Concluding Notes on Cost-Benefit Analysis – EC371, Fall 2011, by Jeremy Smith

(not required)

I. Further Techniques to Estimate Environmental Benefits

1) Defensive/Averting Expenditure and Complementary Demand

- again, since we cannot observe the price and quantity of the resource directly, we can alternatively try to look at a closely-related good for which we can observe the quantity people choose to buy and the prices they pay
- so to value a lake, we can observe people's purchasing behavior of fishing rods and swimsuits (i.e. things that are complementary to the use of the lake) and also estimate how the demand for these goods shifts with improvements to the lake, and then infer the average willingness to pay for lake improvements from the extra benefit that would be derived from these goods with such an improvement
- likewise, we can try to infer the value that the average person would put on an air-quality improvement by examining the average value people put on technology to simulate that improvement (or in other words how much people spend to "avert" or "defend against" the existing level of pollution)
    - (this is related to efforts to calculate welfare losses due to a lack of personal safety by looking at what people spend on locks, security devices and so on – which are sometimes referred to as a category of "regrettable expenditures")
- in general, these methods can be thought of as at best providing a lower bound on the actual willingness to pay of people for the resource in question, since averting behavior is probably limited by technological and practical considerations and there is often only a weak link between a given complementary good and actual use of the resource

2) Hedonics

- the name comes from an older econometric technique (which, if you are interested, is nowadays mostly used in the construction of price and real output series in the national accounts for goods with rapidly-changing quality – but the connection is not important here)
- see also the discussion of the application of this technique to valuing a life below
- the central idea is that some "goods" (especially houses) can be thought of as bundles of more basic characteristics (like floor area, number of rooms, proximity to schools and, of most interest from an environmental perspective, neighborhood air quality and proximity to environmental amenities)

- the method involves associating (econometrically) house prices with these characteristics, and stripping off, if you will, the part of house prices that corresponds with the environmental characteristics of interest
- then this stripped-off information can be used to "back out" an "average" person's demand curve for these environmental characteristics
- one of the first applications used house prices and characteristics in various neighborhoods in Boston along with detailed meteorological data, the latter to infer the typical concentration of industrial particulates in the air in the different neighborhoods over the course of a year
- the strict validity of the technique rests very heavily on the underlying theory and assumptions of perfect housing markets and consumers that are informed about all of the characteristics of houses and neighborhoods, and some other concerns to do both with the theory and econometric practice have been raised; but the technique is nonetheless thought of as being capable of providing good-quality estimates of WTP when applied carefully

II. Conjoint Analysis – An Extension of the Contingent Valuation Technique

- The technique replaces the typical willingness-to-pay question with a question asking the respondent to "vote for" one of a selection of different "states of the world". Each state includes different values for a few selected variables, including, for example, the amount of the resource provided, the average tax rate, the amount of public education expenditure and the level of household income. This might lead to more thoughtful responses than a direct WTP question. Though the required statistical tools are even more sophisticated, the amounts of the variables in the chosen state (compared to those in the states not chosen) can be leveraged to infer the willingness to pay for the resource by an "average" person that is revealed by the respondents' preferred states. The application of this kind of technique to environmental valuation is fairly recent, and it is not clear how well the validity or quality of the estimates has been examined. It appears that the degree of complexity in terms of both survey design and econometric analysis is extremely high. There is a simple discussion in the 8th edition (but not the 7th, which is the one in the library) of the Tietenberg textbook, and I can try to find more information if anyone is interested.

III. Techniques to Value a Statistical Life

1) Human Capital Approach
   - basically, examine earnings for a large sample of people at different stages in their careers, and econometrically estimate life-time total earnings for an "average" person
   - obviously misses a lot of intangible aspects of the value of a life
   - an early approach that is not really used anymore

2) Hedonics

- basically, we want to try to infer how much people value their *own* lives by trying to identify the risk premium they accept as part of their wage for working in a job associated with a certain risk of death
- associate (econometrically) wages with job characteristics including riskiness, then use this relationship to estimate an "average" person's willingness to pay for a given reduction in mortality risk
- this has become the most well-regarded and often-applied technique, in part because it is thought to identify what individuals' personal choices reveal about the implicit value they place on their own lives
- it may nonetheless suffer from some problems, including the fact that the underlying theory assumes that people have full information about the risks they face (research has shown that insisting that some strict aspects of the theory hold may lead to substantial under-estimates of the value of a statistical life: http://papers.nber.org/papers/w14364)
- can also be done by associating the price of cars purchased to characteristics of those cars, including the presence of safety features like side-impact beams and airbags

3) Contingent Valuation

- do a survey to collect a bunch of general information from respondents, and include a question that directly asks the respondent how he/she values his/her own life
- little is known about the quality of estimates of the value of a statistical life drawn from this method, but in the context of the recommendations of the NOAA panel on contingent valuation of environmental resources, it would appear that designing a survey that could be expected to produce bias-free estimates would be virtually impossible

IV. Concluding Points on the Value of a Statistical Life

- Most studies find the value of a statistical life to be between \$3 million and \$10 million, with the most confidence in the \$3-5 million range (these are from a few years ago, so inflation adjusted, they'd be a bit higher).

- Agencies typically have a standard value that they apply in cost-benefit analyses and that is updated every few years (here is an amusing article about the EPA's standard value of a statistical life: http://www.washingtonpost.com/wp-dyn/content/story/2008/07/19/ST2008071900185.html).

- Valuing improvements in health/freedom from injury may also be regarded as an important input to cost-benefit analyses, but it is much more difficult to accomplish.

- The other piece of information required to include the value of lives potentially lost due to implementing a given project – namely, the change in the probability of death associated with a given action – can also be difficult to calculate.
    - For the case of a project that would alter the amount of a chemical present in the air in a given location, such estimates are typically derived from laboratory tests on animals exposed to various amounts of the chemical – but converting the effect on a given species to the potential effect on humans is very sensitive to the scaling factors used (e.g. mass of a mouse compared to mass of an average human; surface area of a rat compared to surface area of an average human), and in general these estimates can be very uncertain.

    - In other situations it can be easier, such as guessing the increased number of deaths per year due to widening a freeway from observing the average number of deaths per year on wider freeways in similar parts of the country.

- There is an important distinction between the value of a "statistical" life (i.e. the value of an anonymous, "average" person's life), the value of a specific person's life and the value of life in general. Since the lives that will typically be saved/lost due to a given project cannot be identified in an *ex ante* sense (e.g. we may know how many highway fatalities to expect in a given year, but we have no way of identifying the precise people who will die before the accidents actually occur), the appropriate concept from the perspective of prospective project evaluation is the value of a statistical life (though, for the case of evaluating the desirability of policy action against climate change, Weitzman has suggested that a proper CBA may need to confront the problem of how to estimate the value of human existence in general, not to mention doing many other things much differently than traditional CBA).

- A particularly thorny issue is how to estimate the value of a statistical life amongst certain age groups (e.g. is an average five-year-old's life more valuable than an average seventy-five-year-old's life?), which might be important to consider if a given project will particularly benefit a certain age group (e.g. reducing particulate pollution might have an important impact on childhood asthma prevalence but little effect on adults). Also, though the techniques discussed above were never really meant to be used as such, some people are tempted to compare the estimated values of statistical lives across countries. We cannot pretend that philosophers, let alone economists, have been able to address such questions satisfactorily.

- Should we even bother trying to calculate the value of a life, given the potential philosophical objections and difficulties, and weaknesses of the estimation methods? Well, if we do not at least try to value lives in order to provide estimates for project evaluations, we must accept that our CBAs will always be drastically incomplete, because they will implicitly treat the value of a statistical life as zero (a definite under-estimate). And taking such an extremely incomplete CBA too seriously may not be much better than just scrapping project evaluation altogether and resorting to making policy on an ad-hoc basis. Further, trying to alternatively evaluate policies competing for limited funds by applying the notion that every life is priceless, besides being futile, ignores the fact that people behave everyday in ways that suggest they make trade-offs between accepting slightly higher risks of death and securing some tangible benefits. My personal opinion is that doing valuation of a statistical life as best as we can presently is much better than the alternative, which is basically to forsake CBA altogether and reduce the economic dimension of policy decisions to a coin toss.

V. Other CBA-Related Issues You May Come Across (not important for our purposes, but may be useful vocabulary) and Closing Thoughts

- If you use real (i.e. inflation-adjusted) dollars to measure costs and benefits, you should use a "real discount rate" (defined as the nominal rate minus the expected rate of inflation by what is known as the Fisher equation after Irving Fisher).
    - We have been ignoring inflation, and implicitly using nominal dollars with a nominal discount rate.

    - The two proper methods (real with real and nominal with nominal) lead to the same results in terms of which project is recommended.

- One should not compare NPVs of projects of different time frames directly.
    - One method to make a correct comparison is to "roll over" the shorter project until the repetition of the shorter project covers the same number of years as the longer project.

    - More commonly, an "equivalent annual net benefit" is calculated for each project.

    - (Sometimes a project with a terminal sale can implicitly be interpreted as an infinite-horizon project, and it's then ok to compare it to another infinite-horizon project. Don't worry, I will never ask you a question that tries to trick you on this matter because it's too complicated for us to have discussed.)

- Projects can also be evaluated on an internal-rate-of-return criterion or a benefit-cost-ratio criterion rather than the NPV criterion, but these have problems and should not be used in some circumstances.
    - The internal rate of return is the discount rate that would lead to an NPV of exactly zero. The criterion is to accept projects for which the IRR is greater than the discount rate actually chosen. The IRR can sometimes be difficult to calculate reliably.

   - The benefit-cost ratio is the ratio of the summation of the present values of all benefits to the summation of the present values of all costs. The criterion is to accept projects for which the BCR is greater than one. The ratio is sensitive to whether damages are included as negative benefits as opposed to positive costs, so its calculation can be subjective.

   - These are equivalent to the NPV criterion when comparing a single project against doing nothing. However, ranking different projects on the basis of higher IRRs or higher BCRs can lead to the recommendation of projects with lower NPVs, or in other words, projects associated with less of an efficiency improvement. It has been argued that this may be appropriate when the projects being considered can be scaled down in response to limited funding.

- There are alternative (and, some would argue, more comprehensive and rigorous) methods to deal with uncertainty besides using expected values.
   - Monte-Carlo analysis involves calculating NPV hundreds, thousands or even millions of times, each time taking a random "draw" for each uncertain variable from its postulated probability distribution, then taking the average of all of these NPVs.

   - Some researchers include "option values" as costs in CBAs. The idea is that, if we were to undertake an irreversible policy action now, it negates our ability to postpone the decision until a future date, when some uncertainty might be resolved and we might have a better way to forecast certain variables and thus produce a higher-quality NPV estimate (while, of course, some additional costs might be borne in the interim as well). The estimation of option values and the evaluation of when this is an appropriate approach are advanced topics, and ones in which there are still many open research questions.

- Cost-benefit analysis is a tool to assist in efficient decision making only.
   - As has been mentioned before, policy makers will typically care about many other things, such as equity, the health of the population, etc., and this can lead to policy at odds with the recommendations of cost-benefit analysts. This can be frustrating, but certainly does not negate the value of having a high-quality cost-benefit analysis as an input to the decision.
      - There is such a thing as "multi-attribute analysis", which takes several objectives and seeks to find the project that comes "closest" to meeting all of these objectives. If efficiency is one of the objectives, then CBA becomes an important input into this more comprehensive decision-making tool.

      - There is also such a thing as "distributionally-weighted" CBA, which attempts to compare projects on an efficiency and equity basis simultaneously by taking into account the varying degrees to which different groups will be affected by a given project.

- Cost-benefit analysis can sometimes involve much more art than science!
    - Given the numerous imperfections and approximations that are unavoidable with the practice of CBA, it can be thought of as a rather blunt tool. But by making careful choices (e.g. on the trade-off between taking on the expense of getting an imperfect estimate of a minor benefit and accepting a less comprehensive NPV estimate by leaving it out), it can nonetheless be a very useful tool.

    - Learning how to make these choices well and to develop NPV estimates that are defensible both to theoretical purists and to those who want a practically meaningful and useful evaluation at a reasonable price and within a reasonable length of time takes a lot of real-life practice. Partly for this reason, the coverage in this unit has at times been purposely vague and not gotten far beneath the surface. But I hope it has at least been a useful survey of many of the most important concepts and issues that will put you in good stead if you ever find yourself contributing to a cost-benefit analysis.

- There is a virtually unlimited demand for cost-benefit analyses, so those with some knowledge and experience are not at great risk of unemployment!
    - Since a decision by President Reagan during his presidency, all proposed federal projects exceeding a certain size are required to pass an “economic justification” test. This decision provided a large impetus for improving and more widely applying “regulatory impact analysis” or CBA techniques. The OMB, CBO and EPA do many CBAs annually, most in-house, while other agencies hire private consultancies. State governments, governments in other countries, multilateral institutions like the World Bank and many non-profits also devote substantial resources to performing CBAs.