PRICING THE PRICELESS:

Cost-Benefit Analysis of Environmental Protection

Lisa Heinzerling
Frank Ackerman

Georgetown Environmental Law and Policy Institute
Georgetown University Law Center

About the Authors

Lisa Heinzerling is a Professor of Law at Georgetown University Law School, specializing in environmental law. Before joining the Georgetown faculty, Professor Heinzerling practiced environmental law in the Massachusetts Attorney General's office and was a Skadden Fellow at a public interest group in Chicago. She clerked for Judge Richard A. Posner of the U.S. Court of Appeals for the Seventh Circuit and for Justice William J. Brennan, Jr., of the U.S. Supreme Court. Professor Heinzerling was editor-in-chief of the University of Chicago Law Review and has contributed to numerous law journals, including the Yale Law Journal, the Harvard Law Review, and the University of Chicago Law Review. She has been a visiting professor at the Yale and Harvard law schools.

Frank Ackerman is an environmental economist and research director of the Global Development and Environment Institute (G-DAE) at Tufts University. He has written widely on the economics of waste and energy, and on alternative approaches to economic theory. He is the author of *Why Do We Recycle? Markets, Values, and Public Policy* (Island Press 1997), and of two books on the economic policies of the Reagan administration. He is also the lead editor of three books in the G-DAE series on "Frontier Issues in Economic Thought," published by Island Press. Before coming to Tufts Dr. Ackerman worked at the Tellus Institute on energy and environmental research, and was a founding editor of *Dollars & Sense* magazine. He has taught economics at Tufts and at the University of Massachusetts, and received his PhD in economics from Harvard University.

The Georgetown Environmental Law and Policy Institute conducts research and education on legal and policy issues relating to protection of the environment and conservation of natural resources. The Institute publishes papers and reports by leading academics and practitioners applying rigorous analysis to contemporary environmental and natural resource issues. The views expressed in Institute publications are those of the authors and do not necessarily reflect the views of the staff of the Institute, its board of advisers, or Georgetown University.

3. What Are the Arguments in Favor of Cost-Benefit Analysis?

Before describing the problems with cost-benefit analysis, it will be useful to set forth the arguments in favor of this type of analysis. Many different arguments for cost-benefit analysis have been offered over the years. Most of the arguments fall into one of two broad categories. First, there are economic assertions that better results can be achieved with cost-benefit analysis. Second, there are legal and political claims that a more objective and more open government process can emerge through this kind of analysis.

Better Results

Economics frequently focuses on increasing efficiency – on getting the most desirable results from the fewest resources. How do we know that greater regulatory efficiency is needed? For many economists, this is an article of faith: greater efficiency is always a top priority, in regulation or elsewhere. Cost-benefit analysis supposedly furthers efficiency by ensuring that regulations are only adopted when benefits exceed costs and by helping direct regulators' attention to those problems for which regulatory intervention will yield the greatest net benefits.

But many advocates also raise a more specific argument, imbued with a greater sense of urgency. The government, it is said, often issues rules that are insanely expensive, out of all proportion to their benefits – a problem that could be solved by the use of cost-benefit analysis to screen proposed regulations. Thus much of the case *for* cost-benefit analysis depends on the case *against* current regulation.

One does not have to read very far into the literature on risk regulation before running across lengthy tables listing the costs per life saved of various federal regulations. The numbers on such tables are fantastic: according to these lists, we are often spending hundreds of millions, and sometimes billions, of dollars for every single human life, or even year of life, we save through regulation.[4]

These estimates of regulatory costs and benefits have become ubiquitous in political debates on environmental law. Scarcely a congressional hearing on this subject occurs in which these kinds of numbers do not figure prominently. Economists routinely cite these estimates as proof of the need for more economic analysis. Browse the web sites of any of a variety of think tanks, and you will find numerous references to the extravagant costs of regulation.

One widely cited study claims that the cost per year of life saved by life-saving interventions varies from zero or negative (some life-saving measures impose no new costs, and may even save money) up to *$99 billion*. The table on the following page is excerpted from that study. (Note, however, that not one of the pollution control measures listed in this table has ever been proposed by the government, much less implemented.)

COSTS PER LIFE-YEAR SAVED OF HYPOTHETICAL POLLUTION CONTROLS AT PAPER MILLS[6]

POLLUTION CONTROL MEASURE	COST PER LIFE-YEAR
Chloroform emission standard at 17 low cost pulp mills	Zero or Negative
Chloroform private well emission standard at 7 papergrade sulfite mills	$25,000
Chloroform private well emission standard at 7 pulp mills	$620,000
Chloroform reduction by replacing hypochlorite with chlorine dioxide at 1 mill	$990,000
Dioxin emission standard of 5 lbs/air dried ton at pulp mills	$4,500,000
Dioxin emission standard of 3 (vs. 5) lbs/air dried ton at paper mills	$7,500,000
Chloroform emission standard of 0.001 (vs. 0.01) risk level at pulp mills	$7,700,000
Chloroform reduction by replacing hypochlorite with chlorine dioxide at 70 mills	$8,700,000
Chloroform reduction at 70 (vs. 33 worst) pulp and paper mills	$15,000,000
Chloroform reduction at 33 worst pulp and paper mills	$57,000,000
Chloroform private well emission standard at 48 pulp mills	$99,000,000,000

Source: Tammy O. Tengs, et al., Five-Hundred Life-Saving Interventions and Their Cost-Effectiveness, 15 Risk Analysis 369(1995).

Numbers like these have been used to argue that current regulatory costs are not only chaotically variable but also unacceptably high. They have even been relied upon to claim that the existing regulatory system actually *kills people* by imposing some very costly life-saving requirements while other, less expensive and more effective life-saving possibilities remain untouched. Indeed, a study drawing upon these data concluded that we could save as many as 60,000 more lives every year with no increase in costs if we simply spent our money on the least rather than most expensive opportunities for saving lives. Relying on this research, John Graham, the current head of the Office of Information and Regulatory Affairs in the Office of Management and Budget and a prominent proponent of cost-benefit analysis, has called the existing state of affairs "statistical murder."[5]

From this perspective, cost-benefit analysis emerges as both a money-saver *and* a life-saver. By subjecting regulations to a cost-benefit test, we would not only stop spending hundreds of millions or billions of dollars to save a single life, we could also take that money and spend it on saving even more lives through different life-saving measures.

That, at least, is the theory. We will argue in the following sections that there are good reasons to question both the theory and the facts it rests on. Nevertheless, the notion that the current system produces crazy, even deadly, rules, and that better economic analysis would avert

this terrible result, remains one of the most persistent arguments offered on behalf of cost-benefit analysis.

Objectivity and Transparency

A second important set of arguments holds that cost-benefit analysis would produce a better regulatory process – more objective and more transparent, and thus more accountable to the public.

The holy grail of administrative law is agency decision making based on objective standards. The idea is to prevent an agency either from just doing anything it wants or, more invidiously, from benefiting politically favored groups through its decisions. Cost-benefit analysis has been offered as a means of constraining agency discretion to avoid these kinds of results.

Another important goal said to be promoted by cost-benefit analysis is transparency of administrative procedures. Decisions about environmental protection are notoriously complex. They reflect the input of biologists, toxicologists, epidemiologists, economists, engineers, lawyers, and other experts whose work is complicated and arcane. The technical details of these decisions often raise important questions about how much scientific uncertainty is too much, which human populations should be protected from illness and even death, and how important the future is relative to the present.

In order for the public to be part of the process of decision making about the environment, these judgments must be offered and debated in language accessible to people who are not biologists, toxicologists, or other kinds of experts. Many advocates of cost-benefit analysis believe that their methodology provides such a language. They also assert that cost-benefit analysis renders decision-making transparent insofar as it requires decision-makers to reveal all of the assumptions and uncertainties reflected in their decisions.

4. Why It Doesn't Work: Fundamental Flaws

As we have seen, cost-benefit analysis involves the creation of artificial markets for things - like good health, long life, and clean air - that are not bought and sold. It also involves the devaluation of future events through discounting.

So described, the mind-set of the cost-benefit analyst is likely to seem quite foreign. The translation of all good things into dollars and the devaluation of the future are inconsistent with the way many people view the world. Most of us believe that money doesn't buy happiness. Most religions tell us that every human life is sacred; it is obviously illegal, as well as immoral, to buy and sell human lives. Most parents tell their children to eat their vegetables and do their homework, even though the rewards of these onerous activities lie far in the future. Monetizing human lives and discounting future benefits seem at odds with these common perspectives.

The cost-benefit approach also is inconsistent with the way many of us make daily decisions. Imagine performing a new cost-benefit analysis to decide whether to get up and go to work every morning, whether to exercise or eat right on any given day, whether to wash the dishes or leave them in the sink, and so on. Inaction would win far too often – and an absurd amount of effort would be spent on analysis. Most people have long-run goals, commitments, and habits that make such daily balancing exercises either redundant or counterproductive. The same might be true of society as a whole undertaking individual steps in the pursuit of any goal, set for the long haul, that cannot be reached overnight - including, for example, the achievement of a clean environment.

Moving beyond these intuitive responses, we offer in this section a detailed explanation of why cost-benefit analysis of environmental protection fails to live up to the hopes and claims of its advocates. There is no quick fix, because these failures are intrinsic to the methodology, appearing whenever it is applied to any complex environmental problem. In our view, cost-benefit analysis suffers from four fundamental flaws, addressed in each of the next four subsections:

- *The standard economic approaches to valuation are inaccurate and implausible.*
- *The use of discounting improperly trivializes future harms and the irreversibility of some environmental problems.*
- *The reliance on aggregate, monetized benefits excludes questions of fairness and morality.*
- *The value-laden and complex cost-benefit process is neither objective nor transparent.*

Dollars Without Sense

Recall that cost-benefit analysis requires the creation of artificial prices for all relevant health and environmental impacts. To weigh the benefits of regulation against the costs, we need to know the monetary value of preventing

the extinction of species, preserving many different ecosystems, avoiding all manner of serious health impacts, and even saving human lives. Without such numbers, cost-benefit analysis cannot be conducted.

Artificial prices have been estimated for many, though by no means all, benefits of regulation. As discussed, preventing the extinction of bald eagles reportedly goes for somewhat more than $250 per household. Preventing retardation due to childhood lead poisoning comes in at about $9,000 per lost IQ point (although, as we will see in Section 5, a much lower price has recently been proposed). Saving a life is ostensibly worth $6.3 million.

This quantitative precision, achieved through a variety of indirect techniques for valuation, comes at the expense of accuracy and even common sense. Though problems arise in many areas of valuation, we will focus primarily on the efforts to attach a monetary value to human life, both because of its importance in cost-benefit analysis and because of its glaring contradictions.

There Are No "Statistical" People

What can it mean to say that saving one life is worth $6.3 million? Human life is the ultimate example of a value that is not a commodity, and does not have a price. You cannot buy the right to kill someone for $6.3 million, nor for any other price. Most systems of ethical and religious belief maintain that every life is sacred. If analysts calculated the value of life itself by asking people what it is worth to them (the most common method of valuation of other environmental benefits), the answer would be infinite, as "no finite amount of money could compensate a person for the loss of his life, simply because money is no good to him when he is dead."[6]

The standard response is that a value like $6.3 million is not actually a price on an individual's life or death. Rather, it is a way of expressing the value of small risks of death; for example, it is one million times the value of a one in a million risk. If people are willing to pay $6.30 to avoid a one in a million increase in the risk of death, then the "value of a statistical life" is $6.3 million.

Unfortunately, this explanation fails to resolve the dilemma. It is true that risk (or "statistical life") and life itself are distinct concepts. But if human life is too sacred to buy and sell, why is it permissible to trade small risks of losing that ultimate value? One-millionth of an immeasurable or infinite value is still immeasurable or infinite, not $6.30.[7]

Human life is the ultimate example of a value that is not a commodity, and does not have a price.

In practice, moreover, analysts often ignore the distinction between valuing risk and valuing life.[8] Many regulations reduce risk for a large number of people, and avoid actual death for a much smaller number. A complete cost-benefit analysis should, therefore, include valuation of both of these benefits. However, the standard practice is to calculate a value only for "statistical" life and to ignore life itself.

The confusion between the valuation of risk and the valuation of life itself is embedded in current regulatory practice in another way as well. The Office of Management and Budget – which reviews cost-benefit analyses prepared by federal

agencies pursuant to executive order – instructs agencies to discount the benefits of life-saving regulations from the moment of avoided death, rather than from the time when the *risk* of death is reduced.[9]

Most spiritual beliefs call on us to value the lives of others - not only those closest to us, but also those whom we have never met.

This approach to discounting is plainly inconsistent with the claim that cost-benefit analysis seeks to evaluate risk. When a life-threatening disease – such as cancer – has a long latency period, many years may pass between the time when a risk is imposed and the time of death. If monetary valuations of statistical life represented risk, and not life, then the value of statistical life would be discounted from the date of a change in risk (typically, when a new regulation is enforced) rather than from the much later date of avoided actual death.[10]

In acknowledging the monetary value of reducing risk, economic analysts have contributed to our growing awareness that life-threatening risk itself – and not just the end result of such risk, death – is an injury. But they have blurred the line between risks and actual deaths, by calculating the value of reduced risk while pretending that they have produced a valuation of life itself. The paradox of monetizing the infinite or immeasurable value of human life has not been resolved; it has only been glossed over.

People Care About Other People

Another large problem with this approach to valuation of life is that it asks individuals (either directly through surveys, or indirectly through observing wage and job choices) only about their attitudes toward risks to themselves.

A recurring theme in literature suggests that our deepest and noblest sentiments involve valuing someone else's life more highly than our own: think of parents' devotion to their children, soldiers' commitment to those whom they are protecting, lovers' concern for each other. Most spiritual beliefs call on us to value the lives of others - not only those closest to us, but also those whom we have never met.

This point echoes a procedure that has become familiar in other areas of environmental valuation. Economists often ask about existence values: how much is the existence of a wilderness area or an endangered species worth to you, even if you will never personally experience it? If this question makes sense for bald eagles and national parks, it must be at least as important when applied to safe drinking water and working conditions for people we don't know.

The difficulty is that the answer to this type of question cannot be deduced solely from your attitudes toward risks to yourself. We are not aware of any attempts to quantify the existence value of the life of a stranger, let alone a relative or a friend, but we are sure that most belief systems affirm that this value is substantial (assuming, of course, that the value of life is a number in the first place).

Voting Is Different From Buying

Cost-benefit analysis, which relies on estimates of individuals' preferences as consumers, also fails to address the collective choice presented to society by most public health and environmental problems.

Under the cost-benefit approach, valuation of environmental benefits is based on individuals' private decisions as consumers or workers, not on their public values as citizens. However, policies that protect the environment are often public goods, and are not available for purchase in individual portions. In a classic example of this distinction, the philosopher Mark Sagoff found that his students, in their role as citizens, opposed commercial ski development in a nearby wilderness area, but, in their role as consumers, would plan to go skiing there if the development was built.[11] There is no contradiction between these two views: as individual consumers, the students would have no way to express their collective preference for wilderness preservation. Their individual willingness to pay for skiing would send a misleading signal about their views as citizens.

It is often impossible to arrive at a meaningful social valuation by adding up the willingness to pay expressed by individuals. What could it mean to ask how much you personally are willing to pay to clean up a major oil spill? If no one else contributes, the clean-up won't happen regardless of your decision. As the Nobel Prize-winning economist Amartya Sen has pointed out, if your willingness to pay for a large-scale public initiative is independent of what others are paying, then you probably have not understood the nature of the problem.[12] Instead, a *collective* decision about collective resources is required.

In a similar vein, the philosopher Henry Richardson argues that reliance on the cost-benefit standard forecloses the process of democratic deliberation that is necessary for intelligent decision-making. In his view, attempts to make decisions based on monetary valuation of benefits freeze preferences in advance, leaving no room for the changes in response to new information, rethinking of the issues, and negotiated compromises that lie at the heart of the deliberative process.[13]

Cost-benefit analysis turns public citizens into selfish consumers, and interconnected communities into atomized individuals. In this way, it distorts the question it sets out to answer: how much do we, *as a society*, value health and the environment?

Numbers Don't Tell Us Everything

A few simple examples illustrate another problem – that numerically equal risks are not always equally deserving of regulatory response. The death rate is roughly the same (somewhat less than one in a million) from a day of downhill skiing, from a day of working in the construction industry, or from drinking about 20 liters of water containing 50 parts per billion of arsenic, the old regulatory limit that was recently revised by the Bush administration. This does not mean that society's responsibility to reduce risks is the same in each case.

Most people view risks imposed by others, without an individual's consent, as more worthy of government intervention than risks that an individual knowingly accepts. On that basis, the highest priority among our three examples is to reduce drinking water contamination, a hazard to which

no one has consented. The acceptance of a risky occupation such as construction is at best quasi-voluntary–it involves somewhat more individual discretion than the "choice" of public drinking water supplies, but many people go to work under great economic pressure, and with little information about occupational hazards. In contrast, the choice of risky recreational pursuits such as skiing is entirely discretionary; obviously no one is forced to ski. Safety regulation in construction work is thus more urgent than regulation of skiing, despite the equality of numerical risk.

In short, even for ultimate values such as life and death, the social context is decisive in our evaluation of risks. Cost-benefit analysis assumes the existence of generic, acontextual risk, and thereby ignores the contextual information that determines how many of us, in practice, think about real risks to real people.

Artificial Prices Are Expensive

Finally, the economic valuation called for by cost-benefit analysis is fundamentally flawed because it demands an enormous volume of consistently updated information, which is beyond the practical capacity of our society to generate.

All attempts at valuation of the environment begin with a problem: the goal is to assign monetary prices to things that have no prices, because they are not for sale. One of the great strengths of the market is that it provides so much information about real prices. For any commodity that is actually bought and sold, prices are communicated automatically, almost costlessly, and with constant updates as needed. To create artificial prices for environmental values, economists have to find some way to mimic the operation of the market. Unfortunately the process is far from automatic, it is certainly not costless, and it has to be repeated every time an updated price is needed.

As a result, there is constant pressure to use outdated or inappropriate valuations. Indeed, there are sound economic reasons for doing so: no one can afford constant updates, and significant savings can be achieved by using valuations created for other cases. In the EPA's original cost-benefit analysis of arsenic (see the arsenic case study, starting at page 17), the estimated value of a case of chronic bronchitis was used to represent the value of a case of nonfatal bladder cancer.

This is not, we hope and believe, because anyone thinks that bronchitis and bladder cancer are the same disease. The reason is more mundane: no one has performed an analysis of the cost of bladder cancer, and even the extensive analysis of arsenic regulations did not include enough time and money to do so. Therefore, the investigators used an estimated value for a very different disease. The only explanation offered for this procedure was that it had been done before, and nothing better was available.

Use of the bronchitis valuation to represent bladder cancer can charitably be described as grasping at straws. Lacking the time and money to fill in the blank carefully, the economists simply picked a number. This is not remotely close to the level of rigor that is seen throughout the natural science, engineering, and public health portions of the arsenic analysis. Yet it will happen again, for exactly the same reason. It is not a failure of will or intellect, but rather the inescapable limitations of

time and budget, that lead to reliance on dated, inappropriate, and incomplete information to fill in the gaps on the benefit side of a cost-benefit analysis.

Summing Up

There are, in short, a host of problems with the process of valuation. On a philosophical level, human life may belong in the category of things that are too valuable to buy and sell. Most ethical and religious beliefs place the protection of human life in the same category as love, family, religion, democracy, and other ultimate values, which are not and cannot be priced.

Absent a credible monetary metric for calculating the benefits of regulation, cost-benefit analysis is inherently unreliable.

It is a biased and misleading premise to assume that individuals' willingness to pay to avoid certain risks can be aggregated to arrive at a figure for what society should pay to protect human life. Risk of death is not the same as death itself, and not all risks can reasonably be compared one to the other. Moreover, the value to society of protecting human life cannot be arrived at simply by toting up individual consumer preferences.

The same kind of problems affect other valuation issues raised by cost-benefit analysis, such as estimating the value of clean water, biodiversity, or entire ecosystems. The upshot is that cost-benefit analysis is fundamentally incapable of delivering on its promise of more economically efficient decisions about protecting human life, health, and the environment. Absent a credible monetary metric for calculating the benefits of regulation, cost-benefit analysis is inherently unreliable.

(main text continued on page 21)