sum accessibility benefits over individuals. This gives the total accessibility benefit of the travel time improvement:

$$B = \sum_n \int_{t_j^0}^{t_j^1} P_j^n \beta_j^n dt \tag{9.}$$

Since the choice probabilities P_j^n depend on t_j, evaluating the integrals is not necessarily easy.[5] But if P_j^n depends approximately linearly on t_j in the interval from t_j^0 to t_j^1, then the integrals can be approximated as follows:

$$B \approx \sum_n \frac{P_j^n(t_j^1) + P_j^n(t_j^0)}{2} \beta_j^n (t_j^0 - t_j^1) \tag{10.}$$

Finally, let's rewrite the formula a little just to clarify what it says. $\sum_n P_j^n(t_j^0)$ is the number of people choosing option j before the improvement; call that F_j^0. Let ΔF_j be the increase in travellers choosing option j, so we have $\Delta F_j = \sum_n P_j^n(t_j^1) - P_j^n(t_j^0)$. $\frac{\Sigma_n P_j^n(t_j^0)\beta_j^n}{\Sigma_n P_j^n(t_j^0)}$ is the average marginal net benefit of the bus travel time improvement (remember that β_j^n measures how much the net benefit of travel option j changes when the bus travel time changes); call this β_j. Further, assume (to reduce the notational complexity) that

$\frac{\Sigma_n P_j^n(t_j^1)\beta_j^n}{\Sigma_n P_j^n(t_j^1)} \approx \beta_j$, i.e. the average marginal net benefit does not change appreciably when some new travellers switch to option j after the improvement.[6] With this simplified notation, we get

$$B \approx \beta_j\left(t_j^0 - t_j^1\right)\left(F_j^0 + \frac{1}{2}\Delta F_j\right) \tag{11.}$$

This is the "rule of a half". The particular expression in (11.) refers to a travel time improvement, since this was our example, but the rule of a half formula can be applied to *any* kind of improvement. For example, consider an increase of the number of shops at the destination from S_j^0 to S_j^1. An analogous derivation shows that the accessibility benefit becomes

$$B \approx \alpha_j\left(S_j^1 - S_j^0\right)\left(F_j^0 + \frac{1}{2}\Delta F_j\right) \tag{12.}$$

This means that we can value this accessibility increase just as easily as a travel time improvement, as long as we can estimate the effect of "more shops at a destination" on travellers' propensity to travel there.

Another example: consider a policy that makes waiting for the bus more comfortable – better benches, lighting and information, perhaps. Assume that this decreases the effect that bus waiting time has on the net benefit: the parameter γ measuring the effect of bus waiting time is reduced from γ_j^0 to γ_j^1. An analogous derivation shows that the increase in overall accessibility of this improvement becomes:

$$B \approx \left(\gamma_j^0 - \gamma_j^1\right) w_j \left(F_j^0 + \frac{1}{2}\Delta F_j\right) \tag{13.}$$

These examples show that essentially any variable that affects accessibility can be captured by the rule of a half, as long as the analyst is able to estimate the effect that the variable has on travel behaviour. Of course, this can present a significant practical challenge. On the other hand, this is necessary for virtually any kind of (quantitative) accessibility analysis. After all, if we have no idea of how a variable affects travel behaviour, there is no hope of saying how it affects accessibility.

If (and only if) we agree that benefits accruing to different individuals should be weighted with the inverse of a person's marginal utility of money, the parameters β and γ will be the familiar "(monetary) value of time" and "(monetary) value of waiting time", while α will be (less familiar) "monetary value of shops". If we choose some other weighting system – that a minute is worth the same for travellers, for example – the "valuations" of benefits α, β and γ will change; but the general principle of the rule of a half still applies. We are just changing the "currency" in which benefits are measured from money to, say, minutes. Analysts or decision makers are free to introduce any kind of weights for different individuals' benefits that they feel are justifiable. How to weight benefits across individuals is one of oldest debates in welfare economics in general and transport economics in particular. The section below contains a very brief summary of the arguments most relevant for transport. A much longer discussion, focusing on whether and how valuations of travel time should be differentiated, can be found in Börjesson and Eliasson (2019).

Practical advantages of the rule of a half

The advantage of the rule of a half is that although it is derived from a very general expression for accessibility, it depends only on quantities that can be measured either directly or indirectly from observable data. Accessibility variables such as travel times t_j, number of shops S_j, waiting times w_j and so on can be measured directly. F_j^0 is the number of people choosing option j – in this example, choosing to go to this particular shopping centre by bus – which is also easy to measure. ΔF_j is the increased number of travellers choosing option j after the improvement, which can usually be estimated approximately. Even better, if we want to evaluate a policy *ex post*, ΔF_j can be measured directly. Finally, the parameters α_j, β_j and γ_j are the (average) effects of changes in the respective variables on choice probabilities, estimated from travellers' observed choices as described above.

A convenient observation is that the term $\frac{1}{2}\Delta F_j$ is often small in practice (although not always). A substantial decrease in, for example, a bus travel time, may increase the number of passengers by, say, 20%. This means that the term $\frac{1}{2}\Delta F_j$ represents an addition of 10% to the total benefits. This shows that if the forecast of new passengers is wrong by some moderate amount, it usually does not have a substantial impact on total benefits.

Another key feature of the rule of a half is that it does not depend on the multitude of other variables that also affect net benefits and hence choices, or on all the other possible travel options. Hence, we do not need a complete description of how net benefits depend on a host of different variables; we only need to be able to isolate the effect of the variable we are improving. Even if this can be difficult in practice, it is infinitely easier than having to specify a complete description of how any conceivable variable affects the net benefit of any option.

Moreover, the formula only depends on the number of travellers choosing the option we are improving, that is, option j. This is also highly important, since the number of "travel options" is almost infinite: any combination of mode, destination, departure time route and so on is a "travel option". What the formula above says is that we do not need to keep track of all the other choices travellers make and how they change; it is enough to focus on the option that we are improving.

Remember that in the derivation of the rule of a half formula, we assumed that there was in principle an extremely large number of travel options, and many variables that affected their benefits, and we allowed the quantification of benefits to be as complicated as needed. But from this very general framework, the "virtual" complexity vanishes, and we are left with an expression that only depends on data that is either directly measurable or is relatively easy to estimate.

Why is it the rule of a half?

A common question is why the benefits of new travellers - ΔF_j – in the formula is only counted as a "half". In other words, why is there a factor ½ in the expression $\frac{1}{2}\Delta F_j$? The mathematical answer is that it follows from the approximation of the integral, in the step from (9.) to (10.). But there is also a natural intuition explaining why the "half" factor is there. The users *not* choosing option j before the change must, by assumption, have some other option that they prefer – otherwise they would also have chosen option j. In principle, some users must be close to indifferent between option j and this other option, so when option j is improved ever so slightly, they switch from the other option to j. Then they get almost the entire benefit of the remaining improvement. Similarly, some users are almost exactly indifferent between j and some other option even when the improvement has been made, but prefer j by a small amount. When they switch to j they therefore get almost no benefit at all – their other option was almost exactly as good, by assumption. Imagine now a continuum of users switching from other options to j. On average, they will get half the benefit of the improvement – some switched to j almost immediately, and some were virtually indifferent between j and some other option even after the improvement.

Assumptions underlying the equivalence

The most noteworthy aspect of the rule of a half is how general it is, in the sense that only some the following, modest assumptions are necessary to employ it:

Travellers choose the travel option that gives them the highest benefit, given their choice set. Their choice set may be constrained by income and time constraints, abilities, car ownership, etc. This is an assumption about people's behaviour, and allows us (in principle) to estimate the parameters of the accessibility measure.

The parameters of the accessibility measure should be derived from people's own implicit preferences, as implied by their choices. In other words, an accessibility increase should be valued in the same way that individuals themselves would value it, and not, for example, as politicians or other stakeholders would do. This is an ethical rather than an empirical statement; it says that the analyses we undertake, and the results they produce, should ultimately be based on people's preferences, rather than what decision makers would like.[7]

The analyst must be able to infer how a variable affects travel behaviour, if a change in that variable is to be studied. Influences of variables that remain unchanged, however, do not need to be known.

The analyst must specify a procedure to aggregate benefits accruing to different individuals, for example by attaching weights to different individuals and then calculate a weighted sum of benefits. The conventional way to do this is to convert all benefits to monetary units, i.e. willingness to pay, and there are some good reasons for this (see below). But the weighting system can be arbitrary, non-linear and very general. For example, it can specify that benefits accruing to persons already enjoying accessibility above a given threshold are valued at zero, or that benefits accruing to low-income groups or specific residential areas are weighted higher.

The choice probabilities P_j^n need to be approximately linear in the relevant interval. In other words, the change in the number of travellers as a response to the change in a parameter must be roughly proportional. However, this is only necessary to go from the user-benefits-integral (9.) to its linear approximation (10.). If the relevant choice probabilities cannot be linearly approximated, then the integral needs to be used instead – but this is doable, although more cumbersome.[8] The most common case when this needs to be done is when a previously non-existent travel option is introduced.

Is willingness to pay a defensible way to aggregate benefits across users?

As pointed out above, it is not necessary to convert benefits into *money* before aggregating across individuals. We could convert benefits into minutes, or do some non-linear transformation of benefits (say, put the value of benefits over a certain threshold at zero) before aggregating. But there are in fact some good reasons to use money as the "currency" in which to measure benefits when aggregating them.

It is true that converting benefits into monetary terms will put a higher value on benefits accruing to rich people, simply because they are able to pay more for the same benefit. See Galvez and Jara-Díaz (1998) and Börjesson and Eliasson (2019) for a formal expression of this argument. However, there are two good arguments for this aggregation method. First, society can arrange some compensation mechanism to compensate any unfair (however that is defined) distribution of benefits, and money is easier to redistribute than, say, time. Second, benefits will in general not stay with the initial recipient; the price system in the economy will redistribute large shares of the benefits from travellers through changes in rents, land prices, wages, fares, prices of goods and services, for example, to land owners, transport operators, firms and taxpayers. And these benefits spreading through the economy are all mediated through willingness to pay, i.e. "real money". An example is transport fares: they are either under public control, in which case they can be used to transfer money back to the traveller population (and in that case those with lower incomes will gain more for the same reasons: they have a higher marginal utility of money), or they are under the control of a commercial enterprise (such as a commercial train operator), and in that case fares will adjust to capture some of the benefits initially accruing to travellers. In both cases – and many other similar ones – compensations, public policies and transfers of benefits are all mediated through prices and money.

The issue of how benefits should be weighted across groups is discussed at length in Börjesson and Eliasson (2019). The discussion there centres on whether the value of travel time savings should be differentiated with respect to, for example, mode, purpose or time of day. That paper shows that it depends on the circumstances, in particular whether travel costs (fares) are under public control and to whom benefits accrue in the long run. In some choice situations, the value of time travel savings should be adjusted for differences in income, but it is important to always take into account differences in marginal utilities of time (e.g. across travel time components, modes and trip purposes). Moreover, the study demonstrates (using Swedish data) that income is *not* the main source of differences in the valuation of time savings; the variation due to differences in marginal utilities of time turns out to be much larger.

Can people really choose the options that are in their best interests?

Basing the valuation of benefits on people's behaviour rests on the assumption that people are actually able to choose the option that is best for them – that brings them the highest benefit or utility. But is this really true? A number of counter-arguments are worth considering.

First, even if individual choices are rational from the point of view of each individual, they may well result in aggregate outcomes that are not in individuals' best interest. This is the tragedy of the commons problem. In transport planning, a common example is when an improvement in road capacity leads to more car trips, which leads to more congestion, which results in the initial benefits of the capacity improvement vanishing, and travel times eventually reverting to what they were before the improvement. Similar situations, where the sum of individually optimal decisions leads to aggregate outcomes that are suboptimal, are widespread: the climate problem is the most important one. But this does not invalidate the possibility of deriving weights for different benefits and costs based on an individual's behaviour; it just means that individuals' decisions must be modified so as to take external effects into account.

Other negative external effects include noise, accidents and emissions. There are also positive external effects, such as the increase in tax revenues when someone goes from unemployment to employment (the benefit for the individual is the after-tax wage but, in addition, net tax revenues increase). All such effects need to be accounted for, and ideally internalised so they are taken into account by individuals, or at least taken into account in policy evaluation. But this is also a separate consideration from the question of whether individuals are able to make decisions in *their own* best interest.

Difficulties arise when it comes to various human cognitive limitations, limited self-control, and sometimes lack of information. The latter is easiest to manage. Cognitive and self-control limitations, however, raise more ethically and philosophically difficult problems. There are a host of studies in psychology, social psychology and behavioural economics that show that people do not always act consistently or in their own best interest, and that their preferences are not always stable (which makes it hard to even define a concept like their own best interest). Some difficult issues from the transport sector include whether people can accurately judge and take into account traffic risks, or the health effects of active modes, living in noisy areas or being exposed to air pollution. Several studies have shown that people are indeed aware of such aspects and take them into account – but whether they do this *accurately* is another question. People may well under- or overestimate the health effects of cycling, for example; we only know that they take them into account (Börjesson and Eliasson, 2012).

However, the null hypothesis must be that people are, in general, able to choose what is best for them. In other words, the burden of proof must lie on the planner or decision maker who argues that he knows what is good for people better than they know themselves (or at least, the way they act themselves). There are in fact quite a few such examples in the transport safety sector, such as compulsory safety belts, since there are some good arguments that people are really bad at judging small risks, with disastrous consequences. But there is not much evidence that people are not able to judge accessibility-related variables, such as travel times versus waiting times or travel costs. This is most likely that the feedback from travel choices is almost immediate: if something takes a long time, or is inconvenient, travellers notice this almost at once. This is very different from other choices such as getting an education or saving for retirement, where feedback may take decades. In such situations, the evidence for people acting against their own best interest (in the long run) is stronger. But deriving parameters of accessibility measures from observed behaviour seems to be well motivated.

Do people really have a choice, or is their behaviour completely determined by constraints?

Another argument against deriving accessibility parameters from observed behaviour is that people do not have any choice. But this is almost never true, except perhaps in the very short run. People decide how to spend their money, time and other resources on a combination of housing (location and size), travel and other types of consumption. There are obvious trade-offs between all of these: by saving on thing, more resources can be spent on something else. People certainly have choice *constraints* – time and money constraints, for example – but that is something else. In the short run, constraints are of course more binding, but in the long run, it seems silly to claim that people do not have any choices at all.

Conclusion

The purpose of this paper was to explain precisely what assumptions need to be made in order to arrive at user benefits as a measure of accessibility and thereby enable a more precise debate about whether these assumptions are reasonable. This makes it easier to specify exactly *which* of these assumptions are not valid (if any). This can lead to replacing these assumptions with a new set and following them to their logical conclusion.

Notes

1 Throughout, the focus will be on personal travel, mainly for ease of exposition. However, the arguments and methods are equally applicable to freight transport.

2 One criticism says that travellers essentially cannot choose, in other words that their choice set is so limited that there is only one travel option; travellers have to travel exactly the way they do. This will be dealt with later however this does not seem to be a realistic description of travel behaviour, at least not in general or in the long run.

3 In particular, it should not be interpreted as a literal description of the actual psychological decision-making process.

4 If we had chosen a more general expression of the net benefit than the simplistic one in (5.), the parameter β_j^n would be replaced by the partial derivative $\frac{\partial v_j^n}{\partial t_j}$, and correspondingly for the following formulas where the parameters α_j^n and γ_j^n appear.

5 If the choice probabilities are so-called logit functions, however, the integral is possible to solve analytically, the result is the well-known logsum formula.

6 This assumption is actually inessential, but it simplifies the formulas and their explanation substantially.

7 This is a deep and profound difference between cost-benefit analysis (CBA), which is ultimately based on citizen preferences, and most (although not all) kinds of multi-criteria analysis (MCA), where weights of various objectives and benefits are derived from decision-makers' (or stakeholders') preferences.

8 In particular, if the choice probabilities are logit functions, the integral will be the logsum.

References

Börjesson, M. and J. Eliasson (2019), "Should values of time be differentiated?" *Transport Reviews*, *39*(3), 357–375, https://doi.org/10.1080/01441647.2018.1480543.

Börjesson, M. and J. Eliasson (2012), "The value of time and external benefits in bicycle appraisal", *Transportation Research Part A: Policy and Practice*, *46*(4), 673–683, https://doi.org/10.1016/j.tra.2012.01.006.

Fosgerau, M., D. McFadden and M. Bierlaire (2013), "Choice probability generating functions", *Journal of Choice Modelling*, *8*, 1–18, https://doi.org/10.1016/j.jocm.2013.05.002.

Galvez, T.E. and S.R. Jara-Díaz (1998), "On the social valuation of travel time savings", *International Journal of Transport Economics*, *25*(1), 205–219.

Reconciling Accessibility Benefits with User Benefits

This paper asks whether transport policy assessments should use accessibility benefits as a key measure instead of user benefits. It argues that both measures are equivalent if accessibility measures are based on transport users' own preferences and if the same principle is used to aggregate benefits. The paper also addresses how distributional questions can be addressed within this approach.

All resources from the Roundtable on Accessibility and Transport Appraisal are available at:
www.itf-oecd.org/accessibility-and-transport-appraisal-roundtable

International Transport Forum
2 rue André Pascal
F-75775 Paris Cedex 16
+33 (0)1 73 31 25 00
contact@itf-oecd.org
www.itf-oecd.org

Published: 11/2020 | Photo credit: Robert_Chlopas/Shutterstock