CAMBRIDGE
UNIVERSITY PRESS

ARTICLE

Inequality and the Value of a Statistical Life

Cass R. Sunstein

Harvard University, Cambridge, MA, USA
E-mail: csunstei@law.harvard.edu

Abstract

A uniform value of a statistical life (VSL) is part of established practice within the federal government. Some people have applauded a uniform VSL on the ground that it respects the equality of persons; takes harm to poor people as seriously as it does harm to wealthy people; avoids expressive harms; and builds appropriate wealth redistribution into regulatory policy. Other people have strenuously objected to a uniform VSL, emphasizing that to reduce mortality risks, poor people are willing to pay less than rich people are, and urging that poor people should not have to pay more than they are willing to pay. Whether a uniform VSL is in the interest of poor people depends on whether we are dealing with subsidies or regulations. In the case of subsidies, a uniform VSL is highly likely to benefit poor people. If we are dealing with regulations, we cannot know whether a uniform VSL helps or harms poor people without knowing the incidence of costs (and benefits).

1. Your money or your life

The value of a statistical life (VSL) is like toothpaste: Many people use it, but no one really loves it.[1] The problem might lie in the (misleading) name; who loves the idea of putting a monetary value on a life, even if it is merely a statistical one? If we spoke more precisely of the value of a statistical mortality risk, the idea might be a little more lovable. But even if so, it would not be a lot more lovable. It is not much fun to think about mortality risks, and if we are turning them into monetary equivalents, well, who could love that?

Still, public officials have to assign monetary values to mortality risks, even if they do so only implicitly. If they are deciding on one level of stringency and not another, they are necessarily deciding that level of stringency is worthwhile (and that another level of stringency is not worthwhile). The question is not whether to have a VSL; it is what VSL to have. Within the federal government in the United States, VSL is the workhorse of cost-benefit analysis, certainly for purposes of regulatory policy; it is the principal driver

---

[1] For a possible exception (with respect to VSL, not necessarily toothpaste), see Viscusi (2013, 2018).

of benefits in multiple domains, whether we are speaking of highway safety, road safety, food safety, cigarettes, or pandemics.[2] That point in turn raises a question: Should agencies use a unitary VSL, or should VSL vary with wealth? Should poor people be taken to have the same VSL as rich people?[3] Within the federal government, the general practice is clear: VSL is unitary.[4]

For orientation, suppose that there are only two kinds of cars: safe cars and safer cars. Suppose also that there are only two kinds of people: poor people and rich people. Safer cars are more expensive than safe cars. Let us suppose that poor people buy safe cars, not safer cars, because they are not willing to pay the selling price of safer cars. Let us suppose too that rich people buy safer cars, not safe cars, because they are willing to pay the premium. Under imaginable assumptions, we might want to subsidize safer cars so that everyone can buy them (though it might well be better to make a lump-sum payment to the poor). But it is not so easy to see why we might want to require everyone, poor and rich, to buy safer cars. Perhaps it could be said that poor people lack relevant information or suffer from some kind of behavioral bias (such as present bias or limited attention), such that a regulatory mandate would be in their interests. But that claim would have to be demonstrated, not simply asserted, and no absence of information or behavioral bias could support the view that poor people ought to be required to spend as much on motor vehicle safety as rich people are willing to spend on motor vehicle safety.

A great deal of empirical work uncovers the unsurprising fact that to reduce statistical risks, poor people are willing to pay less than rich people.[5] The empirical findings are in line with a longstanding claim in utilitarianism, which is that a given unit of money is worth more to a poor person than to a rich person.[6] If the five richest people in the world were given $5000, the gift would matter little if at all to their welfare; if the five poorest people in the world were given $5000, the gift would make all the difference in terms of their welfare. It follows that there is a strong argument in favor of redistribution on simple utilitarian grounds (bracketing incentive effects). Redistribution through taxation might seem the best and most natural response.[7] But if the tax system is unavailable for that purpose, for political or other reasons, perhaps there are other routes.[8]

My central claims here are straightforward. *First*: Whether a uniform VSL is in the interest of poor people depends in the first instance on whether we are dealing with subsidies or regulations. *Second*: In the case of subsidies, the use of a VSL that reflects the population average, or even the number for the top quartile, would be very much in the interest of poor people; in all probability, they would be net gainers. *Third:* In the case of regulations, a uniform VSL is more likely to benefit poor people if they do not pay all or most of the cost; it is more likely to harm them if they pay all or most of the cost. To know the effect of a uniform VSL on poor people in the case of regulations, it is essential to know the incidence of the costs (and also the benefits). These claims bear on the complicated questions of

---

[2] See Hemel (2022).
[3] See Hemel (2022).
[4] See Department of Homeland Security (2021) and Department of Transportation (2021).
[5] See Viscusi (2010) and Greenberg *et al.* (2021).
[6] See Hemel (2022).
[7] See Kaplow and Shavell (1994). On some of the complexities, see Adler (2016), Revesz (2018), and Liscow (2022).
[8] See Liscow (2022).

"distributional weights" on the benefits side of the ledger; they may or may not benefit poor people, depending on what is happening on the costs side (and in particular on whether poor people are paying those costs).

2. Five problems

Let us anchor the discussion in a set of stylized problems, understood to be simplified (but not unrealistic) descriptions of real-life situations.[9]

2.1. Problem 1

The U.S. government is funding various programs to diminish risks, including mortality risks.[10] To receive funding, communities that apply for funding must show that the benefits of their programs justify the costs. Some of those communities are poor; some are wealthy. In poor communities, the average VSL is $1 million. In wealthy communities, the average VSL is $20 million. The U.S. government uses a population-wide average, which is $10 million.

Use of the population-wide average is highly likely to be in the interest of poor communities. The government's policy gives them an in-kind subsidy: They might be willing to pay $10 to avoid a mortality risk of 1 in 100,000, but the government acts as if they are willing to pay $100. Because the government is footing the bill, people in poor communities are unlikely to lose. True, they might prefer the money as an unrestricted cash transfer to the boosted VSL in a risk-reduction program. But they can only gain, and cannot lose, from treating their VSL as if it were $10 million.

The major qualification is that the money for the grant must come from somewhere, and we cannot rule out the possibility that its use here will mean it will be unavailable for some other program or initiative from which poor people would benefit more. Does the program here entail higher taxes on the rich? Spending cuts from programs that help the poor? Deficit spending? If we know that the funding does not entail some loss for the relevant community, or for poor people in general, the optimistic conclusion holds.

For comparison: Suppose that property values, in poor communities, are low, so that such communities cannot show high monetized benefits from protecting property in those communities. Suppose that the government believes that the monetized benefits understate the welfare benefits. In a rich community, the average home might be valued at $3 million, and in a poor community, the average home might be valued at $200,000. The *monetary* loss is much lower in the poor community, but the *welfare* loss from the destruction of the

---

[9] For relevant discussion, see Hemel (2022). Hemel refers to a variety of sources, on which I draw here. See Harrison (1975) (finding that the cost of automobile emission controls as a proportion of income is larger for lower-income households than for higher-income households, while the distribution of benefits by income is less clear); Dorfman (1977) (estimating the benefits of pollution control based on self-reported willingness to pay and finding that pollution control imposes net costs on lower-income households and yields net benefits for higher-income households); Gianessi *et al.* (1979) (finding that air-pollution controls impose net costs on higher-income households and generate net benefits for lower-income households while automobile emissions controls impose net costs on all income groups); Kahn (2001) (finding that "regulation under the Clean Air Act has helped, and not economically harmed, the 'have nots.'").

[10] See McGee (2021).

average home in the poor community might be equal to or greater than the *welfare* loss from the destruction of the average home in the rich community.[11] If the government puts some kind of distributional weight on the average home in the poor community, to capture that welfare loss in the context of a subsidy program, it is helping members of that community.

2.2. Problem 2

The U.S. government is regulating motor vehicles to diminish risks, including mortality risks. The cost of the regulation is borne entirely by consumers.[12] Some of the people who bear the costs of the regulation are poor; some are wealthy. For poor consumers, the average VSL is $1 million. For wealthy consumers, the average VSL is $20 million. The U.S. government uses a population-wide average, which is $10 million.

Use of the population-wide average is not in the interest of poor people.[13] It is as if they are subjected to a forced transfer on terms that they dislike. They are effectively required to purchase more safety than they want. They are willing to pay only $10 to avoid a risk of 1 in 100,000, but the government is requiring them to pay $100. If their VSL is treated as if it were $10 million, they cannot gain; they can only lose. At least this is so if they do not suffer from a lack of information or some kind of behavioral bias.[14]

A notation: If we use a population-wide average in such circumstances, and add a distributional weight to the benefits enjoyed by poor people, we will make things even worse and possibly much worse.[15] Suppose that poor people are willing to pay $10 to avoid a risk of 1 in 100,000. If we treat that risk as if it were (equivalent in monetary value to) a risk of 1 in 50,000, or 1 in 10,000, we will be forcing poor people to spend more than they are willing to pay for a risk of X amount – and thus forcing them to pay much more than they want. That is not a very nice thing to do.

2.3. Problem 3

The U.S. government is regulating an air pollutant to diminish risks, including mortality risks. Some of the people who bear the costs of the regulation are poor; most are wealthy. For poor people, the average VSL is $1 million. For wealthy people, the average VSL is $20 million. The U.S. government uses a population-wide average, which is $10 million.

Use of the population-wide average might well be in the interest of poor people; we cannot know without obtaining more information about the incidence of the costs. It is true that poor people are willing to pay only $10 to avoid a risk of 1 in 100,000. But it is not at all clear that they are being required to purchase more safety than they want. Because most of the people who are paying to reduce the risk are wealthy, poor people might be required to pay less than $10 to avoid risk of 1 in 100,000. If so, they will gain.

These conclusions are tentative, because poor people might turn out to be net losers from the regulation, if they do not enjoy most of the benefits. Suppose that the benefits are

---

[11] See Sunstein (2007).
[12] On the realism of that assumption, see Hemel (2022).
[13] See Hemel (2022).
[14] See Sunstein (2013).
[15] See Hemel (2022).

concentrated among rich people, so that they enjoy the overwhelming majority of the mortality and morbidity gains. We would have to obtain a host of numbers to know, but it is possible that poor people are net losers, even with a VSL of $10 million. Assume, for example, that they pay 49 per cent of the total costs (which, let us suppose, are $800 million), but receive 2 per cent of the total benefits (which are, let us suppose, $1 billion). If so, they are likely to be losers. We could imagine other assumptions that would lead to the opposite conclusion.

2.4. Problem 4

The U.S. government is regulating to protect the safety of construction workers in a well-defined region of the country; the benefits will be enjoyed entirely by those workers. The cost of the regulation would be passed onto relatively wealthy people. For construction workers, the average VSL is $2 million. For those who bear the costs, the average VSL is $15 million. The U.S. government uses a population-wide average, which is $10 million.

Use of the population-wide average is likely to be in the interest of the construction workers. It is true that they are willing to pay only $10 to avoid a risk of 1 in 100,000. But it is not at all clear that they are being required to purchase more safety than they want. Because the people who are paying to reduce the risk are relatively wealthy, construction workers might be required to pay less than $10 – actually $0 – to avoid the risk of 1 in 100,000. If so, they will gain.[16]

2.5. Problem 5

The U.S. government is regulating air pollution to diminish risks, including mortality risks. Most of the people who bear the costs of the regulation are poor; some are wealthy. For poor people, the average VSL is $1 million. For wealthy people, the average VSL is $20 million. The U.S. government uses a population-wide average, which is $10 million.

Use of the population-wide average is unlikely to be in the interest of poor people. They are willing to pay only $10 to avoid a risk of 1 in 100,000, and it appears quite possible that they are being required to purchase more safety than they want. Because most of the people who are paying to reduce the risk are poor, poor people might be required to pay more than $10 to avoid risk of 1 in 100,000. If so, they will lose.

As in Problem 3: These conclusions are tentative, because we do not yet know about the incidence of the benefits. Suppose that the benefits are concentrated among poor people, so that they enjoy the overwhelming majority of the mortality and morbidity gains. We would have to obtain a host of numbers to know, but it is possible that poor people are net gainers, even with a VSL of $10 million. Assume, for example, that they pay 51 per cent of the total cost (which is, let us suppose, $100 million), but receive 98 per cent of the total benefit (which is, let us suppose, $2 billion). If so, they are likely to be gainers. We could imagine other assumptions that would lead to the conclusion that poor people are net losers.

---

[16] I am assuming that there is no other loss for construction workers from the regulation, such as loss of employment, wage reductions, or reductions in benefits provided by employers.

3. Welfare

My focus has been on whether use of a uniform VSL is good for poor people, but that is not, of course, the only question. A VSL could be good for poor people but bad for everyone else. If so, what should be done? One question is how good it is, exactly, for poor people, and how bad it is, exactly, for everyone else. Another question is the size of the relevant populations – how many people are poor, and how many people are not poor? If a policy makes life a great deal better for a large number of poor people, and a tiny bit worse for a small number of rich people, it would seem to be an excellent idea.[17] If a policy makes life a tiny bit better for a very small number of poor people, and a great deal worse for a very large number of rich people, it would not seem to be an excellent idea. Harder cases would involve different magnitudes of effects and different sizes of populations. Suppose, for example, that a policy made life nontrivially better for a significant number of poor people, but also made life nontrivially worse for a significant number of people who are not poor. We would need to have more details to know whether the policy is producing net welfare gains.

Purely on welfare grounds, we should now have something like a general framework with which to approach a host of problems. One implication is that regulations that increase health and safety, or that promote environmental protection, may or may not be more appealing if we include "distributional weights"; a uniform VSL can be seen as a kind of distributional weight, and on certain assumptions, it might hurt rather than help poor people. Another implication is that a uniform VSL, combined with distributional weights for benefits that accrue to poor people, might seem appealing on normative grounds, but might well produce terrible outcomes for poor people.[18]

There are also questions about distribution as such. Suppose (i) there is a net loss in welfare and (ii) poor people gain a great deal and people who are not poor lose slightly more than poor people gain. We might favor that policy on grounds of distributional justice. To know whether we should, we might want to know the magnitude of the loss in welfare; the magnitude of the gain to poor people; and the magnitude of the loss to people who are not poor. And if we are prioritarians, we will give special attention to those at the bottom of the ladder, and accept overall losses if they are the price we have to pay in order to help that group.[19]

References

Adler, Matthew D. 2016. "Benefit-Cost Analysis and Distributional Weights: An Overview." *Review of Environmental Economics and Policy*, *10*(2): 264–285.

---

[17] I am bracketing possible incentive effects: A program that helps the poor, and hurts the rich, might decrease the incentive to be rich rather than poor. But especially insofar as we are speaking of funding programs for (for example) climate-related risks, any such incentive effects are likely to be modest or nonexistent.
[18] See Hemel (2022).
[19] See Adler and Norheim (2022).

Adler, Matthew D. and Ole F. Norheim. 2022. *Prioritarianism in Practice*. New York: Cambridge University Press.

Department of Homeland Security. 2021. Best Practices for the Treatment of a Statistical Life in U.S. Department of Homeland Security Regulatory Analyses. Office of the General Counsel. Available at https://www.dco.uscg.mil/Portals/9/VSL%20Final%20for%20DIST%205April21.pdf (accessed January 17, 2023).

Department of Transportation. 2021. Departmental Guidance on Valuation of a Statistical Life in Economic Analysis. Available at https://www.transportation.gov/office-policy/transportation-policy/revised-departmental-guidance-on-valuation-of-a-statistical-life-in-economic-analysis (accessed January 17, 2023).

Dorfman, Robert. 1977. "Incidence of the Benefits and Costs of Environmental Programs." *American Economic Review*, *67*(1): 333–340.

Gianessi, Leonard P., Henry M. Peskin, and Edward Wolff. 1979. "The Distributional Effects of Uniform Air Pollution Policy in the United States." *Quarterly Journal of Economics*, *93*(2): 281–301.

Greenberg, Kyle, Michael Greenstone, Stephen P. Ryan, and Michael Yankovich. 2021. "The Heterogeneous Value of a Statistical Life: Evidence from U.S. Army Reenlistment Decisions." NBER Working Paper 29104. Cambridge, MA: National Bureau of Economic Research. Lasted revised October 2021. Available at https://www.nber.org/papers/w29104 (accessed January 17, 2023).

Harrison Jr, David. 1975. *Who Pays for Clean Air: The Cost and Benefit Distribution of Federal Automobile Emission Standards*. Cambridge, MA: Ballinger Publishing Company.

Hemel, Daniel. 2022. "Regulation and Redistribution with Lives in the Balance." *University of Chicago Law Review*, *89*(3): 649–734.

Kahn, Matthew E. 2001. "The Beneficiaries of Clean Air Act Regulation." *Regulation*, *24*(1): 34–38.

Kaplow, Louis and Steven Shavell. 1994. "Why the Legal System Is Less Efficient than the Income Tax in Redistributing Income." *Journal of Legal Studies*, *23*(2): 667–681.

Liscow, Zachary. 2022. "Redistribution for Realists." *Iowa Law Review*, *107*(2): 495–561.

McGee, Kelly. 2021. "A Place Worth Protecting: Rethinking Cost-Benefit Analysis Under FEMA's Flood-Mitigation Programs." *University of Chicago Law Review*, *88*(8): 1925–1970.

Revesz, Richard L. 2018. "Regulation and Distribution." *New York University Law Review*, *93*(6): 1489–1578.

Sunstein, Cass R. 2007. "Willingness to Pay Versus Welfare." *Harvard Law and Policy Review*, *1*: 303–330.

Sunstein, Cass R. 2013. "The Value of a Statistical Life: Some Clarifications and Puzzles." *Journal of Benefit-Cost Analysis*, *4*(2): 237–261.

Viscusi, W. Kip. 2010. "The Heterogeneity of the Value of Statistical life: Introduction and Overview." *Journal of Risk and Uncertainty*, *40*: 1–13.

Viscusi, W. Kip. 2013. "The Benefits of Mortality Risk Reduction: Happiness Surveys vs. The Value of a Statistical Life." *Duke Law Journal*, *62*(8): 1735–1745.

Viscusi, W. Kip. 2018. *Pricing Lives*. Princeton: Princeton University Press.