

Euthanizing the Value of a Statistical Life

Trudy Ann Cameron
Department of Economics
University of Oregon
Eugene, OR 97403-1285
cameron@uoregon.edu

*This is an expanded version of a short article that I wrote for the November 2008 Newsletter of the Association of Environmental and Resource Economists (AERE). I am grateful to Rob Stavins for encouraging me to expand upon the points made there and to seek a broader audience for these ideas, and to Lisa Robinson, Kerry Smith, Kip Viscusi, and Richard Zerbe for helpful comments and suggestions. Charlie Kolstad suggested the best title, and Suzanne Leonard's editorial expertise made great improvements to the exposition.

Abstract: For economists, the term “value of a statistical life” (*VSL*) is an eminently reasonable label for the concept it describes. However, outside our discipline, this terminology has been singularly unhelpful. This article argues that there could be a considerable reduction in wasted resources if economists were to change this terminology to something less incendiary, and that this could help to increase the acceptance of benefit-cost analysis as an input to the decision-making process for environmental, health and safety regulations. I propose that we change our standard unit of measurement and replace the *VSL* terminology with “willingness to swap (*WTS*) alternative goods and services for a microrisk reduction in the chance of sudden death”. Analogous terminology would be used for other types of risks to life and health. I also argue that economists’ continual pursuit of a single number for “the” *VSL* is misguided and can be misleading, especially if individual *WTS* is correlated with the magnitudes of the risk changes. Such “one-size-fits-all” *VSLs* also hinder our ability to understand the distributional consequences of risk-reducing policies or interventions. Estimates of aggregate risk reduction benefits need to reflect the particular type of risk reduction as well as the characteristics of the affected populations.

JEL classification numbers: Q51, I18, D61

Keywords: Value of a statistical life, *VSL*, mortality risk, environmental health benefits

Euthanizing the Value of a Statistical Life

[Vizzini has just cut the rope that *The Dread Pirate Roberts* has been climbing.]

Vizzini: He didn't fall? INCONCEIVABLE!

Inigo Montoya: You keep using that word. I do not think it means what you think it means.

- motion picture *The Princess Bride* (1987)

Introduction

There is a wide variety of policies and regulations aimed at protecting human life and health. To analyze the benefits and costs of these measures, we must address society's willingness to pay for a variety of health risk reductions. By convention, economists report estimates of this willingness to pay (*WTP*) in terms of a unit called a "statistical life."¹

While some economists may not be particularly familiar with this terminology, they catch on quickly when the "value of a statistical life" (*VSL*) is explained as a marginal rate of substitution between mortality risk and money (i.e. other goods and services). The *VSL* is a ratio in which the numerator is the marginal utility of a small reduction in mortality risk, r , and the denominator is the marginal utility of a small change in income, Y .² While the *WTP* for an incremental risk change will be small, the numerical value of this *ratio* is very large. For mortality risks (typically the risk of sudden death in the current period), empirical data on the tradeoffs that real people are willing to make often indicate a middle-of-the-road estimate of around \$7,000,000.³

¹ A statistical life is a unit defined as a cross-sectional aggregate of enough individual tiny fractional risk reductions so that the total risk reduction is 1.0.

² This is expressed mathematically as $(\partial U / \partial r) / (\partial U / \partial Y)$, which simplifies to $\partial Y / \partial r$, measured in "dollars per unit of risk reduced."

³ The EPA's current recommendation for the Value of a Statistical Life is that "the central estimate of \$7.4 million (\$2006), updated to the base year of the analysis, be used in all benefits analyses that seek to quantify mortality risk

Although economists understand the *VSL* terminology, non-economists often have difficulty making sense of this concept. The most recent nationwide explosion of indignation over the *VSL* was “detonated” by a July 10, 2008 Associated Press article by Seth Borenstein, titled “An American life worth less today.” The article described a decision by the U.S. Environmental Protection Agency (EPA) to revise downward its standardized estimate of society’s willingness to pay for aggregate mortality risk reductions from environmental policies (i.e., its estimate of “the *VSL*”) from \$7.8 million to \$6.9 million.⁴ While Borenstein interviewed a number of key experts both inside and outside of the EPA, his description of the issues and the data upon which *VSL* calculations are based was oversimplified, forcing readers to invent their own interpretations. And invent, they did.

In the days following publication of the article, which was picked up by local newspapers throughout the country, I collected a large number of comments from a variety of sources, including public comments on news websites. These reactions to the Borenstein article can leave no doubt that the consumers of our research need to be steered away from the mistaken impression that know-nothing bureaucrats presume to decide on behalf of society the intrinsic worth of a human being. The comments also reflect the damage to economists’ reputations from the public’s misinterpretation of the concept of a *VSL*. More importantly, however, these comments suggest that a significant amount of reputational capital and person-hours at various

reduction benefits regardless of the age, income, or other population characteristics of the affected population.” <http://yosemite.epa.gov/ee/epa/eed.nsf/webpages/MortalityRiskValuation.html#currentvsl>. A very thorough and detailed description and history of the EPA’s *VSL* numbers and recent controversies is provided in Viscusi (2009a). See also comments by Carruthers (2009), Fourcade (2009) and Robinson (2009), and a reply by Viscusi (2009b).

⁴ In fact, different offices at EPA use different *VSL* estimates for different policies. To update its numbers to reflect more recent research on the *VSL*, producing the lower *VSL* identified in the Borenstein article, the EPA Air Office apparently “selected as its preferred *VSL* the midpoint of the 25th percentile of the estimates in Mrozek and Taylor (2002) and the 75th percentile of Viscusi and Aldy (2003). This unusual mathematical formulation creates the illusion of precision but lacks any scientific basis.” (Viscusi, 2009a, p. 116).

government agencies is wasted every time another sensational article about the *VSL* appears in the press.

My purpose in writing this article is to discuss in detail what is wrong with the “*VSL*” terminology and economists’ conventional notion of a *VSL*, and to propose some alternative terminology that is both less provocative and more precise. I begin in the next section with a discussion of public misconceptions about “the *VSL*” and then turn to problems with the *VSL* terminology itself. I will contend that the public relations blunder that is “the *VSL*” is largely of our own making, that our profession’s steadfast adherence to the term “value of a statistical life” was as much responsible for the outrage that followed publication of the Borenstein article as was the author’s over-simplified description of the EPA decision. I will also argue that we have been too cavalier with specialized terminology, summarizing demand information in the aggregate, rather than in individual units with which people could more readily identify, and failing to realize that many people have fundamental problems with the idea of “monetization.”

In the following section, I propose what I view as a viable alternative to “the *VSL*”, one that changes our standard unit of measurement and replaces the “value of a statistical life” with “willingness to swap” (WTS) alternative goods and services for a microrisk reduction in the chance of sudden death (or other types of risks to life and health). Next I take up the issue of the aggregation of risks, which has led to our profession’s misguided and misleading pursuit of a single, one-size-fits-all *VSL* and which precludes any assessment of the distributional consequences of risk-reduction policies based on their impacts on individual *net* benefits. I conclude with a discussion of the steps that need to be taken to “euthanize” the *VSL* terminology.

Misconceptions and Criticism of the *VSL*

Attempts at political damage control quickly followed publication of the Borenstein article. The day after the article appeared, the following remarks were attributed to the office of Senator Barbara Boxer (D-CA):

“The EPA’s decision to reduce the value of a human life when they consider the benefits of new environmental regulations is outrageous and must be reversed.” ... “EPA may not think Americans are worth all that much, but the rest of us believe the value of an American life to our families, our communities, our workplaces and our nation is no less than it has ever been. This new math has got to go.” ... “If these reports are confirmed, I will be introducing legislation to reverse this unconscionable decision at the earliest opportunity.”⁵

About a week later, a follow-up article (Fahrenthold, 2008) appeared on the front page of the *Washington Post*, containing the following ideas:

“Someplace else, people might tell you that human life is priceless. In Washington, the federal government has appraised it like a ‘96 Camaro with bad brakes....” “...the Washington bureaucracy takes on a question usually left to preachers and poets.” “...for the first time, the EPA has used this little-known process to devalue life....” ... “By reducing the value of human life, which is really a devious way of cooking the books,...” “...To grasp the mind-bending concept of a Blue Book value on life, government officials say it is important to remember that they are not thinking about anyone specific....” “...an unlikely academic field has grown up to extrapolate life’s value from the everyday decisions of average Americans.” “...lowering the value of life. In some bureaucratic corners of Washington, it is the kind of phrase that nobody blinks at anymore.”

Ordinary people also vented their indignation in a variety of ways, including letters to the editors of newspapers, comments on news media web sites, and even in television news satire.⁶

The wide array of viewpoints, objections, and emotions is striking, as is the number of ways in which the idea of a *VSL* can be misinterpreted. Some people were clearly struggling to

⁵ Naively, I wrote immediately to Senator Boxer’s office using the online link at her official website. I carefully explained why this issue was something her office might want to be sure they understood better, before she introduced new legislation. However, after hitting the “send” button, I was informed automatically that since I was not a resident of California, Senator Boxer wouldn’t be paying any attention to my comments. Viscusi (2009a) describes Senator Boxer’s proposed legislation, the “Restoring the Value of Every American in Environmental Decisions Act” (110th Congress, 2d Session, 2008).

⁶ The on-line supplementary materials for this article present a sample of some 90 public comments. [PUBLISHER- PLEASE INSERT ON-LINE LINK]

understand, while others were completely dismissive. The comments ranged from pious to political, from witty to sarcastic, and from sad to downright angry. Most importantly, however, the comments revealed widespread confusion and misconceptions about the *VSL* concept.

This is not the first time there has been a misinformed public maelstrom over the *VSL*. For example, back in 2003, it was revealed in the press that the EPA had explored the idea of using a *VSL* estimate for seniors that was about one-third lower than for other adults. Then-Administrator of the EPA Christine Whitman faced public outrage about the so-called “senior death discount.” Environmental groups mobilized AARP (formerly the American Association of Retired Persons) to protest the outrageous notion that America’s seniors should be so callously “devalued.”⁷ Whitman did not stay long at the helm of the EPA after this controversy.

The protest over the senior death discount seemed to have as its rallying cry: “We seniors absolutely insist on paying just as much as everyone else!” Having been intrigued for some time by the fervor of this protest, I stopped by the AARP exhibit booth at the meetings of the Agricultural and Applied Economics Association (AAEA) in Orlando, FL, in July 2008, just a few weeks after publication of Borenstein’s AP story. I asked the AARP representative if she knew the definition of the value of a statistical life. She had heard the expression, and knew that the government had tried to put a lower value on the lives of seniors and that the EPA had recently been taken to task for having quietly lowered the overall value of human lives in recent years. However, she couldn’t point me to any AARP literature explaining to its membership the concept of the value of a statistical life and its implications. She was unaware of the Laughland et al. (2007) paper, which explains for a more general audience the role of benefit-cost analysis in regulatory impact analysis, including the implications of using age in the calculations. This

⁷ This mobilization may have been somewhat self-serving. Environmental groups seek to maximize environmental quality, while the AARP seeks to maximize the welfare of seniors. The two different objective functions may not always be fully compatible.

paper was prepared for AARP, and yet it had not filtered down to the front lines, let alone to the general membership of the organization.

A Poor Choice of Words

The term “VSL” highlights a downside to the economist’s habit of eschewing new jargon. This tendency means that we often build technical labels out of commonly used words that we happen to have “lying around the house.” For example, anyone who has taught Principles of Economics knows how much time must be spent explaining exactly what *an economist* means by the words “cost” or “profit.” These words do not have the same meaning for economists as for the person on the street. Non-economists often make no distinction between “cost” and “price,” for example. And when ordinary people use the word “profit,” they typically have in mind accounting profit, rather than economic profit.

By cobbling together the terms “value” and “statistical” and “life,” we have done a grave disservice to both our profession and the policymakers we are often trying to help via our research and analysis. Let’s consider, individually, each of the three main words in “value of a statistical life.”

“*Value*”: By “value,” economists mean willingness-to-pay (or sometimes, willingness-to-accept). By using the term “value” to mean willingness-to-pay for risk reductions, we have created considerable confusion in the public’s mind because the public often equates “value” with “intrinsic worth.”

“*Statistical*”: Economists probably intended the term “statistical” to indicate probabilistic outcomes rather than death with certainty. When summed over enough people, small probabilities of death will yield the expectation of one death in that large group. The trouble is that many intelligent people don’t remember much of what they might have once learned about

statistics. In fact, their knowledge may be based on one of the most widely purchased introductions to statistics for the general reader: Darrell Huff's (1954) "How to Lie with Statistics".⁸ Or perhaps they've heard the catchy quote attributed to Benjamin Disraeli, and popularized in the U.S. by Mark Twain (1907), that mentions "lies, damned lies, and statistics." By introducing the adjective "statistical" into our terminology, we may have courted instant skepticism and distrust of the *VSL*.

"Life": Non-economists think we are valuing one whole, distinct, individual, and identifiable "life", when we are actually seeking to value tiny risk reductions for many different people. It was unwise to decide that the units in which to quote willingness to pay for mortality risk reductions should be aggregated across enough different people so that the sum of their individual risk reductions equals 1.0. It would have been equally easy to quote willingness to pay in terms of the size of the risk reduction that is relevant for the average individual under the policy in question, which would often be some relatively modest number of dollars. Instead, we quote willingness to pay in terms of an arbitrarily huge aggregate risk reduction, which many people confuse with one specific life.

A Poor Choice of Units

Quoting the value of risk reductions in units of statistical lives is akin to quoting the price of milk by the tanker-truck load, rather than by the quart. Most people have a sense of what they might be willing to pay for a quart of milk, at least within an order of magnitude (say \$2.00). Moreover, they are likely to be relatively familiar (and comfortable) with contemplating their personal tradeoffs concerning the decision to buy a quart of milk. However, few people would

⁸ See a discussion in Steele (2005). Huff's book includes such provocative chapter headings as "The Sample with the Built-in Bias," "The Little Figures That Are Not There," and "How to Statisticulate." The author himself acknowledges in the introduction that "It may seem altogether too much like a manual for swindlers." As a result, many readers may be left with the impression that all statistics are purposely designed to mislead.

react well to a requirement that they instead contemplate the tradeoffs involved with purchasing a \$40,000 truckload of milk (especially if it has not been carefully explained that a milk tanker holds 5,000 *gallons*, but that nobody actually expects one individual to buy the whole truckload). This is exactly the same per-unit price as for the quart of milk—just scaled up to a vastly larger quantity than any family would buy on a trip to the grocery store. Any shopper who encountered a price for milk at the supermarket of \$40,000 (per 5,000 gallons) would likely be as much dismayed, befuddled, or outraged as many of the people who contemplate \$7 million for the value of a statistical life.

Ordinary citizens need to be able to consider, dispassionately and objectively, the reasonableness of the evidence on willingness to pay for mortality risk reductions. To help them do this, our evidence about willingness to pay should be reported in a way that is easier to grasp than the concept of “statistical lives” (e.g., in terms of the sizes of risk reductions that are relevant for a single individual or household). To trained economists, the choice to use “statistical lives” as the units for risk reductions may have seemed inconsequential. However, it appears to have been a very bad choice from the perspective of consumers.

The concept of “monetization”

Although it would clearly be an improvement to quote “willingness to pay for risk reductions” in terms of individual risk reductions instead of the “value of a statistical life,” the problems of the terminology go even deeper. As economists, we know precisely what we mean by the term “willingness to pay.” It is the monetized value of the other goods and services that people would be willing to give up in order to get one more unit of the thing in question. As economists, we have also been carefully taught to understand that it is the *real* tradeoffs that matter, and that money merely facilitates transactions. Thus, we need to know what people are

willing to give up in terms of the next best use of the resources that would have to be reallocated to provide or create something (like better health or a safer workplace).

When it comes to health risk reductions, the consumer might have to give up apples or oranges or shoes or cars or any combination of things that would otherwise provide utility if the health risk reduction in question were not chosen. Because we can't add apples and oranges, we use prevailing prices and convert these other things to expenditures. What the consumer is willing to swap to obtain a health risk reduction is measured by the "dollars' worth" of other goods and services he or she would be prepared to do without. This is equivalent to a reduction in income. In our usual fashion, we economists have been stingy with verbiage and describe this process merely as "willingness to pay" for a health risk reduction—or even more briefly as just "monetization."

It is a subtle but very important point that health, itself, is not being directly or arbitrarily monetized. Instead, monetization is the strategy used to convert to a common denominator all those other heterogeneous goods or services willingly forgone to obtain the health improvement. We measure the "amount" of these other things via the expenditure that would be necessary to acquire them, given current prices. In other words, we measure the value of the health improvement to the individual in terms of the "dollars' worth" of other goods and services a consumer would be willing to forgo in exchange for the health improvement.

Unfortunately, we have done too little to educate the general population about what our specialized terms—*VSL*, willingness to pay, and monetization—actually mean. Equally unfortunate is that these terms carry with them a lot of cultural baggage. Economists are able, implicitly, to separate the baggage from the technical meaning, but the general public often cannot.

Cultural baggage and the public's fundamental discomfort with monetization

Americans are skittish about the topic of money. It is one of those things that we are not supposed to discuss in polite company. Why is the topic of money viewed as unsavory? Certain biblical phrases and juxtapositions have had great staying power in our popular culture. For example, the term *filthy lucre* is defined as “money; originally, money obtained dishonestly” (Ammer, 197, p. 207). The Latin word *lucrum*, from which *lucre* is derived, means “material gain, profit” or “avarice.”

There are a variety of translations of the biblical passage in the New Testament that contains “filthy lucre”.⁹ One is “Whose mouths must be stopped, who subvert whole houses, teaching things which they ought not, for filthy lucre's sake.” This particular verse is sometimes summarized as “False teachers who must be stopped.” This is not an auspicious cultural vantage point from which academic economists should be asking people to entertain the concept of “monetization” of statistical lives.

There is also a huge cultural taboo against the notion of “willingness to pay” when the lives of human beings are at stake. There appears to be a short cognitive distance between the “value of a statistical life” and “paying for the life of a human being” or “slavery.” All the repugnant moral connotations of slavery are readily transferred to the idea of paying for a human life. Furthermore, we ask individuals to consider paying for their *own* lives, rather than someone else's! This furthers the sense of moral outrage because people presumably feel an even greater sense of entitlement to their *own* lives.

Finally, consider the anxiety and stress engendered by the traditional “stick-‘em-up” ultimatum: “Your money or your life!” Only a dangerous criminal would propose such a

⁹ The website <http://bible.cc/titus/1-11.htm> appears to be a good source for a variety of parallel translations of the passage in question, Titus 1:11.

choice, and one would have to be looking down the barrel of a gun to consider it. The victim in question is being asked to make a tradeoff, albeit in an increment that is vastly larger than any of those relevant for any normal policy analysis. However, these dramatic and extreme instances of tradeoffs between money and risks to life make an indelible impression on people's psyches and add to their fears.

Rebranding the VSL

Savvy marketers have always known that what you call your product will influence people's demand for it. This phenomenon is reflected in the fact that companies will sometimes legally register a clever trademark (i.e. a name, symbol or other device that identifies a product) before the product itself has even been developed.¹⁰

In other cases, companies are taken to task for giving a product a name that makes it appear more desirable or effective than can actually be proven. Consider, for example, a homeopathic product for the seasonal flu virus that was once called "Flu ResistTM". Several years ago, this product was quietly renamed "Flu ReliefTM", with no other change in its packaging. Most people would probably pay more to *prevent* (or "resist") influenza than to treat it, so the relabeling presumably reflects an adjustment away from a name that exaggerated the desirability of the product, towards one that more accurately reflected its clinically proven efficacy.

Just as the labeling of consumer products is an important determinant of buy-in by the public, so too is the labeling of public policies. During the second Bush administration we were presented with the Clear Skies Initiative and the Healthy Forests Initiative. Who could argue with clear skies or healthy forests? The policies themselves were multi-faceted and would involve a variety of reallocations of resources relative to the status quo, so as usual there would

¹⁰ Of course, there is an expectation that a good-faith effort to develop the intended product is under way.

be some winners and some losers. However, these policy labels were easy to enunciate and to remember, and both had indisputably positive connotations, even though the labels were rather imprecise descriptions of the policies themselves.

Economists have proven themselves to be appallingly poor marketers when it comes to the “Value of a Statistical Life” or “*VSL*.” A vowel would certainly have been helpful, so that the acronym could have been pronounceable as a word. But at least the concept can be reduced to a three-letter abbreviation. However, compared to “Clear Skies” or “Healthy Forests,” our “*VSL*” label is a marketing disaster. Surely, we can do better. We need a label for this concept that is brief, accurate, and at least perceived as being neutral and logical, if not actually appealing, by the non-economist public.¹¹

While you can’t fix a problem simply by relabeling a product as something it is not, it is entirely appropriate to correct a name to more accurately portray what the product is or does. I am reminded of trying to get an online survey instrument through the Institutional Review Board (IRB) at one of my former universities. As an incentive for participation, I was offering each respondent the chance to win a prize for completing the survey. I described this prize competition as a lottery. The IRB contacted the campus legal counsel about this plan and I was immediately notified that what I proposed was illegal, since only the state could sponsor a lottery. Realizing that the legal definition of a lottery specified that people must pay money to participate, rather than simply devoting their time (another scarce resource) to complete a survey, I changed the wording to specify that the respondent would be “entered to win a prize.” The university counsel’s office initially refused to accept this remedy because “just calling it something other than a lottery does not make it legal.” It took more effort than should have been

¹¹ As a product name, *VSL* seems to be right up there with the classic example of the Chevy Nova (which one would expect might be more difficult to market in Mexico—“no va”), or the appetite suppressant candy called *Ayds* that was popular prior to the early 1980s.

necessary to convey the point that the prize competition was never a “lottery” in the first place, in the legal sense, and that just *calling* it a lottery should not make it *illegal*.

I use this anecdote to illustrate the point that even if we attempt to improve our terminology for the value of a statistical life, any one of the unhappy people who commented on the Borenstein article might leap to the conclusion that we are simply “putting lipstick on a pig.” On the contrary, our goal is to prevent the value of a statistical life from being mistaken for a bureaucratic attempt to dictate arbitrarily the worth of a human being. Relabeling the concept more carefully, to describe precisely what it actually means, is entirely appropriate and, in my view, absolutely necessary.

The concept of a micromort

Over two decades ago, Howard (1984, 1989) advocated the concept of a micromort (based on “micro” for millionth and “mort” for mortality) to describe a one-in-a-million chance of death. In most of the literature on the economic valuation of mortality risk reductions, it is assumed that *WTP* is proportional to the size of the risk reduction (as it will be, if the marginal willingness to pay, *MWTP*, is constant). If this proportionality assumption holds, then we are free to scale the dollar amounts of *WTP* to correspond to *any* arbitrarily-sized risk change. The *VSL*, as a ratio that indicates a marginal rate of substitution, is equivalent to a scaling of the monetary value for a modest risk change to a huge 1.00 (or 100%) change in mortality risk. In contrast, willingness to pay for a micromort would scale the monetary value to a tiny 0.000001 change in the risk of death.

It makes sense that the units we use to measure risk changes (and to value them) should be within a few orders of magnitude of the typical sizes of risk reductions contemplated under proposed environmental policies or regulations. Thus it would make sense to quote *WTP* for

mortality risk reductions in 1/1,000,000 units, for example, rather than in 1.00 units (as the *VSL* does). This change in units provides information that is identical to the *VSL*, but with much less potential for misinterpretation by non-economists. Remember that we have absolutely no reliable large samples of data about how much people would pay to reduce their own risk of death from 1.00 to 0.00. Yet this is exactly what is implied by the choice to scale *WTP* to a 100% mortality risk reduction. Our current approach is to observe tradeoffs with respect to small risks, scale them way up to *VSLs* to report them, then scale them right back down again for use with individual risks of sizes that are relevant for policy evaluation. I suggest that we skip the middle step. It is completely unnecessary and is entirely to blame for a whole lot of confusion.

Microrisks as a more appropriate unit

There is considerable merit to the notion that we should separate our primary label for the *size* of the risk reduction in question from a secondary label for the *type* of risk to be reduced. We need to be explicit about both size and type.

If we can pry ourselves away from “statistical lives”, even the concept of “*WTP* for a micromort” may not be sufficiently general, since it would apply only to *mortality* risk reductions. Environmental policies and regulations also reduce the risk of *morbidity* (sickness). One might consider an analogous *WTP* for a “micromorb” for morbidity risk reductions, but unfortunately this would be rather difficult to enunciate. Howard (1984) has also suggested other nomenclature, such as a “microdisability” for risks that lead to disability rather than death.

Our terminology needs to be broad enough to apply to a wide variety of health risk reductions. The more-specific terms mentioned above could be introduced for special cases of morbidity or disability risks. However, Howard’s most generic term is a “microrisk.” Thus we could refer succinctly to a 1/1,000,000 risk reduction as a *microrisk* (μr) reduction for the risk

in question. A 1/1000 risk reduction could be called a *millirisk* (mr) reduction.¹² We could use multiples in between these two terms, if necessary. Mortality rates are often quoted as “per 100,000.” Thus a 1/100,000 risk reduction would be “ten microrisks” ($10\mu r$) and a 1/10,000 risk reduction would be “100 microrisks ($100\mu r$).¹³ Although “WTP for a microrisk reduction” is based on more-standard scientific terminology, it certainly does not roll off the tongue as easily as “VSL.” However, it is less misleading and more precise, and would probably be much less confusing to the public.¹⁴

It is also important to be explicit about the research context that has provided a particular WTP estimate and whether this matches the policy context to which it will be applied.

Willingness to pay for any commodity will depend upon the precise attributes of the commodity in question, as should willingness to pay for reductions in risks to life and health. Thus it is important to include risk attributes for an estimate of WTP for a risk reduction of a given size, for example “a \$10.8 WTP for a microrisk reduction, based on wage-risk tradeoffs by middle-aged males concerning on-the-job fatalities in the current period.” One might easily contend that we should have been using this level of description all along, even with the VSL estimates already in the literature.

Avoiding the idea of “payment”

As discussed above, non-economists are reluctant to accept the technically accurate economic term known as “willingness to pay,” presumably because they are uncomfortable with

¹² The abbreviation adheres to the conventions of the National Institute of Standards and Technology for international System of Units (SI) prefixes for multiples and fractions of units.

¹³ Ten times something is “deka-“ (da) and 100 times something is “hecto-“ (h). So a 1/100,000 risk reduction (ten times a microrisk) could be called a *dekamicrorisk* ($da\mu r$) reduction, and a 1/10,000 risk reduction (one hundred times a microrisk) could be called a *hectomicrorisk* ($h\mu r$). However, this seems like overkill in terms of jargon.

¹⁴ In contexts where one did not wish to digress to explain the definition of a microrisk reduction, one could simply say “a one-in-one-million risk reduction,” or “a 0.000001 risk reduction” if sufficient numeracy can be assumed.

the notion of paying money for their lives or their health. However, policy choice concerns people's willingness to make *tradeoffs*. That is, to gain improvements in their health or longevity, they may have to settle for higher consumer prices, lower wages, lower investment returns, or even higher taxes. These tradeoffs are likely to be unpleasant because they force the individual to do without something else that he or she would otherwise have been able to afford.

Given the adverse reactions of the general public to the notion of being asked to pay for risk reductions, we may need to eliminate the word "pay" altogether. A more accurate term would be "forgo," since it conveys the sense of a tradeoff (i.e., that something else will have to be given up or done without). Forgone opportunities are thoroughly understood by economists and could likely be understood by the general public as well. However, "willingness to forgo" does not lend itself to a helpful acronym.

Proposed replacement terminology

So we need a synonym for "forgo" that conveys the same sense of "to give up in exchange for." One candidate might be "to relinquish." Another is "to sacrifice." But perhaps the most neutral verb, and the one with the fewest alternative definitions, is "to swap", which means "to trade one thing for another" or "to exchange (one thing) for another."¹⁵ "Swap" was originally a horse trader's term, being a shortened version of the phrase "swap a bargain." The word *swap* evolved from the Middle English word *swappen*: "to strike, strike hands (in bargaining),"¹⁶ and from the Gaelic word *suaip*: "to exchange, to barter" (Mackay, 1877, p. 444). The notion of striking hands (which has evolved to a custom of shaking hands) to signify agreement to a willing trade or contract seems very apt for conveying the notion of willingness to

¹⁵ swap. (n.d.). *The American Heritage® Dictionary of the English Language, Fourth Edition*. Retrieved October 22, 2009, from Dictionary.com website: <http://dictionary.reference.com/browse/swap>

¹⁶ swap. (n.d.). *Dictionary.com Unabridged*. Random House, Inc. Retrieved October 22, 2009, from Dictionary.com website: <http://dictionary.reference.com/browse/swap>

pay in contexts where money as a medium of exchange may not be a necessary element of the bargain.

Fortunately, the abbreviation “WTS” also seems to be relatively neutral.¹⁷ There may be some initial confusion for economists concerning the meaning of the acronym, since “S” commonly stands for “supply” or “substitution” or “savings.” But we, as economists, are perhaps better equipped to figure out the WTS acronym in the relevant context than the general public has been at interpreting the “value of a statistical life.”

Willingness-to-swap for a microrisk reduction will not be \$3.14159265358973...

Economists also need to push policy-makers to acknowledge that “the VSL” is not some true-but-unknown fundamental constant of nature that we merely need to measure more accurately. Instead, “the VSL” is the result of attempts to find a convenient one-size-fits-all measure of demand for risk reductions—a number which may or may not be appropriate across all different types of risks or all different affected populations. A single universal VSL number is politically expedient, of course, but it is also likely to be the *wrong* number, unless the risk is “typical” and it affects a population with “typical” demands for protection. I discuss this issue in more detail in the next section.

The Perils of a One-Size-Fits-All VSL

Why do we think we need “the VSL”? Policy evaluators have sought a one-size-fits-all measure of willingness to pay for risk reductions in the form of “the VSL” because in their assessment of social benefits of risk reductions they have traditionally jumped the gun and aggregated across different people too soon in the valuation process.

¹⁷ As an acronym, WTS seems to have as its closest competition the Westminster Theological Seminary.

Premature Aggregation

Ideally, our calculation of the overall benefits of risk reductions should remain disaggregated to the individual level through many steps, starting with the identification of individual risk reductions, through multiplication by the corresponding individual *marginal* willingness to pay amount, to produce individual willingness-to-pay for the particular size of risk reduction in question. Individual marginal costs should also be identified, of course, to permit the calculation of individual *net* benefits. Only then, after the distributional consequences of this pattern of net benefits have been examined, should one aggregate individual net benefits across the affected population to yield an initial (equally weighted) measure of overall net social benefits.¹⁸

In practice, benefit-cost analysts have tended to identify individual physical risk reductions and then jump straight into aggregation. They first sum individual risk reductions across the affected population, which may be a very large number of people. When the sum of individual risk reductions reaches 1.00, one “statistical life” has been accumulated. For the policy in question, the quantity of interest is then the number of “statistical lives” to be “saved.” Of course, no specific life will be saved with certainty by the policy. The actual risk reductions could be on the order of one in ten thousand or even one in a million for each person, which could be viewed by many people as vanishingly small.

Once analysts have calculated the number of statistical lives involved, they need to convert this aggregate risk reduction into an aggregate willingness to pay for this public good. Since the units of risk reduction have been established as “statistical lives,” a corresponding dollar value “per statistical life” is required. This dollar measure cannot reflect individual

¹⁸ Overall net social benefits are expressed as: $\sum_i (MWTP_i - MC_i) \Delta r_i$, where $MWTP_i$ = individual marginal willingness to pay, MC_i = individual marginal costs, and Δr_i = individual risk reductions.

preferences or constraints, because all information about individual heterogeneity has been lost due to premature aggregation, across people, of the individual risk reductions associated with the policy. Aggregate *physical* benefits are measured by the number of statistical lives, and thus a measure of the “social” willingness to pay for those statistical lives is needed to monetize these benefits. This accounts for the persistent desire for a single one-size-fits-all *VSL*—to permit calculation of social benefits by multiplying just two terms: $(\sum_i \Delta r_i) \times VSL$ (where the first term is the aggregate of the individual risk reductions).

Aggregation of risk and then multiplication by the *VSL* may involve minimal bias if everyone faces the same risk reduction and everyone’s willingness to pay for risk reductions is identical. But the risk reductions promised by a policy may be distributed very unevenly across the population. Likewise, individual marginal *WTP* for risk reductions may vary widely across people. Risk reductions and marginal *WTP* amounts must therefore be recognized as jointly distributed random variables, which are potentially correlated, either positively or negatively. When two random variables are correlated, it is not true that the average of their products is equal to the product of their averages. Yet we are implicitly assuming the absence of any correlation when we multiply an aggregate risk reduction by the *VSL*.¹⁹

To illustrate the problem with using a *VSL* when risk reductions may be correlated with *WTP*, consider the case of tighter environmental standards. People in areas already in compliance with the new standard may have higher incomes and may experience no risk

¹⁹ Elementary statistics tells us that the product of two random variables, X and Y , has an expected value given by $E[XY] = E[X]E[Y]$ only when X and Y are independent. If the two variables are correlated, then the expected value of their product depends on the covariance between the two variables: $E[XY] = E[X]E[Y] + Cov[XY]$. In this case, we may consider individual risk reductions as X and individual marginal *WTP* amounts as Y . The total number of statistical lives involved would then be $nE[X]$, where n is the size of the affected population. We want to know the average value of XY which, when multiplied by n , yields overall social benefits.

reductions at all as a result of the policy. People whose neighborhoods are vastly out of compliance may have lower incomes but may reap large risk reductions via the policy because the new standards are binding. If marginal *WTP* for risk reductions increases with income, this would be a situation where individual marginal *WTP* is negatively correlated with prospective individual risk reductions. This means that the estimated overall social benefits would be *overstated* by simply multiplying a population average *VSL* by the aggregate of individual risk reductions. In contrast, estimated overall social benefits would be *understated* in a case where individual marginal *WTP* amounts were positively correlated with the individual risk reductions to be experienced.²⁰

Of course, practitioners must also document policy *costs* in addition to policy benefits. As with marginal *WTP* and risk reductions, individual costs may be correlated with the sizes of individual risk reductions.

Disaggregation and equity assessment

For policy evaluation, the most important adjunct to the efficiency assessment of a conventional benefit-cost analysis is a careful consideration of the distributional consequences of the policy. If we are interested only in overall net social benefits, it is acceptable to calculate aggregate social benefits and aggregate social costs separately, and then to take their difference to yield net social benefits.²¹ However, because the distributions of costs and benefits may differ, equity assessment cannot be done solely in terms of benefits (or solely in term of costs). Individual *net* benefits must be studied. This means that for equity assessments, it is necessary to

²⁰ These two scenarios correspond to Case D and Case E in Appendix A, which provides a concrete numerical example of how premature aggregation can produce misleading results when willingness to pay for a risk reduction is correlated with the size of the risk reduction to be experienced.

²¹ That is: $NSB = SB - SC = \sum_i (MWTP_i \Delta r_{ii}) - \sum_i (MC_i \Delta r_i)$, where *NSB* = net social benefits, *SB* = overall social benefits, and *SC* = overall social costs.

focus attention on the distribution of *individual* net benefits before they are aggregated to produce net *social* benefits.²²

For example, it is possible for a policy to appear progressive in terms of just its benefits. In other words, the policy may confer monetized benefits upon lower-income people, with the benefits representing a larger proportion of income for lower-income people than for higher income people. However, the same policy may also be highly regressive in terms of its costs. That is, lower-income people may end up paying a larger proportion of their incomes to obtain these benefits than would higher income people.²³

Differentiated VSLs (individuation) and equity assessment

Sunstein (2004) makes a persuasive case for the individuation of *WTP* estimates—that *WTP* should be differentiated to reflect differences in risks and in affected populations. Recent evidence concerning this heterogeneity in *WTP* is presented in a special issue of the *Journal of Risk and Uncertainty* (see Viscusi 2010).

The main reason many people object to differentiated *VSL* estimates is that “putting different values on the lives of different people” seems unfair. Some constituencies object to *WTP*-based measures because they tend to suggest greater benefits for people with higher incomes. While it is true that the benefits may be relatively greater for high-income groups because of their greater willingness to pay, high-income individuals may bear an even greater share of the social costs of the risk reduction policy. This means the policy could still be “progressive,” rather than “regressive,” in the sense that the *net* benefits from the policy account for a higher proportion of income for the poor than for the rich. Thus, such a policy could still

²² Net social benefits are the aggregate of individual net benefits: $NSB = \sum_i (MWTP_i - MC_i) \Delta r_i$.

²³ Zerbe (1974) offers some early discussion of the regressivity of some types of environmental improvements. His paper considers optimal jurisdictions for environmental regulation (which is pertinent to the environmental federalism literature).

narrow the gap between rich and poor, even if the benefits to the rich are greater when costs are disregarded.

Any *VSL* estimate reflects an underlying demand function (where individual demand functions are basic to consumer theory in modern economics), so it is unlikely that any such demand function could imply a *WTP* that is a simple constant that can be applied for all types of individuals and all kinds of risks. As for any other good or service, willingness to pay for health risk reductions will depend upon the specifications and the quantity of the good in question, the individual's income, the prices (or rationed quantities) of substitutes or complements, and the individual's preferences. The distribution across individuals of the *benefits* of some type of publicly-provided risk reduction can thus be expected to depend upon (a.) the type of risk and the amount by which it is reduced (or perceived to be reduced) for each individual, (b.) the individual's income, (c.) the availability of other means for each individual to mitigate or compensate for this type of risk, and (d.) the individual's subjective disutility from this type of risk, degree of risk aversion, and discount rate in the case of latent risks that will come to bear sometime in the future.

People need to understand that a policy of “equal protection” via risk-reduction policies—which ignores heterogeneous preferences and incomes and employs a single one-size-fits-all value of a statistical life—is likely to be ethically acceptable only if this protection is purely a gift. Unfortunately, many types of policies and regulations are akin to “unfunded mandates.” Thus, in many important policy and regulatory contexts, people will be required to pay for a substantial share of this extra protection themselves, through higher prices or taxes, lower wages, or lower investment returns, even though they may have other priorities for their scarce resources.

Conclusions: “Euthanizing” the *VSL*

At a recent conference sponsored by the MacArthur Foundation,²⁴ an audience member asked what it would take to increase acceptance of the use of benefit-cost analysis in decisions about the provision of preventive healthcare and the reduction of environmental and safety risks. It seems to me that a necessary (although probably not entirely sufficient) condition for broader public acceptance would be for economists to be able to discuss and report their estimated social benefits of environmental, health or safety policies in a way that neither confuses nor offends non-economists.

This does not necessarily mean that we must change all our current methods for measuring the tradeoffs that people reveal or state they are willing to make with respect to their own health risks—although there certainly are more-general ways to characterize the types of choices people make with regard to altering their risks of experiencing a wide range of adverse future health profiles.²⁵ However, it does seem necessary for us, as a profession, to stop using the term “value of a statistical life” to describe the kinds of numbers we develop. I have suggested that a technically correct measure of what people would be willing to trade for reductions in risk of policy-relevant sizes could be described as “willingness to swap (*WTS*) other goods and services for a microrisk (μr) reduction *in a specified type of risk*.” An abbreviation for this mouthful of terms could be “*WTS*(μr)” or “*WTS* _{μr} ” or simply “*WTS*.” Unfortunately, this is not quite as easy to say as “the *VSL*”. But we could always make it easier by inserting the “i” from “willingness,” which would yield “*WiTS*.”

²⁴ “Unleashing the Power of Social Benefit-Cost Analysis: Removing Barriers,” sponsored by the Benefit-Cost Analysis Center at the University of Washington’s Evans School of Public Affairs and funded by The John D. and Catherine T. MacArthur Foundation, October 19-20, 2009, Washington, D.C.

²⁵ Such a more-general approach has recently been proposed and demonstrated in a series of papers based on Cameron and DeShazo (2010).

While this proposed minimum change is a small adjustment, it may have a very large effect in terms of public relations. To echo an observation by Howard (1984, p. 408) “Although this change is cosmetic only, we should remember the size of the cosmetic industry.”

To properly implement this new terminology, it will also be important to be more specific about our benefits estimates. We should be careful to specify the *nature* of the health threat for which the risk is being reduced, such as “ $WTS_{\mu r}$ (sudden death from heart attack)” or “ $WTS_{\mu r}$ (death ten years earlier than otherwise, after five years of respiratory disease).” The characteristics of the affected individual will also be expected to influence this WTS systematically. Thus a complete set of arguments must include the nature of the risk *and* the characteristics of the affected person or subpopulation. If a benefits transfer is being attempted from one context to another, we should be meticulous about noting any differences between the “study” context and the “policy” context for such a transfer.

Finally, when we report our estimates of $WTS_{\mu r}$, it will be particularly important to not yield to the temptation to cut corners by simply using “dollars” as the units. We should be meticulous about reporting $WTS_{\mu r}$ in “dollars worth of alternative goods and services.” This will be especially important in the abstracts, introductions, and conclusions of our papers, from which some readers are likely to glean their initial impressions of our research. This careful terminology will keep in the forefront of people’s minds the idea that consumers typically need (or will be asked) to swap other desirable things for policy-sized reductions in some specific type of risk to their life or health.

It is a challenge to change a discipline’s long-standing terminology. Even if we start right now, it may take a decade or more before we will cease needing to include a footnote with the instructions: “To achieve a measure that is equivalent to the construct formerly known as

“the value of a statistical life” (or the *VSL*), take the WTS_{μ} measures reported in this table and multiply by 1 million.”

It will also take considerable self-discipline to achieve the conversion. The more “experienced” among us are sufficiently set in our ways that it may be nearly impossible to purge the *VSL* terminology from our casual conversations. Canada’s experience in converting to the metric system, between 1975 and 1983, might be a model for what we can expect. Many older Canadians will never be able to think in terms of litres and kilograms, but the “measuring and weighing” curriculum for schoolchildren provides an opportunity for younger Canadians to “think metric” right from the start.

By analogy, a necessary first step towards euthanizing “the value of a statistical life” (by removing it from our professional vocabulary) would be to ease this unhelpful terminology out of all of our introductory textbooks and our introductory courses. Anyone who has taught this topic knows how easy it is to squander an entire lecture period trying to smooth the ruffled feathers of indignant and idealistic college students who chafe at the idea of monetizing the value of a human being. Students should probably not even encounter the shorthand that is the “value of a statistical life” (if ever) until they thoroughly understand the underlying necessity for tradeoffs in policy-making, and the real empirical evidence about people’s willingness to swap other goods and services for policy-sized risk reductions. In an introductory course, the term should not be introduced at all, and in more-advanced texts, it should warrant no more than a footnote, noting it as an archaic usage. Eliminating the *VSL* terminology would allow those of us who teach to spend scarce lecture time on more substantive concepts, rather than just damage control.

In addition to eliminating “VSL” from our introductory texts and classes, it would help to develop a stable of standard and accessible examples of cases where people are commonly observed to make tradeoffs between risk and money. At the time of this writing, for example, the Toyota Motor Company has dealt with the recall of its vehicles for sudden unintended acceleration and brake problems by embarking on a far-reaching sales program that apparently includes zero percent financing and free maintenance. To the extent that Toyota’s customers are willing to accept small increases in expected mortality from accidents in exchange for a better deal on the purchase of a Toyota, they are voluntarily trading risks and dollars.

While it may be somewhat inconvenient to move away from the *VSL* terminology, by doing so, economists may be able to gain a much better public response to our hard-earned quantitative results. Moreover, government agencies responsible for policy-making related to public risk reduction could improve their public image by avoiding—like the plague—any references to the value of a statistical life. Instead, they should be meticulous in using more-precise descriptions akin to what I have proposed here.

Greater public sympathy (i.e. less misplaced public outrage) should enable our research to have greater influence on the development of rational policies. But this can happen only if we make an effort to keep the highly counterproductive terminology of the value of a statistical life out of our professional writings and our communications with the press and the general public.

Finally, we need to resist the perennial urge to provide single one-size-fits-all estimates for willingness to pay for risk reductions. The willingness to swap other goods and services for specific types of risk reductions can be expected to vary with all of the factors that shift other types of demand functions—income, the prices of substitutes and complements, and tastes or preferences regarding the good in question. Thus, it is not appropriate to average *WTS* estimates,

arbitrarily, across a wide variety of risk reductions affecting different populations. Instead, policymakers should seek to match individual *WTS* estimates as closely as possible to the policy context in which they are to be used. This implies a need for considerably more basic research to help us to better understand the nature of heterogeneity in the *WTS*.

References:

Ammer, Christine (1997) *The American Heritage dictionary of idioms: the most comprehensive collection of idiomatic expressions and phrases*, Boston, MA: Houghton Mifflin.

Amond, Jordan (1985) *Dictionary of Word Origins: A History of the Words, Expression and Cliches We Use*, New York: Kensington Publishing, p. 234.

Borenstein, Seth (2008) "AP IMPACT: An American life worth less today," Thursday, July 10, Associated Press. Archived by the San Francisco Chronicle at "<http://www.sfgate.com/cgi-bin/article.cgi?f=/n/a/2008/07/10/national/a110546D37.DTL> "

Cameron, Trudy Ann (2008) "The Value of a Statistical Life: [They] do not think it means what [we] think it means," *AERE Newsletter* 28(2), p. 36-39.

Cameron, Trudy Ann, J.R. DeShazo (2010) "Demand for Health Risk Reductions," manuscript, Department of Economics, University of Oregon, Eugene, OR

Carruthers, Bruce G. (2009) "Can social science numbers save public policy from politics?" *Regulation & Governance* 3:287-290.

Daniell, David (1994) *William Tyndale: a biography*, New Haven: Yale University Press.

Fahrenthold, David A. (2008) "Cosmic Markdown: EPA Says Life Is Worth Less," *Washington Post*, Saturday, July 19; p. A01)

Fourcade, Marion (2009) "The political valuation of life," *Regulation & Governance* 3: 291-297.

Howard, Ronald A. (1984). "On fates comparable to death". *Management Science* 30: 407–422.

Howard, Ronald A. (1989). "Microrisks for medical decision analysis." *International Journal of Technology Assessment in Health Care* 5: 357–370.

Huff, Darrell (1954) *How to Lie with Statistics* (Illustrated by Irving Geis). New York: W.W. Norton and Company.

Laughland, Drew, A. Myrick Freeman III, Calvin Franz, Aylin Sertkaya (Eastern Research Group, Inc.) and Keith D. Lind (AARP Public Policy Institute) (2007), "Exploring the Role of Cost-Benefit Analysis in Government Regulations," AARP working paper #207-14 (September).

Mackay, Charles (1877). *The Gaelic etymology of the languages of western Europe and more especially of the English and Lowland Scotch of their slang, cant and colloquial dialects*. London: N. Trübner and Co.

Mrozek, Janusz R. and Laura O. Taylor (2002) What determines the value of life? A meta-analysis," *Journal of Policy Analysis and Management* 21: 253-270.

Robinson, Lisa A. (2009) "Valuing lives, valuing risks, and respecting preferences in regulatory analysis," *Regulation & Governance* **3**:298-305.

Steele, J. Michael (2005). "Darrell Huff and Fifty Years of *How to Lie with Statistics*" *Statistical Science*, 20 (3), 2005, 205–209

Sunstein, Cass R. (2004) "Valuing life: a plea for disaggregation" *Duke Law Journal* 54: 385-445.

Twain, Mark (1907). "Chapters from My Autobiography," *North American Review*, No. DCXVIII, July 5. Available at <http://www.gutenberg.org/files/19987/19987.txt>

Viscusi, W. Kip (2009a) "The devaluation of life," *Regulation & Governance* **3**:103-127.

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 (2009b) "Reply to the comments on 'The devaluation of life'," *Regulation & Governance* **3**: 306-309.

Viscusi, W. Kip and Joseph E. Aldy (2003) "The value of a statistical life: A critical review of market estimates throughout the world," *Journal of Risk and Uncertainty* **27**:5–76.

Zerbe, Richard O. (1974) "Optimal environmental jurisdictions," *Ecology Law Quarterly* **4**(2):193-245.

Appendix A: Numerical Example of the Perils of Premature Aggregation

Heterogeneity in risk reductions, combined with heterogeneity in individual willingness to pay for these risk reductions, can produce misleading results if one calculates aggregate benefits prematurely. Premature aggregation means that one first aggregates risk reductions, and then multiplies by an average *VSL* estimate for the population, rather than first calculating individual *WTP* for individual risk reductions and then aggregating these *WTP* measures. Consider the simple numerical example in Table 1. In each of the five cases (A through E), there are one million people in the overall population, divided into two equal-sized groups. In all five cases, based on the data in columns (1) and (2), individual *WTP* for a microrisk reduction averages \$7 across the population. Likewise, based on the data in columns (1) and (3), in all five cases the aggregate risk reduction across the one million people is exactly 1.0 statistical life. Thus if we aggregate the risks first, we will always have one statistical life and if we average the *VSLs* implied by the average *WTP* amounts, we will always get \$7 million. If we aggregate the risk reductions and the willingness to pay information separately, and then multiply, the social benefits will always be \$7 million. But what happens if we do not aggregate prematurely?

In Case A in Table 1, both groups are the same. All one million people in the population share a common willingness to pay for risk reductions that corresponds to a *VSL* of \$7 million ($=\$7/0.000001$), and everyone faces the same size risk reduction (one in a million). Both groups will experience benefits of \$3.5 million, so aggregate benefits produced in Case A will be \$7 million. If everyone is the same, then, it does not matter whether we aggregate risks and multiply by a *VSL*, or preserve disaggregated risks and disaggregated *WTP* amounts and aggregate across the population only after we have multiplied individual risk changes by individual *WTP* to get individual benefits.

In Case B, individual *WTP* for risk reductions differs across the two groups (either because incomes differ or because tastes or risk perceptions differ), but the risk reductions faced by the two groups are identical. In Case C, incomes, preferences and risk perceptions are the same, leading to identical *WTP* for risk reductions, but there are different sized risk reductions for each group. In both cases, if we first multiply individual risk reductions by individual *WTP* to get individual benefits, and then aggregate these benefits, overall benefits will still be \$7 million (because there is no correlation between *WTP* and risk reductions in these two cases).

The results are more interesting for the final two cases, however. In Case D, the group with the *higher* individual *WTP* faces the *smaller* risk reduction, while the group with the *lower* individual *WTP* faces the *larger* risk reduction (i.e. *WTP* is *negatively* correlated with the size of the risk reduction). In this case, if we calculate individual benefits before aggregating, the overall benefit for the low-risk group is \$1.5 million and the overall benefit for the high-risk group is \$3 million (column 5), for total social benefits of \$4.5 million (column 6). Conversely, in Case E, the group with the *lower* individual *WTP* faces the *smaller* risk reduction (i.e. *WTP* is *positively* correlated with the size of the risk reduction). If we again calculate individual benefits before aggregating, the overall benefit for the low-risk group is \$0.5 million and the overall benefit for the high-risk group is \$9 million (column 5), for total social benefits of \$9.5 million (column 6).

The implication of this example is as follows: The conventional strategy—of aggregating physical risk changes into some number of statistical lives, and then multiplying by some one-size-fits-all *VSL*—will tend to (a.) overestimate true social benefits if risk reductions and *WTP* are negatively correlated, and (b.) underestimate true social benefits if risk reductions and *WTP* are positively correlated. This can lead to biases if individual risk reductions and individual *WTP* amounts are non-uniform in the population and vary together in a systematic fashion.

Table 1: Calculation of benefits without premature aggregation.

Cases	(1) Number of people in each group	(2) Individual <i>WTP</i> per 0.000001 risk reduction in each group	(3) Risk reduction for everyone in each group	(5) Benefits by group	(6) Overall benefits
A: Uniform risk reductions, Uniform <i>WTP</i>	500,000 500,000	\$7 \$7	0.000001 0.000001	\$3.5 million \$3.5 million	\$7 million
B: Uniform risk reduction, Differing <i>WTP</i>	500,000 500,000	\$2 \$12	0.000001 0.000001	\$1 million \$6 million	\$7 million
C: Differing risk reductions, Uniform <i>WTP</i>	500,000 500,000	\$7 \$7	0.0000005 0.0000015	\$1.75 million \$5.25million	\$7 million
D: Risk reductions and <i>WTP</i> <i>negatively</i> correlated	500,000 500,000	\$12 \$2	0.0000005 0.0000015	\$1.5 million \$3 million	\$4.5 million
E: Risk reductions and <i>WTP</i> <i>positively</i> correlated	500,000 500,000	\$2 \$12	0.0000005 0.0000015	\$0.5 million \$9 million	\$9.5 million

Note: In all cases, aggregate risk reduction is one statistical life, and the average VSL would be \$7 million.