Oxera

Agenda

Advancing economics in business

Follow the White Rabbit: do people care if their train is late?

The White Rabbit in *Alice in Wonderland* was very upset by his lateness, but are his concerns shared by others? Transport policy is often guided by a desire to improve the predictability of journey times, but how much do people care about this in practice? This article examines the ways this has been assessed and, reflecting insights from behavioural economics, questions the value attributed to reducing travel time uncertainty

People's behaviour changes with uncertainty. They move money out of a bank account and into a government-backed savings account in the face of concerns about a bank's solvency. Firms invest less as uncertainty increases and the value of not investing rises; and people catch an earlier train if they are worried about missing their flight from the airport.

These statements would all seem to reflect typical behaviour, and are themselves reflected in the desire of policy-makers and others to reduce uncertainty. Transport is no exception. For example, a key aim of the Westminster and Scottish governments' plans for the outputs of the countries' railways over the five years from 2009 is to improve reliability. Based on the Public Performance Measure (PPM), the governments are funding improvements to ensure that at least 92% of trains arrive within five minutes (or ten minutes for long-distance services) of the scheduled time.[1]

But what is the evidence underlying this funded improvement? Do people respond to reduced uncertainty in a manner that is consistent with traditional economics and that assumes fully rational behaviour? How does this evidence stand up to the insights from behavioural economics?

This article sets out traditional economic thinking around how people respond to uncertainty. With this in mind, it reviews empirical evidence in relation to the reliability of rail journey times, before questioning this evidence in light of results from other parts of the economy and the insights from behavioural economics. It suggests that, by taking this thinking on board, we can understand people's preferences more in relation to reliability, and their ability to make fully rational decisions when faced with choices. Overall, more evidence is needed before taking substantive policy decisions.

Figure 1 An example of decision-making using expected utility

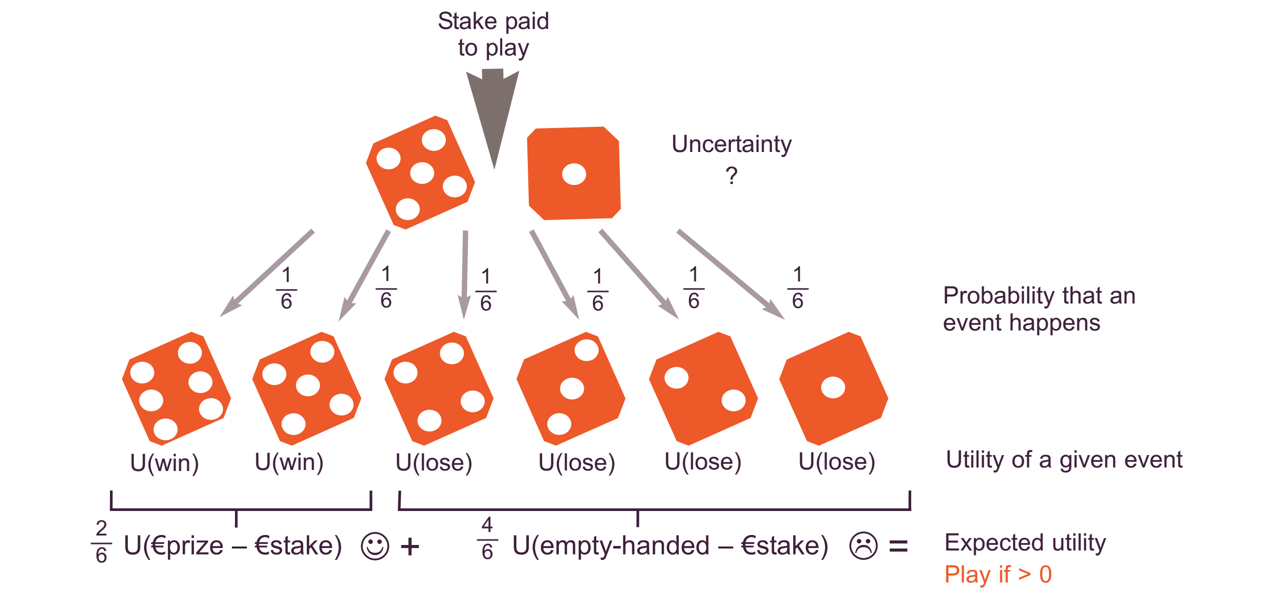

Happiness and risk

The effects of uncertainty on people's behaviour are commonly modelled in economics using 'expected utility' theory. This assumes that people value uncertain situations according to the pleasure (or, less emotively, utility) of the different possibilities that they think could occur, weighted by their assessment of the likelihood of these occurring.

Figure 1 above gives an example of this. An individual is offered the chance to pay a stake in return for which they will, depending on the throw of a dice, win a cash prize (if it is a 5 or a 6), or come away empty-handed (if it is any other number) having lost the stake. For every possible dice roll there is a monetary outcome from which the individual will gain a certain level of utility. If that individual behaves according to expected utility theory, they will play the game provided their expected utility from the bet is positive—the expected utility being equal to the utility of each payoff weighted by the associated probability of it occurring (as shown in Figure 1).

This example may seem very different from many economic decisions. The decision is yes/no, as opposed to there being a range of options; there is a finite range of possibilities that can happen, when in practice these can be infinite; and the probabilities of the different outcomes are known. However, it is possible to extend the framework to deal with these issues, at least in part. The approach is therefore widely used to model behaviour involving uncertainty, including, among other things, gambling, insurance, investment decisions (eg, how much to invest in stocks or bonds), and strategic behaviour (such as negotiations and auctions).

An integral aspect of many applications is that the framework can be used to measure how people's happiness or welfare is affected by the risks to which they are exposed—ie, in technical terms, how risk-averse they are.

Risk aversion

The extent to which people like (or dislike) risk is measured according to whether they wish to be compensated for bearing it or, depending on the circumstances, whether they are prepared to pay to avoid it.

This is usually defined against a benchmark known as the expected outcome of the uncertain situation, which is equal to the possible payoffs that could occur, weighted by the associated probabilities (in the example this is 2/6€prize – €stake). This is generally different from the associated expected utility.[2] Attitudes to risk are commonly classified into three categories. An individual would typically be described as one of the following.

- **Risk-averse: they require compensation for bearing risk**—eg, a person would want to participate in the dice game (ie, their expected utility > 0) only if the expected outcome of the game was positive (2/6€prize > €stake). Effectively, if they played the game repeatedly, the prize for throwing a 5 or a 6 would need to be sufficiently large relative to the stake so that, on average, they would expect to make money.[3] Risk aversion is one of the reasons why the return on holding a portfolio of shares in the stock market is, on average, higher than that for government bonds.

- **Risk-neutral: they require no compensation for the risk**—in the game above, they would be prepared to participate where the expected payoff is zero (2/6€prize – €stake = 0). The effect of the game on the players' incomes would therefore be zero on average, although there would still be uncertainty as to what their income from any one game would be.

- **Risk-loving: they are prepared to pay extra to face risk**—in order to participate they would pay a stake more than the dice game's expected outcome (€stake >2/6€prize). That is to say, they would expect to make a loss on the game if played repeatedly. It is gamblers' risk-loving nature which ensures that gaming organisations such as casinos make money.[4]

Clearly, people do not have utility neurones that give a numerical pleasure level for different outcomes—this is simply intended as a model of their behaviour. Indeed traditional economics would argue that absolute utility—or happiness—is not actually measurable. However, information on individual preferences (that is to say, people's relative valuation of outcomes) can be inferred from empirical data. This can be obtained from surveys, or 'experiments' where individuals' decisions are assessed in a laboratory setting in response to different scenarios, such as the dice game. Estimates of risk aversion can be obtained from the data by finding the type of preferences that best explain their observed behaviour. A number of practical measures of risk aversion that are used are discussed in the box below, including some estimates that have been obtained from different situations.

Valuing reliability

The discussion above sets out an economic framework for understanding how people treat risk in their everyday lives, and also provides some valuations of the degree of risk aversion exhibited in different situations. This section describes an example of this framework in use—rail

Estimates of risk aversion

The extent of risk aversion is usually measured by the premium that an individual is prepared to pay/receive to avoid/obtain a certain level of risk of risk exposure. Standard measures of this are as follows.

- **Absolute risk aversion (ARA)—this is a number that links the premium that an individual would be prepared to pay to avoid a risk to the size of the risk, where that risk is measured by the variance (a numerical measure of the extent of the uncertainty). The greater the measure of absolute risk aversion, the greater the premium that a person would be prepared to pay to avoid a given risk (ie, the more risk-averse they are).**
- **Relative risk aversion (RRA)—a limitation of absolute risk aversion is that it is likely to depend on the income level that an individual has, rather than being a measure of risk aversion that is easily comparable across individuals. Therefore, RRA is often used, which links the size of the premium, relative to income, to the variance of the risk, relative to income.[1] The assumption of constant relative risk aversion (CRRA) is frequently used in economic modelling. In CRRA, individuals' relative risk aversion does not change as their income increases.**

As examples of how risk aversion can vary in different contexts, some estimates of RRA from the academic literature are presented as follows.

- **2.6–10.2 (average 4.37)—in European collateralised debt obligation (CDO) markets.[2]**
- **1.8—3.2—game-show contestants on the Australian version of the *Deal or No Deal* television programme.[3]**
- **1.0—participants in the UK version of *Who Wants to be a Millionaire?* television game show.[4]**
- **0.67—the average of adult Danish participants in experimental tests.[5]**
- **Around 0.5—bidders in an experimental auction.[6]**

Notes: [1] Technically, this is the ratio of the variance to the mean income squared.
[2] Weber, T. (2008), 'Default Risk Premia in Synthetic European CDOs', Working Paper, Universität Konstanz, p. 22.
[3] De Roos, N. and Sarafidis, Y. (2006), 'Decision Making Under Risk in *Deal or No Deal*', University of Sydney and CRA International working paper, p. 25.
[4] Hartley, R., Lanot, G. and Walker, I. (2005), 'Who Really Wants to be a Millionaire? Estimates of Risk Aversion from Gameshow Data', Warwick Economics Department Research papers (2005), p. 42.
[5] Harrison, G.W., Lau, M.I. and Rutström, E.E. (2007), 'Estimating Risk Attitudes in Denmark: A Field Experiment', *Scandinavian Journal of Economics*, **109**:2, June, pp. 341–68.
[6] Goeree, J.K., Holt, C.A. and Palfrey, T.R. (1999), 'Quantal Response Equilibrium and Overbidding in Private-value Auctions', *Journal of Economic Theory*, **104**, pp. 247–72.

passengers' willingness to pay to increase journey time predictability—and presents recent research suggesting that the degree of risk aversion exhibited by rail passengers in surveys is much lower than that seen in situations presented in the box above.

Conceptual framework

Analysis in this area has to consider two aspects of behaviour—how people respond to train service delays, and how they form expectations about the degree of delay that will occur during their journey. The combination of these two aspects will lead people to leave earlier than they otherwise would to reach their destination on time. In addition to this response, research has found that people also place a negative value on variability in their journey times per se—measured either by the mean or standard deviation of delays. Both responses will be more substantive the greater the risk aversion of the passenger in question.

Researchers have chosen to represent the way in which delays affect demand by treating lateness minutes as if they were additional to journey time, potentially with a further weighting to reflect the inconvenience associated with the delay. For example, if every minute late were valued by passengers as if their journey were three minutes longer, a five-minute delay on a 30-minute service would be treated under this framework as if the journey were actually 45 minutes. For demand forecasting purposes, the impact on demand of a reduction in delay by five minutes on average per journey is converted into a reduction in overall journey times of 15 minutes (using the weighting of three), and then a demand increase using an elasticity to journey time changes.

The way researchers—for example, Bates et al. working for the British rail industry's Passenger Demand Forecasting Council (PDFC)—develop this evidence is through the use of market research, where passengers (in this study recruited from trains, but interviewed at home) are offered a choice between routes with similar average delays to their rail route, and those offering better or worse punctuality, but with higher or lower fares.[5] These responses are then aggregated across the survey sample, with econometric techniques used to elicit the additional minutes that passengers perceive are added to their journeys by delays.

Empirical estimates

The Bates et al. study described above produces a range of estimates of the journey time penalty that passengers—in the world of these models—add onto the actual time of the rail journey when making decisions. As expected, this varies according to journey purpose, with passengers travelling to the airport treating delays as if they had made the journey very much longer, for example. Of the two effects noted above—the response to the mean or standard deviation of delay, and catching an earlier train than would otherwise be the case—the former accounts for nearly 84% of the currently recommended journey time 'penalty' in demand forecasting guidance, with the latter making up the difference.[6]

What does this evidence say about peoples' risk aversion in relation to train delays? The results are surprising. A recent study has taken another look at the Bates et al. data and estimated coefficients of absolute risk aversion displayed by the participants in the market research.[7] While estimates in other situations described above range from 0.5 to 10.2, depending on the context, the researchers have discovered that the Bates et al. participants displayed risk aversion levels of only 0.12—at least an order of magnitude lower than most estimates in other contexts.

Are the valuations biased?

The previous section demonstrated that the degree of passenger risk aversion, in relation to rail performance, implied by the answers provided to market researchers, is quite low compared with the degree of risk aversion seen in other situations. So do people really not care if their train is late? If this is the case, why have the Westminster and Scottish (and, presumably, the Japanese and Swiss) governments invested heavily in improving punctuality and reliability on the railways, especially as there are likely to be situations where rail delay as a proportion of total door-to-door journey time is small?

Perhaps, though, the degree of risk aversion implied by these surveys is not what it seems. A few possible explanations are provided below.

– **Self-insurance**. By catching a train (or trains) earlier than they would otherwise if the train were guaranteed to arrive on time, people can insure themselves against arriving late at their next appointment. The cost of this insurance is the value of the time they would otherwise be spending doing something else. If the survey evidence were capturing people's ability to self-insure in this manner, it is plausible that, with some of the performance risk insured against, the degree of risk aversion relative to other situations may be expected to be lower. The policy response to this behaviour is, presumably, to invest in punctuality and reliability in rural areas where trains are less frequent (although the numbers of potential beneficiaries are likely to be smaller), and also on routes where a high proportion of passengers have high values of time (eg, airport routes or routes carrying large numbers of business travellers[8]). The counterpoint to this argument is that the evidence suggests that only a small proportion of the impact of delays on passengers is attributable to people having to catch an earlier train.

– **Not 'in the moment'**—the Bates et al. study used survey data collected in interviews in participants' homes, with respondents recruited from a variety of train services. Market research of the sort carried out to support willingness-to-pay valuations of service attributes in a sector—be it rail in relation to train performance, or water in relation to sewer flooding—is likely to find it difficult to capture the real attitudes of people to a negative outcome. The survey is unlikely to get people 'in-sample' who have missed their flight as a result of a delayed train, or whose house has been flooded with sewage as a result of a burst pipe. Perhaps more importantly, it is difficult to prompt people to remember (or imagine) their emotions in these situations when responding to such a survey.

– **'Survey rationality'**—on a related point, respondents may answer survey questions more 'rationally' than they actually behave in real life. Behavioural economics shows that, in many situations, people are 'reference-dependent' and 'loss-averse'. In a train punctuality context, it might therefore be expected that, around some 'reference' journey time that passengers become accustomed to, people dislike lateness (loss) much more than they value a train being ahead of schedule (gain). People are also more likely to dislike lateness if they have experienced it very recently. However, in a relaxed atmosphere for completing market research in relation to train performance—of a railway station, or at home, perhaps—people may have more time to reflect on the one issue of avoiding unreliable trains than they otherwise would have in the course of a day. Emotive, intuitive decision-making may not be captured adequately, so, arguably, these circumstances are unlikely to lead to people placing very high values on avoiding rail unreliability, unless their experience of the day has been extremely bad (in which case the survey may not be undertaken anyway).

– **Evidence from revealed-preference data**—one way of skirting the issue of people's inability to be 'put in the moment' when surveyed about train reliability is to estimate valuations based on actual behaviour. Oxera has estimated elasticities to train delays in previous studies for the GB rail industry, finding that the relative valuations of avoiding delay suggested by stated-preference studies are borne out.[9] However, the elasticities estimated were slightly higher than those implied by previous survey-based approaches, suggesting that:[10]

– either some of the potential biases (suppressing valuations) highlighted above arising from survey-based approaches were avoided, and a better representation of passenger responses to unreliability were obtained; or

– the assumptions (which treat unreliability as an extension to journey time, with a multiplier reflecting inconvenience caused, and then using a journey time elasticity) used to obtain elasticities implied by the previous survey evidence are incorrect.

Conclusions and a way forward

Overall, the evidence base in relation to the valuation by users of improvements in train performance is not as compelling as may have been hoped for given the investment required to maintain and improve levels of punctuality. Arguments seem to be required instead that demonstrate the impact on business and modal shift of the funded improvements.

However, the behavioural economics literature is providing new hypotheses for economists to test that are intended to reflect expected disparities between people's actual behaviour and what they say they will do when responding to a survey. For example, a survey that prompts people to think back over recent experiences in relation to train punctuality before answering questions on how they would trade off price, punctuality and journey duration may take issues of reference dependence, loss aversion, and emotive decision-making into account more effectively.

Perhaps a more radical idea would be to consider techniques from the psychology literature and develop choice experiments where individuals use real cash provided to them with which to buy rail tickets under different rationales for the rail journey (having to catch a plane, visiting an elderly relative, meeting a new client). The experiment could then be repeated over a number of iterations where participants receive data (which would be derived from actual train performance) on their chosen train's performance and then, if they deem it necessary, adjust their behaviour for future 'rounds'. This exercise might then omit biases that could be present in a number of the analytical approaches used above, although it is clear that participants might also, in such a 'clinical' situation, focus more on decision-making in respect of train punctuality than they actually would do in reality.

---

[1] See 'The First ScotRail Passenger Rail Franchise: Prepared for the Auditor General for Scotland', November 2008; and Select Committee on Transport', 'Memorandum from Network Rail (RWP 04)', available at: http://www.publications.parliament.uk/pa/cm200708/cmselect/cmtran/219/219we07.htm.

[2] If people valued uncertainty by the expected outcome, this would imply that each extra euro they obtained would provide the same level of additional pleasure, whereas, in practice, the effect on happiness of gaining another euro may be different when the individual already has: €100, €1,000, €1,000,000, etc. An example of this is as follows: the expected outcome of the game is (2/6)€prize – €stake. If the prize is increased by a series of additional €1, the expected payoff becomes (2/6)(€prize + €1 + €1 + €1 …) and each extra €1 increases the expected payoff by exactly (2/6)€1, but the extra pleasure derived from each additional unit of prize money may not stay the same as the prize money increases, so the expected utility may be different.

[3] These are given in the context of someone deciding to enter into a risk. Where an individual has little or no choice in the matter, the risk-averse person would be defined as someone prepared to pay to avoid uncertainty, even if it has no effect on their expected income.

[4] Here, gamblers act rationally, but have risk-loving preferences. If gamblers also misperceive the probabilities involved in the game, or feel that they have control over random outcomes, they might also exhibit risk-taking behaviour due to errors in decision-making rather than just their underlying preferences.

[5] Bates, J., Polak, J., Jones, P. and Cook, A. (2001), 'The Valuation of Reliability for Personal Travel', *Transportation Research*, Part E, **37**, pp. 191–229.

[6] The recommendation in the current version of the *Passenger Demand Forecasting Handbook* (PDFH) is that each minute late is equivalent to an additional three minutes' journey time. Of the three additional minutes, 2.5 are due to the response to average delay, while the remainder is accounted for by displeasure with having to catch an earlier train than would otherwise be the case.

[7] Liu, X. and Polak, J. (2008), 'Alternative Specifications of Attitudes to Risk in Expected Random Utility Models of Risky Choice', Papers and Proceedings of the European Transport Conference, Leeuwenhorst.

[8] While noting the result of a recent study that 'briefcase travellers' consider themselves to be 98% as productive on board a train as in the office, which decreases the cost of unreliability, depending on the nature of the appointment being travelled to. See Fickling, R., Gunn, H., Kirby, H., Bradley, M. and Heywood, C. (2008), 'The Productive Use of Rail Travel Time and Value of Travel Time Saving for Travellers in the Course of Work', Papers and Proceedings of the European Transport Conference, Leeuwenhorst.

[9] Oxera (2006), 'Rail Market Speeds of Adjustment', June, prepared for the Passenger Demand Forecasting Council.

[10] This is complicated by the fact that the approach used for translating valuations from stated-preference studies into demand responses leads to implied elasticities that are highly specific to the circumstances being analysed.

If you have any questions regarding the issues raised in this article, please contact the editor, Derek Holt: tel +44 (0) 1865 253 000 or email d_holt@oxera.com

Other articles in the February issue of *Agenda* include:

- **capital budgeting at banks: the role of government guarantees**
  Viral Acharya and Julian Franks, London Business School
- **diversion ratios: why does it matter where customers go if a shop is closed?**
- **ducting the issue: what role might duct access play in an NGA environment?**

For details of how to subscribe to *Agenda*, please email agenda@oxera.com, or visit our website

www.oxera.com