Rethinking the Value of Life: A Critical Appraisal of the Value of a Statistical Life

The Center for
Growth and **Opportunity**
at Utah State University

Rethinking the Value of Life: A Critical Appraisal of the Value of a Statistical Life

Author:
James Broughel[a]

January 2020
Policy Paper 2020.002

The Center for Growth and Opportunity at Utah State University is a university-based academic research center that explores the scientific foundations of the interaction between individuals, business, and government.

We support research that explores a variety of topics from diverse perspectives. Policy papers are published to stimulate timely discussion on topics of central importance in economic policy and provide more accessible analysis of public policy issues.

The views expressed in this paper are those of the author(s) and do not necessarily reflect the views of the Center for Growth and Opportunity at Utah State University or the views of Utah State University.

[a] Senior Research Fellow, Mercatus Center at George Mason University

* The author would like to thank Hayden Warlick and Blake Hoarty for helpful research assistance that contributed to this report.

The Center for Growth and Opportunity at Utah State University

Contents

Introduction

When economists perform a cost-benefit analysis (CBA) for regulations aimed at addressing life-threatening risks, they will often assign a dollar value to each life they expect to save. This allows them to compare the benefits of risk reduction to the monetary costs associated with complying with regulations, in order to assess whether on balance regulations should proceed. The most common way to assign a dollar value to a life in CBA is by looking at what currently living individuals are willing to accept in terms of monetary payment for accepting increases in mortality risk. Armed with this information, an analyst can deduce an estimate of how much people value their own lives in dollar terms. This "implicit value of life" measure, which is known as the "value of a statistical life," or VSL, is used extensively at US federal regulatory agencies.[1]

Despite fairly widespread acceptance of the VSL among economists, there are a number of complications associated with the metric that, although sometimes acknowledged in the academic literature, remain unresolved. This article reviews these complications and argues that it is time to revisit them in a more forthright manner. Taken together, estimates of life's value besides the VSL should be considered when estimating the value of mortality risk reductions in CBA. While the focus of this article is on the VSL, many of the problems identified here extend to the more general practice of "monetizing"—or assigning dollar values to—goods not traded on any market using consumer willingness to pay or accept measures.

To be clear, the problems with the VSL are not empirical; they are theoretical. It is not uncertainty with respect to how the VSL is measured that is the concern of this article, but rather whether the VSL itself is a theoretically sound way to value human lives in CBA. An impressive empirical literature has blossomed in recent decades, largely as data became available that allowed for easy estimation of the implicit value of life. However, the corresponding theory to justify the use of estimates from this literature in CBA is seriously underdeveloped.

In fact, many, if not all, of the value-of-life estimates used in CBA prior to the VSL's widespread adoption may be superior from the perspective of economic efficiency. This article, for example, proposes a financial approach to valuing lives, such that the marginal contribution to capital that corresponds with extending a life becomes the primary focus. Such an approach overcomes many difficulties found with the VSL and is also likely to be easier to implement in practice.

Although a financial approach faces some challenges of its own and may be somewhat counterintuitive to economists, its chief advantage is that it better accounts for the opportunity cost of resources than does the VSL. A financial approach is also likely to lead to greater overall reductions in risk due to its emphasis on wealth maximization, thereby increasing the resources available to devote to risk reduction (or indeed any other societal goal).

The article concludes with discussion of the policy relevance of these findings, focusing on the implications of switching from the VSL to a financial approach in the context of federal air quality regulations. Air pollution regulations are among the most consequential issued by the federal government, and the VSL turns out to be particularly ill-suited for use with these rules. At the very least, the VSL should be compared against other metrics, but it may also be time to reconsider the VSL altogether.

1 For evidence of the widespread use of the VSL at federal agencies, see agency guidelines on valuing mortality risk reductions. For example, US Department of Transportation, *Revised Departmental Guidance 2016: Treatment of the Value of Preventing Fatalities and Injuries in Preparing Economic Analyses*, 2016; US Environmental Protection Agency, *Valuing Mortality Risk Reductions for Policy: A Meta-Analytic Approach*, 2016; US Department of Health and Human Services, *Guidelines for Regulatory Impact Analysis*, 2016.

1. Overview of the VSL

Economists typically estimate the VSL using data on the salary premium commanded for a given increase in the level of job risk, although surveys are also sometimes used.[2] These data are available from labor markets, where individuals who work in risky professions are paid "compensating differentials" (i.e., excess wages) for taking on higher levels of job risk relative to their safer options. Implicit in the decision to accept money for risk is the value a worker assigns to his or her own life. For example, if 100 dollars is the minimum a worker is willing to accept for taking on an increased chance of death of one-in-100,000 in his lifetime, implicitly the worker values his expected remaining life at $100 / 0.00001 = $10 million.

If the worker is rational, one would expect that what he is willing to accept in terms of cash for risk should roughly equate with what he is willing to pay to reduce risk more generally on the margin.[3] This provides the basic logic for transferring labor market estimates of the VSL to other contexts, such as when valuing prevented fatalities achieved by regulations.

The VSL is not a constant, but rather it tends to vary across individuals and also across risks.[4] Despite this heterogeneity, US regulatory agencies tend to use an average VSL for the population, rather than a unique VSL for each person whose life they expect to save. Use of an average is partly for practical reasons, to avoid an unrealistic amount of work. It also helps to avoid the thorny issue of assigning different values of life to different people or groups of people. Richer individuals and countries tend to have higher VSLs than poorer ones,[5] and middle-aged and older adults tend to have higher VSLs than younger adults,[6] although estimates vary.[7]

Differences in VSLs across individuals and countries are largely due to disparities in wealth, as those with more wealth naturally have more to spend on risk reduction as they have greater means to do so. A 2017 study by Viscusi and Masterman estimated the average VSL for Ethiopia to be around $100,000, while the corresponding value for Norway was more than $16 million.[8] Some government analyses allow the VSL to vary by age group. For example, in 2017, President Trump's Council of Economic Advisors estimated the fatality cost of the opioid crisis to have been $432 billion in 2015, using age-varying VSLs.[9]

Such studies are generally a minority, however, as government and academic CBAs using a population-average VSL are more common. This raises complications, however, as an average VSL, rather than an individual or group-specific VSL, breaks the empirical link between values used in analysis and people's own preferences, and it is often argued that a key advantage of willingness to pay measures is that they are based on people's own values.[10] Use of an average runs the risk of making many people worse off from policy. For example, if regulations target groups such as low-income individuals who are willing to pay less

2 See Cass R. Sunstein, *Valuing Life: Humanizing the Regulatory State* (Chicago: University of Chicago Press, 2014); W. Kip Viscusi, "The Value of Risks to Life and Health," *Journal of Economic Literature* 31, no. 4 (1993): 1912–46.

3 There are reasons why WTP and WTA values can differ, including behavioral biases. For simplicity, this article will assume the two values are generally the same.

4 W. Kip Viscusi, *Pricing Lives* (New Jersey: Princeton University Press, 2018).

5 W. Kip Viscusi and Clayton J. Masterman, "Income Elasticities and Global Values of a Statistical Life," *Journal of Benefit-Cost Analysis* 8, no. 2 (2017): 226–50.

6 Joseph E. Aldy and W. Kip Viscusi, "Age Differences in the Value of Statistical Life: Revealed Preference Evidence," *Review of Environmental Economics and Policy* 1, no. 2 (July 1, 2007): 241–60.

7 Alan Krupnick, "Mortality-Risk Valuation and Age: Stated Preference Evidence," *Review of Environmental Economics and Policy* 1, no. 2 (July 1, 2007): 261–82.

8 Viscusi and Masterman, "Income Elasticities."

9 US Council of Economic Advisors, *The Underestimated Cost of the Opioid Crisis*, 2017.

10 For example, M. Jones-Lee and coauthors defend willingness to pay metrics on the grounds that "social decisions should, so far as possible, reflect the interests, preferences and attitudes to risk of those who are likely to be affected by the decisions." M. W. Jones-Lee, M. Hammerton, and P. R. Philips, "The Value of Safety: Results of a National Sample Survey," *Economic Journal* 95, no. 377 (1985): 49–72.

for the benefits of regulation than the average across society, these groups may be harmed if they still have to bear a significant portion of the regulatory costs.[11]

2. The Opportunity Cost of Funds

One fairly well-known problem with the VSL relates to difficulties that arise when trying to assign a value to life as life nears its end. In other words, how much should society be willing to pay to extend life by very short durations? While this issue may appear to be specific to age, it is part of a more general problem related to opportunity cost, namely, that the opportunity cost of resources to a single individual may be much different than the opportunity cost of those same resources to society (i.e., the sum total of all individuals).

Consider the example of an extravagant billionaire who has no heirs and is near the end of his life. He is willing to spend his entire fortune to stay alive as long as possible, even if it means spending every cent he has to stay on continued life support. Should the government also be willing to spend billions to see this man's life extended? What if all the billionaire's friends are willing to chip in?

Clearly if the government uses these individuals' preferences as a guide, the answer would seem to be unambiguous that, yes, welfare is improved if the government spends billions to extend the man's life by hours, minutes, and perhaps even seconds. After all, he is willing to pay as much for the benefit, so society should presumably be willing to pay as much too.

But it just cannot be true that it is reasonable to spend vast amounts to extend life by trivial amounts. The billions spent to extend one man's life could be used to save hundreds of other lives in some other context, and for longer durations, and people would be willing to pay for those benefits too.

It should be obvious that the opportunity cost of funds is much lower to the man who has very little time left than it is to other members of society.[12] Another way of stating this is that the billionaire's preferences are not aligned with the social interest. This is an extreme example, but the basic principle applies to consumption in almost any form. From an individual's perspective, one could be willing to spend a considerable amount to, say, go on an exorbitant luxury vacation, but from the perspective of all other members of society, those resources would be better off put to some other use. At the very least, if invested, the resources could generate the wealth to fund even more luxury vacations for more people.

This problem becomes more evident as life nears its end; but in fact, the problem extends across one's entire lifespan because individuals' generally value their own consumption higher than its opportunity cost to society.[13]

Some have proposed using the "value of a statistical life year," or VSLY, as an alternative to the VSL in some cases.[14] The VSLY is a measure of the value of one additional year of life, taken by dividing the VSL by an expected number of life-years remaining. This would tie the value of life closer to life expectancy. But the VSLY approach has problems too.

First, if analysts abandon the VSL and use the VSLY exclusively, this again breaks the empirical link between the values used in analysis and people's own values, since there is little to no empirical evidence supporting the view that people are willing to pay the same amount for each additional year of expected remaining life.[15] If analysts instead alternate between the two metrics depending on the age of policy beneficiaries, any cutoff age after which the VSLY substitutes for the VSL will be arbitrary.

11 See Cass R. Sunstein, "Are Poor People Worth Less Than Rich People? Disaggregating the Value of Statistical Lives" (Washington, DC: AEI-Brookings Joint Center for Regulatory Studies, 2004).

12 This general issue is discussed at length in Joanne Linnerooth, "Murdering Statistical Lives...?," in *The Value of Life and Safety*, Michael Jones-Lee (ed.) (Amsterdam: North-Holland, 1982).

13 Two related reasons for this are that individuals and society have different time horizons and different rates of discount.

14 See Viscusi, *Pricing Lives*; Cass R. Sunstein, "Lives, Life-Years, and Willingness to Pay," *Columbia Law Review* 104 (2004): 57; Michael J. Moore and W. Kip Viscusi, "The Quantity-Adjusted Value of Life," *Economic Inquiry* 26, no. 3 (1988): 369–88.

15 Aldy and Viscusi, "Age Differences."

Even if the VSLY was a conceptually appropriate solution to the problems related to opportunity cost raised here, it is not clear how politically feasible it is to use it, since much higher values are currently used in analysis. The Environmental Protection Agency during the George W. Bush administration tried to use a discounted VSL for citizens older than 65, as part of its Clear Skies Initiative.[16] In response, some critics claimed that a "senior death discount" was being applied to the lives of the elderly.[17]

The VSLY may well be superior to the VSL in many instances, but given the outrage that arose in response to the Bush administration's attempts to adjust the VSL, it seems unlikely a future administration will make any similar adjustments to the VSL soon.

3. Identifiable and Statistical Lives

Proponents of the VSL sometimes state that they are putting a dollar value on risk, not on actual people's lives,[18] but this is semantics. When the government estimates a rule will save two people, it multiplies the VSL by two. This is mathematically the same as aggregating what two million people would be willing to pay to eliminate a lifetime death risk of one in a million for each of them. There is no mathematical difference between applying a dollar value to risk and applying it to someone's actual life. In fact, risk levels only change when regulations reduce mortality (i.e., save someone). If no lives are saved, there are no corresponding changes in risk for a population, and there are no benefits to which one can attach a dollar value. Nor is it true the VSL is only applied to small risks, since a death, by its very nature, is a large event.

What makes lives statistical in the sense that economists use the term is whether they are identifiable or not.[19] If a child falls down a well, the identity of the child whose life is at stake is known. With government regulation, it is often unclear whose life is at stake. Those lives are statistical.

Empirically, however, it turns out that people are often willing to pay much more to save known individuals than to save unknown statistical people,[20] despite the fact that in either case, someone's actual life is on the line.[21] VSL advocates generally respond by arguing the VSL should only be applied to statistical lives,[22] but this raises the question of whether it is efficient for the same life to be valued differently depending on whether the identity of the person is known. After all, it would seem the opportunity cost of the resources required to save them has not changed.

4. The Time Value of Money

When putting a dollar value on a life, or indeed on any nonpecuniary benefit, a general kind of problem emerges, which is that it often makes sense to delay public investments because this will improve the cost-benefit calculus.

This problem actually has two parts. First, money, unlike lives, can be invested in an account and earn interest. Suppose a project costs $1 million dollars and would save two lives, but doubling spending to

16 US Environmental Protection Agency, *Technical addendum: Methodologies for the benefit analysis of the Clear Skies Initiative*, 2002.
17 Cindy Skrzycki, "Under Fire, EPA Drops the 'Senior Death Discount,'" *Washington Post*, May 13, 2003.
18 For example, the Office of Management and Budget states in its regulatory instructions to agencies that, "Some describe the monetized value of small changes in fatality risk as the 'value of statistical life' (VSL) or, less precisely, the 'value of a life.' The latter phrase can be misleading because it suggests erroneously that the monetization exercise tries to place a value on individual lives. You should make clear that these terms refer to the measurement of willingness to pay for reductions in only small risks of premature death. They have no application to an identifiable individual or to very large reductions in individual risks. They do not suggest that any individual's life can be expressed in monetary terms. Their sole purpose is to help describe better the likely benefits of a regulatory action." See US Office of Management and Budget, *Circular A-4: Regulatory Analysis*, September 17, 2003.
19 James K. Hammitt and Nicolas Treich, "Statistical vs. Identified Lives in Benefit-Cost Analysis," *Journal of Risk and Uncertainty* 35, no. 1 (August 1, 2007): 45–66.
20 Ibid.
21 Joanne Linnerooth, "Murdering Statistical Lives...?,"; Lisa Heinzerling, "The Rights of Statistical People," *Harvard Environmental Law Review* 24 (2000): 189–207.
22 This is the approach taken in Office of Management and Budget guidelines, for example. See footnote 18 above.

$2 million would save four lives. There is a reasonable argument for investing the first million in financial markets (instead of spending it immediately on the life-saving project), until it grows into $2 million. This way more lives can be saved. The paradox is that it seems rational to keep the $1 million down payment invested indefinitely as more and more lives can be saved the longer we wait.

The second part of this problem is that when analysts value resources based on what individuals are willing to pay for them, resources' value will tend to rise over time due to economic growth, since willingness to pay for normal goods rises with income. Therefore, in choosing between two equally expensive projects, one that saves a life today and another that saves a life tomorrow, the second project will be the better option because the life in the future can be expected to be worth more.

It has been argued that discounting future benefits and costs, including discounting future lives, helps resolve this waiting problem.[23] In fact, some have argued that failing to discount in CBA would mean that the current generation should starve itself for the benefit of the future (presumably because 100 percent of income should be invested each period).[24] But this is problematic. First, it seems strange if the purpose of the social discount rate is to resolve anomalies that arise when valuing resources in a particular manner. If other valuation techniques are available that avoid such anomalies, why not use these alternatives instead? The practice of discounting the lives of our children and grandchildren is controversial, and if this practice could be abandoned by abandoning the VSL, that could be a worthwhile tradeoff.

Second, there is no guarantee that a social discount rate even resolves this problem, which can arise with or without a social discount rate. If willingness to pay is generally rising over time as a result of economic growth, then if the rate at which willingness to pay is rising is equal to or greater than the social discount rate used in analysis, the waiting problem reemerges even with discounting.

Finally, from an efficiency standpoint, this delay "problem" may not be a problem at all, but instead a recommendation. It is simply an implication of the time value of money, which says:

> If you are promised $1 of cash flow, for certain, ten years from now, this dollar would not be as valuable to you as $1 is today because, if you had the $1 today, you could invest it and end up with more than $1 in ten years.[25]

Financial analysis compares cash flows with other cash flows, but it is easy to see that if cash flows are compared to flows of consumption, an implication of the time value of money is that it is best to delay consumption in order to invest, if your objective is to maximize wealth. This is because consuming means forgoing interest. Some refer to this as an "anomaly,"[26] but if it is truly an anomaly to recommend delaying consumption in order to maximize wealth, is it not also an anomaly to prefer cash flows as soon as possible in order to maximize wealth?[27]

None of this means that the current generation should starve. Some consumption contributes to production and investment, and is therefore efficient. Without her morning breakfast, the factory worker has no energy and is little use as a laborer. Without spending on room and board, the college student does not accumulate human capital that raises her productivity later in her adult career. Furthermore, at some point adding more investment becomes counterproductive, since so many resources will have to be devoted to replacing depreciating capital each year that consumption levels actually fall.

---

23 W. Kip Viscusi, John M. Vernon, and Jospeh E. Harrington Jr., *Economics of Regulation and Antitrust*, 4th ed. (Cambridge, MA: MIT Press, 2005); W. Kip Viscusi, "Rational Discounting for Regulatory Analysis," *University of Chicago Law Review* 74, no. 1 (2007): 209–246.

24 Cass Sunstein and David Weisbach attribute this claim to Tjalling Koopmans. See David Weisbach and Cass Sunstein, "Climate Change and Discounting the Future: A Guide for the Perplexed," *Yale Law & Policy Review* 27, no. 2 (December 2, 2015).

25 Frederic Mishkin, *The Economics of Money, Banking, and Financial Markets*, 11th ed. (London: Pearson Education, 2016), 112.

26 For example, US Office of Management and Budget, *Circular A-4: Regulatory Analysis*, September 17, 2003; Viscusi, "Rational Discounting."

27 Note that it can be perfectly rational and utility-maximizing from an individual's perspective to forgo interest in order to consume. What the time value of money implies is that this choice, even if sensible from an individual's perspective, is not the wealth-maximizing choice.

Nor will a failure to discount necessarily lead to "policy paralysis" as some have claimed. For example, Kip Viscusi has argued, along with coauthors, that a zero discount rate leads to what he calls a "permanent cost slam dunk."[28] He states, "in a world with no discounting, a $1 loss forever will have an infinite present value, swamping any finite amount of benefits and leading to policy paralysis."

It is certainly true that, without discounting, ongoing costs, such as those from foregone investments, would have an infinite present value. But ongoing benefits also have an infinite present value. The policy paralysis Viscusi describes is a function of comparing ongoing costs to fleeting benefits from consumption. The issue he identifies is easily resolved if the analyst simply values benefits in terms of the cash flows generated. Then—due to the time value of money—the sooner a benefit arrives the better, and the policy paralysis can be overcome.

5. Individual and Social Values

Economists tend to be very sympathetic to the idea of using willingness to pay as a basis for measuring the benefits of public policies. They are also very attached to using people's preferences as a guide for policy. As a result, some economists might be understandably uncomfortable with a financial approach to valuing benefits and costs. However, this approach is not so far removed from conventional economic theory as it might at first appear.

Market data used to estimate the VSL necessarily come from those who are alive at the time a study is conducted, so these studies do not by definition take into account the preferences of people in the future who are not yet born. It is easy to imagine that future generations, if they could trade in our markets, would be willing to pay something for actions that benefit them, as well as be willing to pay to stop actions that harm them. However, there is a missing market that prevents them from doing so.

A more comprehensive measure of willingness to pay would take into account the preferences of future generations. This means the total, or social, value of a resource might differ substantially from the value to those who are currently living,[29] and that what current individuals are willing to pay to reduce risk may not be consistent with what the sum total of all individuals would be willing to pay.[30] This suggests data collected from labor markets and then used to estimate the VSL may be misleading. Even if these data capture the private benefits to individuals from reducing risk, they won't capture the social, or total, value.

Like with other problems, this issue is not unique to the VSL, but rather is a general kind of problem that arises when using the willingness to pay of current citizens to value benefits. There is no reason to believe that individuals' preferences accord with the interests of society more generally.[31] As a result, their measured willingness to pay will have to be adjusted.

What economists seem to have done is confused *what is* (i.e., how markets actually operate) for what *ought to be* (i.e., what is economically efficient). Markets allocate resources to those currently living citizens who are willing to pay the most for them, but there is no reason to believe this is wealth-maximizing, i.e., efficient, once one considers a time frame that extends beyond the current generation.

28 W. Kip Viscusi, Joel Huber, and Jason Bell, "Responsible Precautions for Uncertain Environmental Risks," *Journal of Benefit-Cost Analysis* 10, no. 2 (2019): 296–31.

29 This is true even if citizens are maximizing their own utility.

30 It is even conceivable that people in the future would be willing to pay present citizens to accept *more* risk, for example in exchange for more investment and economic growth.

31 See footnote 13 above.

6. A Financial Approach to Valuing Lives

This article has identified at least four problems with the VSL. These relate to overlooked opportunity cost, differential treatment of known versus statistical lives, "anomalies" arising as a result of the time value of money, and a deviation between individual and social willingness to pay values.

Along each of these dimensions, a financial estimate of life's value is superior to the VSL. First, a measure of life's value connected to one's productive contributions better accounts for the opportunity cost of resources than does the VSL. The focus of CBA becomes generating as much wealth as possible, ultimately enabling the most lives to be saved.

A policy of wealth maximization also handles challenges arising at the end of life. Unlike the VSL, with a financial approach the lives of the young will generally be valued greater than the lives of the elderly, since the young have more years of productive life remaining. This conforms with common sense ethics.

There is even reason to believe such an approach is less controversial to apply in practice, because the analyst applies dollar values only to goods and services and not to people. This would allow analysts to avoid the controversial practice of discounting the lives of future citizens. Courts and other entities sometimes use a variant of the financial approach when determining awards to survivors in wrongful death suits or in response to tragedies. For example, the September 11th Victim Compensation Fund made payouts to victims' families based on projections of the lost lifetime earnings of the deceased.[32] While it wouldn't be fair to call these judgments uncontroversial, they have proven less politically toxic than applying discounts near the end of life.

A financial approach also resolves discrepancies that arise between what people are willing to pay to save known lives versus statistical lives. With a financial approach, the life of a person with a certain demographic profile would have the same expected value regardless of whether the person's identity is known or not.

Next, the most obvious way to avoid the anomaly that life-saving projects should be continually delayed is to value a life in terms of the net financial cash flows it generates. Since cash is preferred sooner rather than later, due to the time value of money, the earlier a life is saved in general the better (assuming contributions are positive).

Finally, a financial approach overcomes issues related to the divergence between individual and social values. Citizens currently alive are naturally short-sighted as they have finite lifespans. As a result, they are likely to give too much weight to temporary fleeting benefits that come in the form of consumption, relative to investments with large dividends, but that pay off in the distant future. Society as a whole has no such time preference, however, and once one considers that people in the future would be willing to pay something for investments that benefit them, it becomes clear that their willingness to pay eventually swamps that of current citizens.

Taken together, these arguments suggest that a financial approach to valuing lives has significant appeal, but such an approach also raises some concerns that are worth noting. First, some individuals do not earn incomes but do create financial costs when they die. For example, grandparents provide valuable services when they babysit their grandchildren, but they aren't paid for them. Replacing the services grandparents provide costs something in financial terms, so replacement cost might be the right way to value lives financially. In his book, *Stubborn Attachments: A Vision for a Society of Free, Prosperous, and Responsible Individuals,* George Mason University economist Tyler Cowen argues that replacement cost may be the

32 Kenneth Feinberg, *Who Gets What: Fair Compensation after Tragedy and Financial Upheaval* (New York, NY: PublicAffairs, 2012).

correct approach.[33] Replacement cost has at times been used by the federal government, most notably in the military.[34] These examples suggest the concept has some mainstream support.

Perhaps the greatest concern with a financial cost of death approach is that it seemingly leaves out aspects of life that are not financial. However, the utility we experience during our lives is fleeting. This may seem tragic, but on the other hand, what we produce during our lifetimes has the potential to continue to bear fruit long after we are gone. If our time horizon extends into the distant future, it is our contributions to the capital stock, to production, and to the body of human knowledge, that will live on long after we are gone.

Consider the workers who built the Brooklyn Bridge. Such men no doubt made valuable contributions to their families and to their communities while they were alive. These contributions were important, but they were also temporary. And yet the contribution of these workers to the economy continues to add growing value to this very day and will continue long into the future as well.

7. Policy Implications for Clean Air Act Rules

The value of life chosen in CBA is critical because it will help determine whether public policies generate more in terms of benefits to society than costs. If the value of life used is inconsistent with a goal of maximizing economic efficiency, this can have widespread ramifications for public policy.

For example, the Environmental Protection Agency's (EPA) air quality rules are among the most expensive issued by any agency in the federal government, and they provide a useful case study of how a financial approach to valuing lives might work in practice, and also how this might change policy recommendations substantially.

First, there is an ongoing debate about how the benefits from air quality regulations should be valued. Much of the debate centers around the health benefits associated with reducing fine particulate matter (commonly referred to as PM2.5). Because these benefits are often not directly targeted by regulation, but are instead a coincident benefit that occurs indirectly as a consequence of targeting other pollutants, like mercury or carbon dioxide, some economists have wondered whether such indirect benefits should be counted in analysis.[35] Meanwhile, another debate rages about whether PM2.5 is really as dangerous to public health as some regulatory agencies claim it is.[36]

This article will not address these controversies. Instead, for the sake of argument, it will take for granted that the health effects of PM2.5 are real and harmful, and that indirect coincident benefits should be counted in analysis alongside other direct costs and benefits. Even taking these factors as given, there remains the question of whether PM2.5-related benefits are valued appropriately in the EPA's regulatory impact analyses, since those values rely in major part on the VSL.

In writing on the health effects of PM2.5, one recent study noted that "about 80 percent of the burden of air pollution is borne by the elderly."[37] Furthermore, "the aggregate mortality burden of PM2.5 is concentrated among the elderly with five to ten years of remaining life expectancy, followed by those with two

33 Tyler Cowen, *Stubborn Attachments: A Vision for a Society of Free, Prosperous, and Responsible Individuals* (San Francisco: Stripe Press, 2018).
34 H. Spencer Banzhaf, "Retrospectives: The Cold-War Origins of the Value of Statistical Life," *Journal of Economic Perspectives* 28, no. 4 (November 2014): 213–26.
35 One issue is the asymmetric counting of indirect benefits without consideration of indirect costs. See Susan Dudley et al., "Consumer's Guide to Regulatory Impact Analysis: Ten Tips for Being an Informed Policymaker," *Journal of Benefit-Cost Analysis* 8, no. 2 (2017): 187–204.
36 For a small sample of studies challenging prevailing views about PM2.5, see Louis Anthony Cox, "Miscommunicating Risk, Uncertainty, and Causation: Fine Particulate Air Pollution and Mortality Risk as an Example," *Risk Analysis: An Official Publication of the Society for Risk Analysis* 32, no. 5 (May 2012): 765–67, 768–70; Louis A. Cox, Douglas A. Popken, and Paolo F. Ricci, "Warmer Is Healthier: Effects on Mortality Rates of Changes in Average Fine Particulate Matter (PM2.5) Concentrations and Temperatures in 100 U.S. Cities," *Regulatory Toxicology and Pharmacology* 66, no. 3 (August 1, 2013): 336–46; Ke Zu, Ge Tao, Christopher Long, Julie Goodman, and Peter Valberg, "Long-Range Fine Particulate Matter from the 2002 Quebec Forest Fires and Daily Mortality in Greater Boston and New York City," *Air Quality, Atmosphere & Health* 9, no. 3 (April 1, 2016): 213–21.
37 Karen Clay and Nicholas Z. Muller, "Recent Increases in Air Pollution: Evidence and Implications for Mortality" (NBER Working Paper 26381, National Bureau of Economic Research, Cambridge, MA, October 2019).

to five years remaining"[38] In other words, EPA air rules are a slightly less extreme example of the problem identified with the extravagant billionaire above. Given that current estimates of the VSL used in regulatory impact analysis are in the range of $10 million dollars, and that these values are attached to the lives of the elderly, it seems likely that the EPA's benefit estimates are off. Therefore, it's worth considering how alternative metrics might change policy conclusions.

As discussed above, one alternative to the VSL is the VSLY. One recent estimate of the VSLY was $369,000,[39] which, if multiplied by five years of remaining life for a typical person exposed to air pollution, would be equal to a little under $2 million, or less than one fifth of a typical VSL value. A recent estimate from Hugonnier, Pelgrin, and St-Amour put the average "human capital value of life" in the United States at $421,000, or roughly one-twentieth of the average VSL for the United States.[40] The human capital approach values lives in terms of the stock of human capital people possess, which gradually depreciates to zero by the end of life. A value for the elderly would presumably be close to zero for this reason. This change alone would likely wipe out most of the benefits of EPA air regulations.

Another approach was taken when making payouts from the September 11th Victim Compensation Fund, which based compensation on the foregone earnings of the deceased. The average payment between 2001 and 2003 was around $2 million,[41] still far below the VSL. Further, these payments often related to individuals in the prime of their working-age lives, many of whom worked in the financial sector and may have had higher-than-average earnings compared to Americans more generally.

An approach based on human capital or foregone earnings has a long history of being criticized by economists.[42] However, it is important to note that there is a difference between the financial cost of death approach as advocated in this paper and a human capital or earnings approach. The approach advocated here is an assessment of how much money preventing a death saves, which is a proxy for the amount of productive capital a life is expected to contribute on the margin. Such an approach may only be loosely connected to one's future earnings. For example, if a firm loses a worker who passes away unexpectedly, the firm will still have to compensate some other worker who presumably replaces the deceased worker. Thus, the resource cost to society from losing a worker often differs substantially from the worker's lifetime earnings.

In general, earnings should probably be viewed as an upper bound on the replacement cost of a worker, except in cases where there are significant external benefits to society not captured in an employee's earnings.[43] The opportunity cost of replacing a worker may be fairly low in cases where search costs are low or if a replacement is drawn from a pool of the unemployed. At times, the cost of death may even be negative, for example with a retiree who consumes or has expensive medical bills, but doesn't earn. This has direct relevance to the EPA's air quality regulations, since presumably many of the beneficiaries of these regulations are retired.

According to the Bureau of Labor Statistics, the total average annual consumption expenditure made by a consumer unit aged 65 years and older was $49,542 for 2017; for those 65 to 74 years old it was $54,997; and for those 75 years and older it was $41,849.[44] Notably, these expenditures exceed the average

38 Tatyana Deryugina, Garth Heutel, Nolan H. Miller, David Molitor, and Julian Reif, "The Mortality and Medical Costs of Air Pollution: Evidence from Changes in Wind Direction," *American Economic Review* 109, no. 12 (2019): 4178–219.

39 See Kip Viscusi and Thomas Kniesner, "The Value of a Statistical Life," in *Oxford Research Encyclopedia of Economics and Finance* (Oxford: Oxford University Press, forthcoming), citing US Department of Health and Human Services, *Guidelines for Regulatory Impact Analysis,* 2016.

40 Julien Hugonnier, Florian Pelgrin, and Pascal St-Amour, "Valuing Life as an Asset, as a Statistic, and at Gunpoint" (HEDG Working Paper 18/20, University of York, 2018).

41 CNN, "September 11th Victim Aid and Compensation Fast Facts," July 3, 2019.

42 Jacques Drèze, "L'Utilité Sociale d'une Vie Humaine," *Revue Française de Recherche Opérationnelle* 6 (1962): 93–118; Thomas Schelling, "The Life You Save May Be Your Own," in *Problems in Public Expenditure Analysis*, ed. Samuel B. Chase, Jr. (Washington, DC: Brookings Institution, 1968), 127–62; E. J. Mishan, "Evaluation of Life and Limb: A Theoretical Approach," *Journal of Political Economy* 79, no. 4 (1971): 687–705.

43 This might be true, for example, of an artist like Beethoven. His contributions to society to-date have no doubt far exceeded the value of his lifetime earnings.

44 US Bureau of Labor Statistics, "Table 1300. Age of Reference Person: Annual Expenditure Means, Shares, Standard Errors, and Coefficients of Variation, Consumer Expenditure Survey, 2017," in *Consumer Expenditure Survey*, September 2018.

pre-tax income of consumer units in these age brackets, suggesting that seniors are consuming more they earn. The public sector also incurs costs supporting the elderly. For example, the 2016 Medical Expenditures Panel Survey put the average medical expenditure for a 65-year-old at $7,305, for a 75-year-old at $10,511, and for an 85-year-old at $12,703.[45] Much of the healthcare costs associated with the Medicare program also come in the last year of life.[46]

What about the contributions of seniors? One study found that, for the year 2006, the average annual production value of an American aged 60–64 was $43,565, for 65–69 it was $26,139, for ages 70–74 it was $21,111, for 75–70 it was $17,698, and for 80 and over it was $13,375.[47] Thus, as medical expenditures tend to rise at the end of life, production value tends to fall.

None of this should be particularly surprising, as these are simple facts of life. Furthermore, air pollution has other consequences, which should be considered. For example, it may increase infant mortality or asthma rates.[48] And yet, as hard as it may be for regulatory agencies to admit it, PM2.5-related emissions may be cost-saving on balance insofar as they reduce net consumption expenditures with respect to the elderly. This creates an ethical as well as an economic dilemma: all of us want to support seniors, but resources spent supporting them have an opportunity cost. Current federal regulatory impact analyses do not account for this opportunity cost; instead they rely on the VSL.

8. Discussion and Conclusion

The VSL has many shortcomings, many of which extend to the general practice of monetizing nonmarket goods using measures of current consumer willingness to pay. Taken together, these problems suggest use of the VSL produces inefficient policy recommendations, and this has direct relevance to some of the most significant and expensive regulations issued by the federal government.

By contrast, a financial approach, based upon the amount of money that is saved by extending a life, is in a good position to provide a robustness check against the VSL or to supplant it altogether in benefits analysis. A financial approach better accounts for the opportunity cost of resources, conforms with the common sense notion that the young have more productive years of life remaining than the elderly, does not require different valuations for identifiable and statistical lives, and, for practical purposes, may be easier to allow to vary across individuals, and thus be more economically efficient.

Why have the problems with the VSL, many of which have been known for a long time, gone unaddressed for so long by scholars? One reason is likely because regulatory agencies may prefer the VSL. The VSL tends to be much higher than competing estimates, making it easier for agencies to justify expensive regulatory interventions. Another reason is likely because the literature on the value of life tends to emphasize measurement over theory. Early discussions among economists about the value of life were theoretical and philosophical,[49] but out of these discussions no real consensus emerged. By contrast, once data became available, especially related to workplace fatalities, economists quickly developed the empirical tools to estimate the VSL with impressive precision. Its dominance emerged in the academic literature

45 US Department of Health and Human Services, Agency for Healthcare Research and Quality, *Medical Expenditure Panel Survey*, 2016.
46 Ian Duncan, Tamim Ahmed, Henry Dove, and Terri L. Maxwell, "Medicare Cost at End of Life," *American Journal of Hospice & Palliative Care* 36, no. 8 (August 2019): 705–10.
47 Scott Grosse, Kurt Krueger, and Mercy Mvundura, "Economic Productivity by Age and Sex: 2007 Estimates for the United States," *Medical Care* 47, no. 7 (July 2009).
48 Kenneth Y. Chay and Michael Greenstone, "The Impact of Air Pollution on Infant Mortality: Evidence from Geographic Variation in Pollution Shocks Induced by a Recession," *Quarterly Journal of Economics* 118, no. 3 (August 1, 2003): 1121–67; Michael Guarnieri and John R. Balmes, "Outdoor Air Pollution and Asthma" *Lancet*. 383, no. 9928 (May 3, 2014): 1581–1592.
49 See for example: John Broome, "Trying to Value a Life," *Journal of Public Economics* 9 (1978): 91–100; E. J. Mishan, "Evaluation of Life and Limb: A Theoretical Approach," *Journal of Political Economy* 79, no. 4 (1971): 687–705; Jacques Drèze, "L'Utilité Sociale d'une Vie Humaine," *Revue Française de Recherche Opérationnelle* 6 (1962): 93–118; Thomas Schelling, "The Life You Save May Be Your Own," in *Problems in Public Expenditure Analysis*, ed. Samuel B. Chase, Jr. (Washington, DC: Brookings Institution, 1968), 127–62; Joanne Linnerooth, "Murdering Statistical Lives...?," in *The Value of Life and Safety*, Michael Jones-Lee (ed.) (Amsterdam: North-Holland, 1982).

around that time but without the corresponding theory to justify its use in CBA. The empirical cart seems to have gotten ahead of the theoretical horse in this case.

Regulation has the potential to produce dramatic improvements in the lives of citizens, but only if it contributes more to societal well-being than it creates in terms of harm. In recent decades, the VSL has emerged as one of the most widely used tools for assessing these economic tradeoffs inherent when regulating, but the theoretical problems with it are myriad, suggesting its use is producing misleading policy recommendations. At the very least, these issues should generate discussion. In the end, however, the VSL may simply be an analytical contrivance too problematic to provide a meaningful answer to the enduring question of how to value a human life.