Fallacy 3: Older People Are Less Valuable

Regulations save lives. When decisionmakers ask whether a particular regulation is justified, they are often swayed by the number of lives it saves. How they measure the benefit of a saved life, then, is critical. The standard way to measure the benefit of a saved life is to look at people's willingness to pay to avoid health and safety risks, which is then used to estimate the value of a statistical life. Recently there have been efforts to substitute the value of *life-years* for the value of statistical life. Use of life-years, however, leads to the third fallacy in cost-benefit analysis—that older people are less valuable than younger people, in proportion to their life expectancies. Under the life-years methodology, because younger people will on average lose more life years when they die, their lives are assigned a much greater value—the life of a 40-year-old is seven times more valuable than that of a 70-year-old. This outcome is both inconsistent with economic theory and flatly contradicted by empirical data on how people value risk.

The origin of the life-years approach is typically traced to an article by Michael J. Moore and W. Kip Viscusi entitled *The Quantity-Adjusted Value of Life*.[231] In that study, Moore and Viscusi introduced the idea that the value of a statistical life "cannot be divorced from the duration of life involved since lives are extended, not permanently saved."[232] They argued that "in the case of fatalities, a young person loses a much greater amount of lifetime utility than does an older person."[233]

Compared with the average subject of a wage-risk study, who is roughly 40 years old with a 35-year life expectancy,[234] a beneficiary of a regulation might be 70 years old and in poor health, with a life expectancy of around five years. Proponents of the life-years method argue that it is inappropriate to apply a value of statistical life estimate derived from studies of 40-year-old workers to elderly individuals who have significantly fewer remaining life years.

Proponents of the life-years method use a constant per life-year value, so that all life-years are valued equally no matter when they occur during the life cycle. They derive that number by taking the value of a statistical life and dividing it by the average life expectancy of the subjects of wage-risk studies. The application of a constant per life-year value across the life cycle results in smaller values being assigned to the elderly or the unhealthy. For example, using a $180,000 life-year value, the life of an elderly (or unhealthy) individual with a remaining life expectancy of five years would be valued at $900,000. Therefore, under the life-years approach, the value assigned to the life of the elderly individual would be only one-seventh the $6.3 million valued assigned to an individual of average age.

WHY IT'S BAD: THE CASE OF SMOG

To illustrate the contrast between the statistical life and the life-years methods, consider it in the context of ground-level ozone regulation. Ground-level ozone, commonly known as smog, leads to many terrible respiratory conditions, including asthma, permanent lung damage, and death. Besides these health effects, it also has negative environmental effects, including harm to plant life and decreased visibility. For these reasons, ozone levels have been regulated since 1970.

Under the Clean Air Act, the Environmental Protection Agency (EPA) sets permissible ozone levels. The agency is required to review these levels on a periodic basis to keep them current with developing science. At this writing, the EPA is conducting such a review. Several studies since the last review in 1996 have shown that the harm of short-term exposure to ozone has been underestimated. New data on the effects of ozone exposure on, in the jargon of the EPA, "additional respiratory-related endpoints, newly identified cardiovascular-related health endpoints, and mortality," have caused the staff of EPA to recommend consideration of a stricter standard that would require a lower level of ozone.[235]

Regulating ozone has not been easy. Ground-level ozone is generally created through a chemical reaction in which volatile organic compounds and nitrogen oxides mix in the presence of sunlight. These compounds have myriad sources, including motor vehicle exhaust and gasoline vapors. Controlling the pollution associated with automobiles is famously difficult. Although pollution per mile driven has fallen in the last thirty years, the number of miles driven has risen,[236] leaving ozone levels dangerously high in many parts of the country.[237] Engineering solutions to these problems have been scarce and costly.[238]

A more stringent regulation from the EPA, then, would be an enormous headache for the affected industry. But how can industry complain about a headache that saves lives? One answer is that the benefit of the headache is less than it might seem to be, because "saving lives" should be understood not as saving individual lives but as saving life-years.

Use of this method could scuttle or weaken new ozone standards. Individuals with existing respiratory problems will likely be the chief beneficiaries of a stricter regulation. As a population, these individuals will have a lower life expectancy than the average American worker, either because of advanced age or ill health. Adopting a life-years methodology would prevent some of the most vulnerable people in our society from being counted the same as everyone else, potentially blocking regulations that can save their lives.

If short-term exposure to elevated ozone levels results in the deaths of individuals with an average life expectancy of five years, then the benefit of the regulation, calculated by the life-years method, will be one-seventh that of the benefit calculated with the value of a statistical life estimate: $900,000 per life. A half-billion-dollar benefit is transformed into a $70 million benefit with a wave of the life-years wand. Compared to the statistical life method, the estimated regulatory benefit is reduced by as much as 85 percent. Thus, life-years, used in this and similar contexts, can have extremely important consequences—the difference between smog-induced deaths and cleaner air.

LIFE-YEARS ANALYSIS IN THE REGULATORY STATE

Life-years analysis has become an important part of several recent cost-benefit analyses conducted under the George W. Bush administration. John Graham was its main advocate, and with coauthor Tammy Tengs, had previously championed the life-years method in an influential criticism of federal regulations.[239] As Office of Information and Regulatory Affairs (OIRA) director, Graham prodded the EPA to use "alternative benefit calculations" placing the life-years calculations alongside calculations based on the value of statistical lives. As a result, the EPA incorporated this technique in several important cost-benefit analyses.[240]

The use of life-years analysis became the subject of significant public debate when critics in the media and in the environmental and senior citizen communities charged that the valuation techniques used by the EPA for the Clear Skies Act cost-benefit analysis amounted to a "senior death discount."[241] Highly unfavorable stories ran in the *New York Times*, the *Wall Street Journal*, and the *Washington Post*. The criticism eventually forced the Bush administration to drop some of the more controversial techniques. Even during the political hoopla over the senior death discount, however, John Graham only somewhat revised his views on the life-years methodology, conceding that it needed to be improved, but disagreeing with the idea that it needed to be scrapped altogether.[242]

Despite the outcry, the life-years method has not been abandoned by the Bush administration. It remains an important and influential way of thinking about cost-benefit analysis both at the Office of Management and Budget (OMB) and within administrative agencies. As recently as September 2006, members of the EPA Science Advisory Board's Committee on Environmental Economics discussed the utility of life-years analysis for EPA cost-benefit analysis, failing to come to a consensus conclusion.[243]

WHY IT'S WRONG

The life-years method does not flow from either sound economic theory or good facts. The approach is fundamentally inconsistent with the important tenet of economic theory in which value is determined by the willingness to pay.[244] Under that tenet, the economic value of mortality risk reductions should be determined by how much an individual would voluntarily exchange for the reduction. It would only be economically defensible to decrease the value assigned to mortality risk reduction to account for age if one's willingness to pay decreases as one ages.

But the life-years method ignores willingness to pay as a proxy for value, and instead assumes a downward linear relationship between a person's age and the value of that person's life. This assumption is inconsistent with the standard economic observation that individuals generally assign greater value to goods that are more limited in supply. The technique uses a constant per life-year value, so that all life years are valued equally no matter when they occur during the life cycle.

As people age, they can anticipate fewer future life years. Because of this scarcity, we might expect that they would value their future life years more highly than younger people would. By assuming that no difference exists between the values a 40-year-old and a 70-year-old would attribute to an additional year of life, the life-years method overlooks the effect of scarcity on valuation.[245] By ignoring the effect of scarcity and focusing regulatory efforts on reducing risk for young and healthy people, the life-years method delivers regulatory benefit to those who value it least. This approach takes the standard economic logic of "willingness to pay" and stands it on its head. Generally, the most efficient system is the one that moves resources to the people that value them most. The life-years method accomplishes exactly the opposite. Moreover, across a certain age range of their lives, as people grow older, they have more income and wealth. It is well established that the willingness to pay to avoid risk is highly correlated with income. The greater affluence of middle-aged individuals (at least preretirement) thus suggests an increase in the willingness to pay, contrary to the prediction of the life-years method.

Various models—all ignored by life-years advocates—seek to determine how the value of risk reduction might change with age. No clear answer has emerged.[246] Some models predict that as the probability of death increases, so does the willingness to pay to avoid risk,[247] because people cannot take money to their graves. In other models, increases in background risk, which occur as people age, decrease the willingness to pay for a specified risk.[248] Other models are simply ambiguous.[249] It is possible that none of these models captures the whole story. What is important, however, is that no plausible economic model offers even lukewarm support for the diminishing linear relationship between life expectancy and willingness to pay that undergirds the life-years method. The life-years method, then, is entirely without theoretical justification.

Second, this method is devoid of empirical support and inconsistent with existing empirical work. Relevant studies have found that the willingness to pay does not resemble the constant age-dependant discount postulated by proponents of the life-years method. The results of these studies have varied somewhat. Some studies have found that the willingness to pay for risk reduction is independent of age. A 2002 stated preference study by Anna Alberini, Maureen Cropper, Alan Krupnick, and Nathalie Simon[250] found no statistically significant decline, with age, in the willingness to pay for risk reduction in the U.S. population.[251] Furthermore, the study found that the willingness to pay did not decline among people with lower life expectancy attributable to diagnoses of cancer, or heart or lung disease. On the contrary, people with these conditions tended to be willing to pay more for risk reduction. Other studies have found that the willingness to pay increases with age. A revealed preferences study done by V. Kerry Smith, Mary Evans, Hyun Kim, and Donald Taylor Jr., using data from the Health and Retirement Study found that the oldest and most risk-averse individuals required "significantly higher, not lower" compensation in order to take on greater risk.

Still other studies have found an inverted-U-shaped relationship between the willingness to pay and age, in which the willingness to pay increases early in life, levels off in the middle, and then drops off near the end.[252] One of these studies, by Thomas Kniesner, W. Kip Viscusi, and James Ziliak[253] found that the inverted-U-shape was lopsided, with willingness to pay shooting up during the younger years, peaking at around age fifty, and then declining at a much slower rate.

In a large review of value-of-statistical-life estimates from over sixty studies in ten countries, Viscusi and Joseph Aldy looked at how age tended to impact the willingness to pay. They found ambiguous results, with some studies showing no significant relationship between age and the willingness to pay, while others suggested that age tended to decrease the willingness to pay. The studies that showed the willingness to pay decreasing with age, however, are suspect because of the deeply unintuitive finding that people are risk-loving, meaning that they were willing to pay to take on risk rather than demanding a premium.[254]

Although the empirical work on the relationship between age and willingness to pay has yet to provide a clear explanation about how age affects the willingness to pay to reduce risk, it clearly disproves the life-years hypothesis. All studies show that the life-years method is empirically unjustified and will lead to the systematic underestimation of the regulatory benefits of important programs. Even the estimates showing a relationship between age and willingness to pay do not support as steep and constant a decline as is assumed by the life-years method. All of the empirical evidence shows that this method produces extremely significant underestimates of the value of risk reduction to elderly individuals and individuals with low life expectancies due to poor health.

Use of the life-years technique, then, divorces cost-benefit analysis from its economic foundations. Because cost-benefit analysis claims to identify economically efficient regulation, its methodologies must be economically coherent—which the life-years technique is not. Use of this technique, or some other alternative valuation that is not grounded in economic theory and reality, amounts to smuggling noneconomic conclusions about how people "should" value risk into cost-benefit analysis. This perverts the meaning of the technique and ultimately renders its conclusions confused and useless.

WILLINGNESS TO PAY AND THE USE OF AVERAGE VALUES

The correlation among willingness to pay, age, wealth, and other individual characteristics raises important questions about how to calculate the benefits of life-saving regulations. Even if empirical work revealed a rock-solid relationship among certain individual characteristics, including age and willingness to pay, the question of whether regulators *should* take those differences into account, rather than using average values, would remain open.

For example, as just discussed, income can be expected to correlate quite closely with a willingness to pay for risk reduction.[255] Ethnicity and race may correlate as well.[256] It is not clear, however, that dividing the beneficiaries of regulations through categorization according to income or race would be ethically defensible—or, for that matter, legal. The same holds true for age; it is not clearly morally acceptable to target a specific subpopulation and reduce its estimated value of risk reductions without doing the same for other demographic factors.

The standard willingness-to-pay methodology uses average values for risk reductions, and does not break down these values according to demographic subpopulations. This practice, however, is also open to criticism. Because it is unwilling to allow the valuation of a statistical life to reflect people's actual preferences, a cost-benefit analysis using an average value of a statistical life is unmoored from its economic justification.

In the aggregate, however, using the average value of a statistical life can produce reasonably efficient levels of regulation. More finely tuned valuations might result in less error if the population affected by a regulation had markedly different valuations. Nevertheless, those errors will occur on "both sides"—*i.e.*, over- and under-estimations—so that, overall, regulation should be neither too strict nor too weak. Although some inefficiently weak or strong regulations will pass a cost-benefit analysis test, using the average value of a statistical life does not create a systematic bias, either in favor of or against regulation.

If the average is used for the value of a statistical life, people with high willingness to pay receive "too little" regulation, people with low willingness to pay receive "too much" regulation, and the average person gets a "just right" level of regulation. In many cases, it will not be possible to finely calibrate regulation, so this phenomenon is simply a byproduct of the inability to individually tailor regulations.

Even if regulations do affect subpopulations rather than the population as a whole, the use of an average value allows us to avoid making troubling race-gender-income-age-based categorizations. Furthermore, the use of average values tends to result in a form of regulatory wealth transfer whereby those with less wealth (who, therefore, have less willingness to pay) get more regulatory benefit than they might bargain for, and those with more wealth get somewhat less than they would prefer. So long as the costs of the regulation are not borne by the direct beneficiaries—as they almost never are—the result is a progressive distribution of social goods—in this case regulatory benefits—that is not normatively troubling. Although redistribution through regulation is generally not the most efficient means of achieving egalitarian results,[257] it can be justified if reliance on average values avoids morally troubling categorizations. Thus, even if there were support for the relationship between willingness to pay and age contained in the life-years approach, which there is not, it might still not be appropriate to use life-years valuations for regulatory purposes because these valuations single out the elderly over other groups for negative treatment.

CONCLUSION

The life-years method, which leads to the fallacy that older people are less valuable than younger people, should be abandoned. The current state of the science is the value of a statistical life, based on revealed-preference studies. The effects on age of the willingness to pay are interesting, and deserve future study. Even if age can be shown to have consistent effects on the willingness to pay to avoid risk, however, many thorny issues remain. Those issues include the fundamental unfairness of treating lives differently, or how to treat children (who, because they have no assets, have little willingness to pay). We currently avoid these problems by using average values. Before we abandon that strategy, proponents of new methodologies have the burden to show that their new techniques address these issues at least as well as the current method.

There is another important lesson to be learned from life years. With their campaign against the "senior death discount," environmentalists actively confronted a methodological issue within cost-benefit analysis with very significant success. They forced the George W. Bush administration to revise its position, illustrating that if progressive groups get involved in the conversation over how cost-benefit analysis is conducted, they can prevail, especially if they stand on the moral and rational high ground, as they did on this issue. The debate over life years, while far from over, shows that proregulatory interests can win important victories by fighting these fallacies in the trenches.