Valuing Life and Health

Econ/Envi 213, Prof Jacobson

In any analysis that tries to quantify costs or benefits of environment-related actions, it is nearly inevitable that some of the important consequences we must value involve human life and health.

This is an ethically tricky area—most people are (to a lesser or greater extent) concerned by the idea of putting a dollar value on human life and health, or more generally with putting these "sacred" things on the scales in any way a cost benefit analysis. The counterargument is that when we do this it renders explicit a tradeoff we are inevitably making implicitly whenever we make a policy decision in these arenas, and that it is better to have transparency and use arguably-objective measures in analyzing these tradeoffs as opposed to trusting to whatever someone claims is the most ethical choice in the situation. I've thought about it a lot and I'm personally deeply ambivalent; it's up to you where you come down on it.

Regardless of how you feel about cost benefit analysis and this kind of valuation, you should understand the main methods for valuing risks to life and health because they are widely used. I give an overview of them in this document, with a focus on the environmental context.

Valuing Life

Let's work from an example. Air pollution increases population mortality, i.e., air pollution kills people. It does this in a statistical way: we can't say that any individual person's death from heart disease or COPD was directly caused by a higher ozone concentration in the relevant geographic area, but we do know that the rates of these afflictions in the population, and thus the death rate in the population, will be increased by the pollution. Therefore, we know that if we reduce air pollution we reduce mortality. In other words, air pollution regulation saves lives – albeit statistical lives, rather than identifiable individuals. This is, of course, also true for many other kinds of regulation, including those regarding seat belts, driving speed limits, water pollution, and controls on toxic substances.

If these things kill people, why do we allow them? First, in many of these cases, efforts to reduce mortality risks to zero would be incredibly costly and limiting. For example, risks from car accidents could be eliminated if we eliminated cars, but cars provide a lot of benefits to people that we're clearly not ready to give up. Second, often the reduction of one risk requires the increase of another risk: e.g., if we increase body scans at the airport, we reduce the risk of deaths from terrorism but increase the risk of deaths from radiation-induced cancer or car accidents, as people shy away from flying and drive instead (Blalock et al., 2009). Indeed, in some cases the money we spend to save lives in one place (e.g., by installing filters to reduce air pollution) is taken from a budget where its next best use would have been to save lives in another way (e.g., hiring inspectors to ensure tainted foods are removed from the food supply).

In other words, we are obviously willing to make some sacrifices to save some lives, but it is demonstrably untrue that people are willing to pay any cost to reduce population mortality (i.e., to save a (statistical) life).

In this spirit, we would like to find a way to put a dollar value on risks to life so we can have some idea what tradeoffs people are willing to make. As with other valuation tools, the idea is to derive society's value for the thing (in this case, mortality reductions) from our best estimate of individuals' value for it.

So in this case, the goal is not to come up with a dollar value for the life of an individual person (call him Fred); the goal is to come up with an amount people seem to be willing to pay to reduce their mortality risk and from that to infer what we as a whole seem to be willing to pay to reduce population mortality by one death. (By convention, this is one death over the course of a year.)

This concept is called the Value of a Statistical Life (VSL). Some economists have suggested that to be clearer about the fact that we're not talking about Fred's life but a change in population mortality risk, we should use a different word, like "micromort" (as argued in Cameron, 2010). There's a mathematical sense in which the two ideas are identical—if we're increasing deaths by one, why should it matter whether the person who dies is identifiable or not? However, while people seem to be willing to trade off population mortality risk against other things, we do seem to have a visceral ethical reaction to valuing Fred's life. And my guess is that, if put to the test, we probably would be willing to pay more to save an identifiable victim's life than to reduce population mortality by one person. (Economists generally consider this a behavioral bias.)

VSL is very important in policy. Many government agencies have a VSL estimate that drives their cost-benefit analyses or policy studies. For example, the US Environmental Protection Agency uses a value of $7.4 million in 2006 dollars (US Environmental Protection Agency, 2012), which is $9.38 million in 2018 dollars. The US Department of Transportation uses a value of $9.6 million in 2016 dollars (US Department of Transportation, 2016), which is $10.18 in 2018 dollars. When these agencies analyze policies that have been, or may be, enacted, they use these numbers to value changes in risks to lives. For example, the EPA's assessment of a revised air pollution rule (the Cross State Air Pollution Rule) found that the rule provides much larger benefits than costs, and this conclusion is largely driven by reductions in mortality risks (US Environmental Protection Agency, 2011).

So, again, we try to value life in this way using a tradeoff grounded in people's preferences, or our best estimate of preferences. There is some work using stated preference methods to estimate VSL—basically surveying people or conducting contingent valuation studies—but the estimates that people pay the most attention to are revealed preference methods because hypothetical bias may inflate or add noise to the stated preference measures. The revealed preference methods we use are basically hedonic methods that see how much people are willing to pay to reduce risks to their own lives or how much money people need to be paid to accept risks to their own lives.

The idea is that if we can find out the exact amount you are willing to pay (or accept) to avoid (or allow) an increase in risk of your death of Δp, then we can extrapolate to figure out how much you'd pay (or accept) to avoid (allow) certain death, assuming that the relationship between willingness-to-pay (WTP) or willingness-to-accept (WTA) and probability is linear:

$$VSL * \Delta p = WTP \quad \text{or} \quad VSL * \Delta p = WTA$$

Therefore if the WTP or WTA amount is $\$\Delta X$, then we can say that

$$VSL = \frac{\Delta X}{\Delta p}$$

Now, we generally don't see the exact amount that people are willing to pay to avoid a risk or that they are willing to accept to allow a risk. The best we can do is come up with boundaries for a range of a person's VSL. I'll demonstrate in the following paragraphs.

The willingness-to-pay methods often look at purchases of safety devices. The easiest example is fire safety. People generally buy smoke alarms. Say that fire alarms have an annualized cost of $20 including purchase and upkeep, and say also that they reduce risk of death by 1 in 100,000 per year. By buying the smoke alarms, people are demonstrating that they value their lives enough to make that tradeoff. So these people must value their lives *at least as much as*:

$$VSL \geq \frac{\$20}{1/100{,}000} \rightarrow VSL \geq \$2{,}000{,}000$$

Imagine, on the other hand, that people are generally unwilling to buy sprinkler systems. These people must not value their lives enough to make that tradeoff. Say that a sprinkler system has an annualized cost of $1,000 but it reduces the risk of death by 1 in 10,000 per year. Then they must value their lives by *less than*:

$$VSL \leq \frac{\$1{,}000}{1/10{,}000} \rightarrow VSL \leq \$10{,}000{,}000$$

Therefore we can say that for these people, the VSL lies between $2 million and $10 million.

I have seen estimates of willingness to pay measures of VSL based on safety features on cars, air quality (though people may value air quality for other reasons), and sunscreen, among other things.

The willingness to accept measures use the same idea, but here we're looking at people's choices to reduce their own safety for some money gain. The classic way to do this is to look at "compensating wage differentials," which you may have learned about in previous econ classes. In a general sense, the compensating wage differential is the increase in wage (as compared to an otherwise similar job) that is attributable to some unattractive feature of the job. You can think of hazard pay as an explicit manifestation of this, but it's often not that explicit, and indeed, compensating wage differentials can compensate for any unpleasant (or pleasant) feature of a job, including stress, fun, smell, etc. With regard to life risks, the compensating wage differential is the increase in wage that's attributable to the increase in mortality risks. For example, imagine workers can choose between job A and job B. If both jobs are identical in all other ways but job A (say, meat packing or lumberjacking) increases the risk of death while job B (say, office work) does not, then job A must have a higher wage or else no-one would be willing to take that job. The wage must be higher to compensate for the risk. So if we look at wages of different jobs and we control for all other aspects of the jobs, we should then be able to infer the population's VSL. Here we assume that the magic of labor market competition makes it so that the wage exactly compensates for the difference in risk, so that we can directly estimate VSL (not just bounds):

$$VSL = \frac{\Delta wage}{\Delta risk}$$

In a this vein, economists studying the willingness of Alaskan fishermen (like the people in *The Deadliest Catch*) to go out and fish, given the level of danger presented by the weather on a given day, estimated a VSL of $4-$4.67 million (Schnier et al., 2009).

Of course, there are a lot of problems with the practice of estimating VSL and with the ethics of the details. I'll just list a few here. First, is it right to extrapolate a life value from a willingness to trade off for very small changes in probabilities? This is particularly questionable when our estimates of all of the relevant components are poor—since the probabilities we're talking about are tiny, small differences in probabilities can yield large differences in VSL estimates. Also, the estimation is bedeviled by all sorts of problems—are people ignorant of the risks or overconfident? Are they extra-sensitive to some risks? Do the specific people who enter into the risky jobs do so because they have (or believe they have) a lower probability of death in those jobs than other people do? How comparable are the jobs we're comparing?

Next, whose willingness to trade off money for risk should we be looking at? It will generally be true that poorer people and minorities will demonstrate a lower VSL as compared to richer people and non-minorities. Poorer countries also have much lower VSL than rich countries; as discussed in Cropper (2009), recent estimates for China were about $220,000, which is something like 40 times less than the US VSL's I showed you above. It is obvious that a person's willingness to pay must be affected by their income, and willingness to accept probably will be as well. So it might seem empowering to allow people's own choices to guide the policies that their governments use, but on the other hand that will naturally lead to more sacrifice of life in some areas than others.

Also, should there be some relationship between the value of life and a person's age? Younger people have more years ahead of them; they should be willing to pay more to preserve those years as compared to the amount an older person is willing to pay. It seems that people's VSL as expressed at least through wage choices increase with age for a range and then drop again eventually (Aldy and Viscusi, 2007). However, to the extent that this causes us to apply anything like a "senior discount" to some mortality cost estimates, many people find this very troubling.

An alternative way to value life, which I mention by way of contrast, is the "human capital" approach. This is not used very much in economics, largely because it does not reflect our "consumer sovereignty" approach (the idea of building society's values out of people's preferences). However, it's similar to the way in which damages are calculated in legal cases. The idea of this approach is to say that if Fred dies, he no longer is working and producing what he was producing before. Therefore, the cost of his premature death is the value of his lost production (which should be equal to the income he would have earned had he not shuffled off this mortal coil). Note that according to this method, young people's lives would have a *much* higher value than old people… in fact, retired people's lives would have no value… and people who have disabilities that limit their ability to work for pay would also be without value… so this approach is more problematic than the VSL approach.

Valuing Health

Economists have done far less work studying how to value non-fatal health risks, probably because it's much harder. The proper term for these health damages is "morbidity," which generally refers to disease or other factors that make a person be in less than good health. Because it is so hard to value, morbidity is frequently ignored in analyses.

One difficulty is in defining degrees of morbidity. Obviously we care more about a health affliction that causes Fred to be bedridden for half of the remaining days of his life than about one that gives him the sniffles for a week. So we need to be able to estimate health effects both in terms of prevalence and severity (whereas, despite what they say in *The Princess Bride*, there's only one level of being dead).

One category of morbidity measurement methods, the "cost of illness" approach, estimates the costs of resources and opportunity costs associated with an illness. Note that this is very different from the idea of looking at people's preferences. One way to do this is to simply add up healthcare costs: doctors, equipment, medications, and that sort of thing. This may seem intuitive, but there's no reason to think that this corresponds to people's value (willingness to pay or willingness to accept) for the health condition.

A related approach is to look at the work-days lost and to multiply that by the person's wage or productivity for those days. This is a very instrumental conception of a person's value, just like the "human capital" idea I described for valuing life, and is problematic in the same ways. However, lost work-days is one of the measures I see used most often—probably because it's the easiest thing to calculate.

Notice that these methods do not consider the discomfort or unhappiness of being sick or injured, where presumably a method that derives value from a person's preferences would.

Another category of morbidity measurement methods tries to use individual choices to reveal preferences. In these methods, like in VSL estimation, estimates are drawn from people's own willingness to pay to avoid morbidity or willingness to accept to experience morbidity. We look at people's willingness to make expenditures or accept lower wages to reduce morbidity risks just as we do for mortality risks. However, these are often averting "expenditures" (like not going outside, not doing a particular task), and the costs of these are very hard to estimate because they are often not market purchases.

There are also all sorts of impediments to measuring these sorts of values. Remember that people don't pay for much of their own health care in a direct fashion; most people pay for insurance in advance and insurance pays health care costs when they arise. Employers also provide sick leave benefits so that people don't bear all of the costs of missing work. So in these ways, people may not directly bear all of the financial costs of ill health—some are shifted onto other members of the insurance pool or other employees of the same business. And of course, it's not clear how good people are at assessing the risks to their own health.

Another method that comes from health economics uses a measure of health called the quality-adjusted life year (QALY). When QALY is applied to a medical treatment or to a policy that will improve population health, it is a measure of how much improvement in health is expected and over what period (including an increase in years lived). Researchers perform surveys to learn about people's preferences for such an improvement, and then generally translate that into dollar amounts. QALY is not strongly grounded in economic theory, naturally requires the use of imprecise measures of morbidity, and uses stated rather than revealed preferences; but it does give us a way to get a handle on people's preferences for avoiding all elements of being sick or injured.

One thing to bear in mind is that none of our measures consider the fact that people reduce the likelihood and severity of illness and injury through various kinds of averting measures, from wearing sunscreen to wearing helmets to avoiding going outside on bad-air-quality days. If these are costly measures we should include their costs in our measures (i.e. we should add them in), because these avoidance costs are real costs that people are bearing.

References

Aldy, J.E., Viscusi, W.K., 2007. Age Differences in the Value of Statistical Life: Revealed Preference Evidence. Review of Environmental Economics and Policy, 1(2), 241-260.

Blalock, G., Kadiyali, V., Simon, D.H., 2009. Driving fatalities after 9/11: a hidden cost of terrorism AU - Blalock, Garrick. Applied Economics, 41(14), 1717-1729.

Cameron, T.A., 2010. Euthanizing the Value of a Statistical Life. Review of Environmental Economics and Policy, 4(2), 161-178.

Cropper, M., 2009. Measuring the Costs of Air Pollution and Health in China. Resources, Fall 2009.

Schnier, K.E., Horrace, W.C., Felthoven, R.G., 2009, The Value of Statistical Life: Pursuing the Deadliest Catch. Center for Policy Research, Maxwell School, Syracuse University, Center for Policy Research Working Papers: 117, pp. 52 pages.

US Department of Transportation, 2016, Guidance on Treatment of the Economic Value of a Statistical Life (VSL) in U.S. Department of Transportation Analyses - 2016 Adjustment.

US Environmental Protection Agency, 2011, Fact Sheet: The Cross-State Air Pollution Rule: Reducing the Interstate Transport of Fine Particulate Matter and Ozone.

US Environmental Protection Agency, 2012, Frequently Asked Questions on Mortality Risk Valuation, http://yosemite.epa.gov/ee/epa/eed.nsf/pages/MortalityRiskValuation.html.