

Cass R. Sunstein*

The value of a statistical life: some clarifications and puzzles

Abstract: Many people have wondered why the US government conducts cost-benefit analysis with close reference to the value of a statistical life (VSL). It is helpful to answer that question by reference to the “Easy Cases,” in which those who benefit from regulatory protection must pay for it. In such cases, WTP is usually the right foundation for VSL, because beneficiaries are hardly helped by being forced to pay for regulatory protection that they believe not to be in their interests. In the Easy Cases, arguments from both welfare and autonomy support the use of WTP and VSL (with potentially important qualifications involving imperfect information and behavioral market failures). The analysis is less straightforward in harder cases, in which beneficiaries do not pay for all of the cost of what they receive (and may pay little of that cost). In such cases, arguments from welfare and autonomy might not lead in any clear direction. In the harder cases, regulation might be justified on welfare grounds even if the cost-benefit analysis (based on VSL) suggests that it is not. In principle, a direct inquiry into welfare (the master concept) would be preferable to use of cost-benefit analysis. In the harder cases, distributional considerations might also count in favor of proceeding (as prioritarianism suggests). But at the current time, direct inquiries into welfare consequences and into distributional effects are challenging in practice, and hence regulators should generally rely on cost-benefit analysis, making welfarist adjustments, or adjustments based on distributional considerations, only in compelling cases.

Keywords: cost-benefit analysis; value of a statistical life; willingness to pay.

*Corresponding author: Cass R. Sunstein, Harvard Law School, Areeda Hall 225, 1563 Massachusetts Avenue, Cambridge, MA 02138, USA, e-mail: csunstei@law.harvard.edu

1 Introduction

Why should government conduct cost-benefit analysis with close reference to the value of a statistical life (VSL)? What is the argument for relying on people’s willingness to pay (WTP) in regulatory policy? Why should government care about WTP at all?

It is useful to answer these questions by focusing on the “Easy Cases,” in which each individual who benefits from regulatory protection must pay all of the cost. In such cases, WTP is usually the right foundation for VSL, because beneficiaries are hardly helped by being forced to pay for regulatory programs that they do not believe to be in their interests. The major qualifications here involve lack of information and bounded rationality, potentially taking the form of “behavioral market failures.” But if those qualifications are put to one side, the argument for using VSL is quite powerful in the Easy Cases on grounds of both welfare and (less obviously) autonomy. Indeed, that argument calls for a high degree of individuation with respect to VSL, in recognition of the fact that people’s WTP for risk reduction, and hence VSL, varies both across risks and across individuals. To the extent that the empirical evidence is secure, it makes sense to have varying VSLs both for different kinds of risks (such as cancer deaths) and for different social groups (such as children), certainly in the Easy Cases.

The analysis is much less straightforward in harder cases, in which beneficiaries pay only a fraction of the cost of what they receive. In such cases, beneficiaries may well be net winners from regulation. They are receiving benefits, and they are paying less than the full cost of those benefits. Of course there will likely be losers as well, and the regulation might prove unjustified on balance – but unlike in the Easy Cases, arguments from welfare and autonomy may not lead in any clear direction. As we shall see, it is possible that regulation is justified on welfare grounds even if the cost-benefit analysis (based on VSL) suggests that it is not. The reason is that in terms of welfare, the winners may win more than the losers lose (Adler, 2011).

Suppose, for example, that as a result of some regulation, poor workers gain a great deal in terms of safety while (wealthy) consumers have to pay more for the goods that the workers produce. Even if the monetized costs exceed the monetized benefits, the welfare gains to workers might be higher than the welfare losses to consumers. Because welfare is the master concept, and because cost-benefit analysis is simply an imperfect means of assessing welfare effects, this point cannot be ignored, at least as a matter of principle.

In such cases, it is also possible that regulation is justified on redistributive grounds, if it helps those who need help while hurting those who are quite well-off. It is important to see that the best response to unjustified inequality is a redistributive income tax, not regulation – which is a crude and potentially counterproductive redistributive tool (Adler, 2011; Kaplow and Shavell, 1994). But suppose that we are dealing with the harder cases and that a nation lacks an optimal income tax and seeks greater redistribution. If so, it is possible (though far from inevitable) that regulation is justified – and that the use of a VSL that exceeds the WTP of the beneficiaries will produce desirable redistribution (and

also be justified on welfare grounds). I outline the circumstances in which this conclusion might hold.

The conclusion is that in the Easy Cases, the argument for using VSL, based on WTP, is fairly secure. But in the hard cases, the use of VSL may produce a cost-benefit analysis that fails to point in the right directions. At the same time, it is exceedingly difficult to measure welfare directly, and regulation is at best an imperfect redistributive tool. In the hard cases, regulators should depart from the outcome of cost-benefit analysis, based on VSL, only when there is compelling reason to believe that the regulation is nonetheless justified on the ground that it will either promote welfare or achieve important distributive goals.¹

2 WTP: theory and practice

Let us begin with a very brief account of existing practice (Cropper, Hammitt, and Robinson, 2011; Lindhjem, Navrud, Braathen, and Biauxque, 2011; Robinson and Hammitt, 2013; Sunstein, 2014; Viscusi, 1992). We are dealing, of course, with regulations that reduce mortality risks, typically in areas that involve health, safety, and the environment. Other kinds of social welfare regulations – involving, for example, social security, civil rights, and privacy – are not involved, though for those who are concerned with cost-benefit analysis, they create their own kinds of puzzles.

To produce monetary amounts for statistical risks, agencies rely on two kinds of evidence (Thaler and Rosen, 1976). The first and most important involves real-world markets, producing evidence of compensation levels for actual risks (Viscusi, 2010; Viscusi and Aldy, 2003). In the workplace and in the market for consumer goods, additional safety has a price; market evidence is investigated to identify that price. The advantage of such real-world markets is that under certain assumptions, they will reveal people's actual preferences, especially when large populations are aggregated. In part for this reason, real-world markets provide the foundation for actual government practice (DOT, 2013). One disadvantage of these studies is that the underlying decisions are "noisy." When people take jobs or purchase products, it is not easy for them to isolate the particular component that is attributable to mortality risks. There are related questions, taken up below, about whether people are sufficiently informed and whether their decisions are fully rational.

¹ The empirical analysis of the distributional effects of regulations remains in its early stages, and hence the discussion here is largely theoretical, with occasional references to the incipient empirical literature.

The second kind of evidence comes from contingent valuation studies, which ask people how much they are willing to pay to reduce statistical risks (Williams, 2013). For example: *How much would you be willing to pay to avoid a 1/100,000 risk of getting cancer as a result of arsenic in drinking water?* The advantage of contingent valuation studies is that they can isolate people's willingness to pay to avoid mortality risks. The disadvantage is that the questions are hypothetical and unfamiliar, and there are many reasons to wonder whether they provide an accurate measurement (Diamond and Hausman, 1994).

The relevant risks usually are in the general range of 1/10,000–1/100,000. The calculation of VSL is a product of simple arithmetic. Suppose that workers must be paid \$900, on average, to assume a risk of 1/10,000. If so, the VSL would be said to be \$9 million. Note in this regard that the Department of Transportation, building on the relevant literature, adopted a revised VSL estimate of \$9.1 million in 2013, with suitable adjustments for future years (DOT, n.d.). That estimate fits with existing understandings in the technical literature, which has become increasingly refined over time.

For a few of the foundational labor market studies on which agencies have long relied, consider the following Table 1 (EPA, 2000):

A large advantage of labor market studies of this kind is that they avoid the lively disputes over the use of “willingness to pay” or “willingness to accept” (WTA) in regulatory policy (Korobkin, 2003). In both experiments and the real world, people tend to demand more to give up a good than they are willing to pay to obtain it in the first instance – a disparity that seems to complicate efforts to assign monetary values to regulatory benefits, including mortality and morbidity. If people are willing to pay \$25 to eliminate an existing risk of 1/100,000, but demand \$100 to incur a new risk of 1/100,000, then it is difficult to know how to proceed for purposes of monetary valuation of risks. Should agencies use \$25, \$100, or some

Table 1 Early labor market studies on the value of life.

Study	VSL (in US\$)
Kniesner and Leith (1991)	0.7 million
Smith and Gilbert (1984)	0.8 million
Dillingham (1985)	1.1 million
Marin and Psacharopoulos (1982)	3.4 million
V.K. Smith (1976)	5.7 million
Viscusi (1981)	7.9 million
Leigh and Folsom (1984)	11.7 million
Leigh (1987)	12.6 million
Garen (1988)	16.3 million

intermediate figure? Fortunately, this problem dissipates in the context of labor market studies. If workers who face a risk of 1/10,000 are paid \$600 more for doing so, and if workers who refuse to face such a risk are paid \$600 less, then it is irrelevant whether agencies speak in terms of WTP or WTA. And indeed, there appears to be no difference between the two in this context (Kniesner et al., 2012).

Despite the widespread use of VSL, based on the relevant research, there remains considerable controversy about whether the resulting figures actually capture people's informed choices, and whether other methods might be preferable (Adler, 2011; Ashenfelter and Greenstone, 2002; Bronsteen, Buccafusco, and Masur, 2013; Dorman, 1996; Droman and Hagstrom, 1998). I will not explore that controversy in detail here (though some of the discussion will bear on it). The basic goal is to see whether VSL, based on WTP, is the right basis for policy *assuming that we have properly identified WTP*.

3 Easy cases

For the sake of simplicity, assume a society in which people face multiple risks of 1/100,000, and in which every person is both adequately informed and willing to pay no more and no less than \$60 to eliminate each of those risks. Assume too that the cost of eliminating these 1/100,000 risks is widely variable, ranging from close to zero to many billions of dollars. Assume finally that the cost of eliminating any risk is *borne entirely by each individual who benefits from eliminating that risk*. For example, people's water bills will entirely reflect the costs of a policy that eliminates a 1/100,000 risk of getting cancer from arsenic in drinking water. If the per-person cost is \$100, each water bill will be increased by exactly that amount. We are assuming, then, that each individual is willing to pay no more and no less than \$60 to eliminate each risk of 1/100,000, and we are assuming as well that the cost of eliminating each risk will be fully borne by each individual.

3.1 Welfare and autonomy

With these assumptions, the argument for using WTP to calculate VSL is straightforward. Regulation amounts to a forced exchange.² It tells people that they must

² Of course it is true that under the simple economic model, with perfect competition and full information, we would expect the market to provide these benefits if those affected are willing to pay the costs. The analysis of Easy Case, as of harder ones, depends on an antecedent judgment that there has been some kind of market failure (perhaps including a "behavioral market failure," see Sunstein, 2013a).

purchase certain benefits for a certain amount. Why should government force people to pay for things that they do not want?

1. *Welfare*. Begin with welfare (Adler, 2011, 2013; Graham, 2013). By hypothesis, a forced exchange on terms that people dislike will make them worse off. It will require them to buy something that they do not want, undoubtedly because they want other things more. They might want to use the relevant money not to eliminate a mortality risk of 1 in 100,000, but to buy food or education or medical care, or to eliminate a mortality risk of 1 in 20,000 or 1 in 10,000. At first glance, use of WTP, on the assumptions that I am making, seems hard to contest. The argument for respecting WTP is a form of simple, *ex ante* Paretianism. In free societies that are concerned with people's welfare, we should begin by asking people what they want, and if people do not want certain goods, we should presume that they know their own priorities. A forced exchange will decrease their welfare. Indeed, a forced exchange would violate Mill's Harm Principle without apparent justification (Mill, 1859).

It is of course natural to ask about the meaning of the term "welfare." Does it mean happiness, narrowly conceived? Might it include whatever makes lives good and meaningful, even if happiness, strictly speaking, is not involved? For present purposes, we can understand the term "welfare" very broadly, to capture whatever informed people care about. With that broad understanding, the capacious idea of welfare can be separated from the narrower one of utility, and it can include many elements of life that may not count as "happiness." (It is true that if we adopt a conception of welfare that does not capture what choosers care about, and that refers to an objective account of goods, independent of the concerns of informed choosers, the use of WTP will seem puzzling.)

For purposes of evaluating regulation, it does not matter if the existing distribution of income is unjust or if poor people are, in some intelligible sense, coerced to run certain risks (as they might be if they live in bad or desperate conditions and have few, and bad, opportunities). The remedy for unjust distributions, and for that form of coercion, is hardly to require people to buy regulatory benefits on terms that they find unacceptable. Suppose that people are willing to pay only \$60 to eliminate a 1/100,000 risk because they are not rich, and that if they had double their current wealth, they would be willing to pay \$120. Government does people no favors by forcing them to pay the amount that they would pay if they had more money. It does not help them. On the contrary, it hurts them.

In the Easy Cases, it is true but irrelevant that willingness to pay is dependent on ability to pay. Suppose that certain people are willing to pay a very

small amount to eliminate a 1/100,000 risk, not because they would not obtain much welfare from eliminating that risk, but because they have little money to spend. It might be thought that the welfare gain is not captured by the very small amount that they are willing to pay. That thought should not be used to collapse two different questions. (a) Should government force people to spend more than the very small amount that they are willing to pay, because the welfare gain would not be trivial? (b) Should government itself, though a compelled expenditure from consumers or taxpayers, be willing to spend more than poor people are willing to pay to reduce a risk, because the welfare gain would not be trivial?

The answer to (b) is not clear, and if government is concerned with welfare, it might well answer in the affirmative. But (a) is a quite different question. Unless there is a problem of lack of information or bounded rationality, the answer to (a) is clearly no. The welfare gain would not justify the welfare loss. That is the beauty of the WTP approach, which offers an automatic test of the welfare consequences of regulation.

2. *Autonomy.* Perhaps regulatory policy should not be based on welfare; perhaps it is unclear what “welfare” really means. Even if so, WTP might be defended instead on the ground of personal autonomy (Adler, 2011; Dworkin, 2002; Sunstein, 2013a). On this view, people should be sovereign over their own lives, and government should respect people’s choices about how to use their limited resources (again, as long as those choices are informed).

When people decline to devote more than \$60 to the elimination of a 1/100,000 risk, it is because they would like to spend the money in a way that seems to them more desirable. If regulators do not use people’s actual (informed) judgments, then they are insulting their autonomy. If people in a free society are entitled to have a kind of mastery over the conduct of their own lives, then they should be permitted to make such allocations as they choose. It is usual to justify use of WTP on welfare grounds, but the same approach may be at least equally defensible as a means of respecting the autonomy of persons.

3.2 Disaggregating VSL?

For VSL, government agencies and departments now use a population-wide number. In general, they do not distinguish among either risks or persons. But the arguments thus far suggest that in the Easy Cases, agencies should disaggregate VSL across both risks and persons, and should not rely on a population-wide average or median (Sunstein, 2013a,b; Viscusi, 1992). In principle, every

individual in society has a separate VSL, and each of these varies across risks (Kling, 2011).³

Suppose, for example, that people are willing to pay no more than \$50 to avoid a 1/100,000 risk of dying in a car crash, but that they are willing to pay up to \$100 to avoid a 1/100,000 risk of dying of cancer.⁴ If government uses a WTP for both risks of \$75, it will force people to pay more than they want to avoid the risks associated with car crashes, and less than they want to avoid risks of cancer. Why should government do that? And if the argument is convincing in this example, it should apply in numerous cases in which WTP and hence VSL vary across mortality risks. The central question would not be conceptual but instead empirical: How does VSL vary, depending on the nature of the particular risk at issue?

With respect to persons, the central idea is that different people should be expected to have a different WTP to avoid mortality risks, and to the extent feasible, regulators should not use a “mass” VSL, but should instead attempt to individuate (Viscusi, 1992). This argument is more controversial, among other things because it might well treat children differently from adults (Leung and Guna, 2006; Williams, 2013) and the elderly differently from the young (Aldy and Viscusi, 2007; Kling et al., 2011),⁵ and because it would appear to treat poor people as less valuable than rich people. Because they have less money, poor people would have a lower VSL than wealthy people. [Similarly, poor areas, and poor nations, would have a lower VSL than wealthier ones (Miller, 2000).] But so long as we are dealing with Easy Cases, differences appear to be appropriate here as well.

The reason is not that poor people are less valuable than rich people. It is that *no one, rich or poor, should be forced to pay more than she is willing to pay*

³ See Kling et al. (2011, p. 7): “Because appropriate valuation of reductions in mortality risk generally requires information on how VRR [Value of Risk Reduction] varies among individuals and with risk characteristics, the SAB recommends that EPA orient its approach toward (a) recognizing the conceptually appropriate method to estimate population benefits and (b) developing a set of estimates of VRR for policy-relevant cases characterized by risk and individual characteristics (or a function relating VRR to risk and individual characteristics).”

⁴ There is a question whether a “cancer premium” is justified by the existing literature. For discussion, see US Environmental Protection Agency (2010) and Kling (2011).

⁵ See Kling et al. (2011, p. 8): “The SAB acknowledges that heterogeneity in WTP across types of risks will be more palatable to some audiences than heterogeneity across affected subpopulations. In the past, for example, the Agency has been criticized for considering VRRs that differ by individuals’ age. This is, however, a failure of communication rather than any theoretical ambiguity about whether economics admits for different demands by different types of people. It can be difficult to convey the distinction between the “intrinsic value of different human beings” and the “different WTP of people in different circumstances.” However, this difficulty does not justify using the wrong benefit measure for proposed policies.”

for the reduction of risks. In fact this idea embodies a norm of equality (and the right one). If poor people are unwilling to pay much for the reduction of serious risks, and if government wants to help, the appropriate response is not a compelled purchase, but a subsidy, perhaps in the form of cash payments. We might even say that part of the right conception of risk equity is that people should not be compelled to pay more than they are willing to pay to reduce risks (unless there is a lack of information or a problem of bounded rationality) (Adler, 2008).

Suppose, for example, that each member of a group of relatively poor people, earning less than \$30,000 annually, is willing to pay only \$45 to eliminate a risk of 1/100,000 – one-half, suppose, of the nation's population-wide median of \$90. Should regulators require every citizen, including those in the relatively poor group, to pay \$90? Government should not force poor people to pay more than their WTP to eliminate statistical risks; forced exchanges of this kind do poor people no good and some harm.

It is tempting to defend a uniform VSL, one that does not distinguish between rich and poor, on the ground that it embodies the right conception of risk equity, treating every person as equal to every other person and redistributing resources in the direction of poor people. But this is an error. A uniform VSL, taken from a population-wide median, does not produce redistribution toward the poor, any more than any other kind of forced exchange. Government does not require people to buy Volvos, even if Volvos would reduce statistical risks. If government required everyone to buy Volvos, it would not be producing desirable redistribution. A uniform VSL has some of the same characteristics as a policy that requires people to buy Volvos. In principle, the government should force exchanges only on terms that people find acceptable, at least if it is genuinely concerned with their welfare. That principle is the correct conception of risk equity.

Note, once again, that the argument for using WTP does not imply satisfaction with the existing distribution of wealth. The problem with forced exchanges is that they do nothing to alter existing distributions. In fact they make poor people worse off, requiring them to use their limited resources for something that they do not want to buy.

To be sure, there are strong pragmatic questions about use of a more disaggregated WTP (Robinson and Hammitt, 2011). Regulators may well lack reliable information about how to disaggregate, whether the question involves risks or persons. The findings of existing studies may not be sufficiently clear for official use. Increased differentiation, for risk characteristics as well as for individuals outside of the labor market (particularly the very young and the very old), is not likely to be possible without greater reliance on stated preference research,

since it allows the researcher to tailor the scenario to particular types of risks and affected populations. But we have seen that there is good reason to question the reliability of stated preference research.

Moreover, many regulatory programs necessarily affect large groups of people at the same time, and it is not possible to tailor the level of stringency to different subgroups, let alone individuals. But if the analysis of the Easy Cases is correct, these concerns involve limitations of knowledge and feasibility; they do not suggest that disaggregation would be a mistake in principle.

3.3 Are there Easy Cases?

Do the Easy Cases seem implausibly unrealistic? In many contexts, they are not likely to match reality. Air pollution regulations are especially important, because they count for a large percentage of both the benefits and the costs of all federal rules, and there is evidence that the costs of such rules are not fully borne by its beneficiaries. Under the Clean Air Act, relatively poor people get disproportionate benefits, and they do not bear their full costs (Kahn, 2001). But for workers' compensation regulation, the situation is very different. With the enactment of workers' compensation programs, nonunionized workers faced a dollar-for-dollar wage reduction, corresponding almost perfectly to the expected value of the benefits that they received (Hedbergian, Gray, and Morgan, 2007). For drinking water regulation, something similar is involved. The entire cost of regulation is passed onto consumers in the form of higher water bills (Sunstein, 2002a).

To be sure, affected industries may respond to the costs of regulation in many ways. Companies could increase prices, decrease wages, and/or decrease firm owners' profits. If air pollution controls are imposed on coal-fired power plants, some of the costs are likely to be borne by individuals other than energy consumers. If regulations are designed to increase worker safety at construction sites, multiple people will pay the cost. In addition, market distortions, such as regulation or taxation of wages, complicate the relationships among the variables.

As noted, each individual's WTP is likely to vary depending on his or her personal characteristics as well as the characteristics of the risk reduction (including the magnitude of the reduction he or she experiences). Even if the costs of safer drinking water are fully reflected in price changes, or if the costs of worker safety are fully reflected in wages, the price or wage change experienced by each individual is not likely to be equal to his or her WTP in light of the diversity of individual preferences.

Nonetheless, the Easy Cases do find a number of real-world analogues. And even where beneficiaries do not pay the full cost of what they obtain, they might pay a substantial portion of it, and for such cases, the analysis of the Easy Cases is at least a place to start.

4 Objections

There are several plausible objections to the use of WTP to calculate VSL, even in the Easy Cases. They point to some important qualifications of the arguments thus far and suggest some puzzles that deserve continuing empirical and conceptual attention.

4.1 “Miswanting”

The first objection is that people may suffer from a problem of “miswanting” (Gilbert and Wilson, 2000; Wilson and Gilbert, 2003). They want some things that do not promote their welfare, and they do not want some things that would promote their welfare. In some settings, people’s decisions appear not to make them happier, even when alternative decisions would do so (Dunn and Norton, 2013). *Predicted welfare*, or welfare at the time of decision, may be very different from *experience welfare*, or welfare as life is actually lived (Kahneman and Varey, 1991). If this is so, then WTP loses much of its underlying justification. People’s choices do not actually promote their welfare (Kahneman, 2003; Kahneman and Varey, 1991; Kahneman, Wakker, and Sarin, 1997; Sunstein, 2013a,b). If government can be confident that people are not willing to pay for goods from which they would greatly benefit, perhaps government should abandon WTP.

A more specific concern is that people’s preferences may have adapted to existing opportunities, including deprivation (Adler and Posner, 1999; Elster, 1985). Thus Tocqueville writes, “Shall I call it a blessing of God, or a last malediction of his anger, this disposition of the soul that makes men insensible to extreme misery and often gives them a sort of depraved taste for the cause of their afflictions?” (de Tocqueville, 1848, p. 317). Perhaps people show a low WTP for environmental goods, including health improvements, simply because they have adjusted to environmental bads, including health risks. Perhaps people’s WTP reflects an effort to reduce cognitive dissonance through the conclusion that risks are lower than they actually are (Akerlof and Dickens, 1984). It is not a lot of fun to think that you face serious dangers, and some people undoubtedly develop an

unduly optimistic account of their actual situation (Sharot, 2010) (a problem to which I will return).

In some contexts, the idea of miswanting raises serious problems for neoclassical economics and for unambivalent enthusiasm for freedom of choice (Conly, 2012). As Daniel Kahneman and Carol Varey have explained, “if people do not know their future experience utilities, or if their preferences of the moment do not accurately reflect what they do know, a case can be made for using experience utility rather than preference as the unit of account in utilitarian calculations.” (Kahneman and Varey, 1991, pp. 128–129). If the basis for use of WTP is welfare, there is a real difficulty, because use of WTP may be imperfectly connected with promoting people’s welfare.

In making these claims, I do not mean to take a contentious position on the nature of welfare. People want their lives to go well, but their understanding of what it means for their life to go well are diverse, and include a diverse array of goods. Hedonic states are important, but they are hardly all that matters. People choose certain activities not because they are fun or joyful, but because they are right to choose, perhaps because they are meaningful. People want their lives to have purpose; they do not want their lives to be simply happy (Benjamin, Kimball, Heffetz, and Rees-Jones, 2012). People sensibly, even virtuously, choose things that they will not in any simple sense “like” (Karlson, Loewenstein, and McCafferty, 2004; Ubel and Loewenstein, 2008).

For example, they may want to help others even when it is not a lot of fun to do that. They may want to do what they are morally obliged to do, even if they do not enjoy it. An important survey suggests that people’s projected choices are *generally* based on what they believe would promote their subjective well-being – but that sometimes people are willing to make choices that would sacrifice their happiness in favor of an assortment of other goals, including (1) promoting the happiness of their family, (2) increasing their control over their lives, (3) increasing their social status, or (4) improving their sense of purpose in life (Benjamin, Kimball, Heffetz, and Rees-Jones, 2012). The point here is not that people seek to promote a narrow conception of welfare, but that whatever their preferred conception, they make mistakes, and these mistakes can be implicated in the use of WTP. On welfarist grounds, there is also an objection that in the Easy Cases, the argument depends on a form of ex ante Paretianism, but that ex post Paretianism might be preferable, at least if we place a large emphasis on equality (Adler, 2011).

With respect to miswanting, autonomy is implicated as well. The idea of autonomy requires not merely respect for whatever preferences people happen to have, but also for preferences that are actually informed, and for social conditions that allow preferences to be developed in a way that does not reflect coercion or injustice (Elster, 1985). With respect to some risks, the relevant preferences are

nonautonomous. Consider the fact that many women face a risk of male harassment or domination (or even violence) under circumstances in which they believe that little can be done – and hence adapt (Khader, 2011).

In the context of ordinary regulatory policy, however, this objection may have more theoretical than practical interest. Typically regulation involves the reduction of low-level mortality risks (say, 1/100,000). In the abstract, there is no reason to believe that the use of people's WTP (say, \$90) is a product of adaptive preferences or a problem of miswanting. It is true that we cannot rule out the possibility that with respect to mortality risks, preferences have adapted in this way (Akerlof and Dickens, 1984). And when WTP does result from adaptive preferences, the judgment about the Easy Cases must be revised. But in the real world of regulatory practice, there is no reason to think that problem arises often, or that it is sufficient to “impeach” the evidence on which regulators rely in using VSL.

4.2 Information and behavioral market failures

A closely related objection would point to an absence of information and to bounded rationality, meant as an umbrella concept for a wide range of findings from behavioral economics (Kahneman, 2011; Sunstein, 2013b; Thaler and Sunstein, 2008). We can use the term “behavioral market failures” to refer to a set of problems that may make markets work imperfectly, including unrealistic optimism, myopia, and self-control problems.

Here is a way to make the point. Imagine a population of people. Let us call them Simpsons (after the character Homer Simpson in the television show of that name) (Thaler and Sunstein, 2008). Simpsons make choices, but the choices reflect systematic errors, in the sense that they are unrealistically optimistic, neglectful of the long-term, and reckless. The Simpsons will have an identifiable WTP to avoid mortality risks (and other risks). By hypothesis, the WTP and the corresponding VSL will be low. But the fact that it is low does not mean that for the Simpsons, the government should use a low VSL in regulatory policy. What matters is the welfare of the population, and the Simpsons' welfare is not adequately capturing by the Simpsons' WTP. Regulators should use preferences that are informed and rational, and that extend over people's life-histories. The Simpsons' preferences do not satisfy those criteria.

No nation is the Simpsons, but as behavioral economists have shown, people often have difficulty dealing with low-probability events (Thaler and Sunstein, 2008). If people are not aware of the risks that they actually face, or if they have a poor understanding of such risks, their WTP, as measured by market evidence, might be too low. Suppose that workers receive a \$90 wage premium for risks of

1 in 100,000. What does this mean, concretely? Are workers actually trading off risks and money? Do they even know about the relevant probabilities? If these questions do not have clear answers, the market evidence might not be reliable. (Note that this is not an objection to basing VSL on informed WTP; it is merely a concern that existing evidence does not allow us to be certain that we are eliciting informed WTP.)

Or perhaps the availability heuristic will lead people to underestimate mortality risks. If people cannot recall a case in which some activity produced illness or death, they might conclude that a risk is trivial even if it is not. Perhaps market evidence will reflect such mistakes. Or perhaps the same heuristic, and probability neglect (Sunstein, 2002b), will lead people to exaggerate risks, producing a WTP that is wildly inflated in light of reality. And if people are unable to understand the meaning of ideas like “one in fifty thousand,” or to respond rationally to such ideas, then there are serious problems with relying on contingent valuation studies to produce WTP.

It is also possible that people’s WTP reflects excessive discounting of future health benefits. If workers are disregarding the future, or applying an implausibly high discount rate, then there is a good argument for not relying on their WTP. In the context of climate change, for example, the temporally distant nature of the harm might well lead to insufficient concern for a potentially catastrophic risk. The same is true for less dramatic risks that people face in their daily lives. Young smokers, for example, undoubtedly give too little attention to the long-term health risks associated with smoking. Those who choose a poor diet and little exercise often fail to consider the long-term effects of their behavior. Self-control problems are an important part of bounded rationality. If a low WTP shows a failure to give adequate attention to the future, then there is reason not to use WTP.

To be sure, a dollar today is worth less than a dollar tomorrow, in part because a dollar today can be invested and made to grow. For money, some kind of discount rate makes a great deal of sense. And for rational reasons, people might prefer welfare today to welfare tomorrow; for example, there is at least some chance of death tonight, which argues for welfare today. The question of how rational people distribute welfare (as opposed to money) over time does not admit of an easy answer. But if people care very little about their future selves, and are willing to impose a great deal of future suffering in return for a small benefit in the present, something has likely gone wrong. An appealing welfarist approach emphasizes preferences that are fully informed and fully rational, and that extend over life-histories (Adler, 2011). If people’s choices do not satisfy these constraints, they are “impeached” from the standpoint of welfare.

When a behavioral market failure is involved, appropriate adjustments should be made to WTP, and the VSL that emerges from WTP should be corrected accordingly. It is possible, of course, that across large aggregations of workers, behavioral market failures are not a serious problem, and hence existing numbers are trustworthy. As noted, no nation consists of Simpsons. But further conceptual and empirical work needs to be done on these issues.

4.3 Rights

A quite different objection would point to people's rights. Perhaps people have a right not to be subjected to risks of a certain magnitude, and the use of WTP will violate that right. It is tempting to think that whatever their WTP, human beings should have a right not to be subject to risks above a particular level. Imagine, for example, that poor people live in a place where they face a 1/20 annual risk of dying from water pollution. That risk is intolerably high. It makes sense to say that the government, or the international community, should take steps to reduce that risk even if the relevant population is poor, even if people are willing to pay only \$1 to eliminate it, and even if the per-person cost is \$2.

As an abstract claim about people's rights, the objection may be correct. Something has gone badly wrong if people are exposed to serious risks and their WTP prevents them from doing anything in response. It would be foolish to suggest that WTP is determinative of the appropriate use of government resources. (Would it make sense to say that government would give poor people a check for \$100 only if they were willing to pay \$100 for the check?) And in many cases, people are subject to risks whose magnitude is indeed a violation of their rights. But for several reasons, this point has little force against my conclusions for the Easy Cases.

The initial problem with this objection is that in the cases under discussion, rights of this kind are usually not involved; we are speaking here of statistically small risks. Suppose that this initial response is unconvincing⁶ and that rights are indeed involved. If so, there is a still more fundamental response. When rights are involved, the proper response is not to force people to buy protection that they do not want, but to provide a subsidy that will give them the benefit for free, or enable

⁶ There are, of course, complex questions about the relationships among rights, wealth, and risk. In one view, these variables cannot be separated from one another, and whether people have a "right" to freedom from certain statistical risks depends on an assessment of the level of resources in the relevant society. In that respect, some analysis of social welfare may be a precondition for any judgment about rights.

them to receive the benefit at what is, for them, an acceptable price. Nothing here is meant to deny the possibility that government should provide certain goods via subsidy, or indeed that subjection to risks above a certain level is a violation of rights. The question instead is one of regulation under the stated assumptions. So long as that is the question, use of WTP does not violate anyone's rights.

4.4 Democracy and markets

An independent objection would stress that people are citizens, not merely consumers. It would urge that regulatory choices should be made after citizens have deliberated with one another about their preferences and values (Sen, 2002). The argument against forced exchanges treats people as consumers; it sees their decisions about safety as the same as their decisions about all other commodities. For some decisions, this approach is badly misconceived. Well-functioning constitutional systems promote deliberative democracy, and many social judgments should be made by citizens engaged in deliberative discussion with one another, rather than by aggregating the individual choices of consumers.

Consider some examples:

- The permissible level of race and sex discrimination is not set by using market evidence, or contingent valuation studies, to see how much people would be willing to pay to discriminate (or to be free from discrimination). Such discrimination is banned, even if discriminators would be willing to pay a lot to avoid associating with members of unpopular groups. Through democratic processes, citizens have decided that certain forms of discrimination are illicit, whatever people's WTP.
- The prohibition against sexual harassment does not emerge from consulting people's WTP. Many harassers would be willing to pay something, perhaps a great deal, for the privilege of harassing. In imaginable circumstances, the harassers' WTP might exceed their victims' WTP to prevent harassment. Nonetheless, harassment is forbidden. One reason is that a goal of the civil rights laws is to alter existing preferences and beliefs, not entrench them.
- Laws that forbid cruelty to animals, and that impose affirmative duties of protection on human beings, stem not from WTP, but from a belief that morality justifies such laws. When laws require protection of animals against cruelty or suffering, it is not decisive that those who are regulated may be willing to pay a significant amount to avoid the regulation. Of course the cost of the regulatory burden might play an important role in deciding whether to impose it. But the underlying moral judgment is rooted in a belief in the prevention of suffering that does not essentially turn on WTP.

Stressing the limits of any approach that takes “preferences” to be the foundation of regulatory policy, Amartya Sen emphasizes that “discussions and exchange, and even political arguments, contribute to the formation and revision of values.” (Sen, 2002, p. 287). He urges that in the particular context of environmental protection, solutions require regulators “to go beyond looking only for the best reflection of given individual preferences, or the most acceptable procedures for choices based on those preferences.” (Sen, 2002, p. 289).

Certainly if we bracket some questions about the real world of democratic policymaking, we should acknowledge that Sen’s claims are both fundamental and correct. They point to some serious limitations on the use of WTP. But it is important not to read such objections for more than they are worth. In trading off safety and health in their private lives, people do not have static values and preferences. Much of the time, human choices are a product of reflection, even if choosers are simply acting as consumers. Reflection and deliberation, including reflection and deliberation with other people, are hardly absent from the market domain. To be sure, moral questions should not be resolved by aggregating private WTP. Sometimes people’s preferences, even though backed by WTP, are morally off-limits (consider sexual harassment), and policy should not take account of them. In addition, people may be unwilling to pay a great deal for goods that have strong moral justifications; animal welfare is a potential example. In these circumstances, the market model is inapplicable and WTP reveals very little.

But what about the Easy Cases? Do these arguments suggest that government should override individual choices about how much to spend to eliminate low-level risks, even when those choices are adequately informed? For environmental protection generally, it is indeed important to go beyond “the best reflection of given individual preferences.” But this point does not mean that people should be required to pay \$100 to eliminate mortality risks of 1/100,000 when they are willing to pay only \$75. If people’s WTP reflects an absence of information, bounded rationality, or insufficient deliberation, then it is important for other people, in government and elsewhere, to draw attention to that fact. And in some cases, a low WTP might be overridden on the ground that it is rooted in errors. But these points should not be taken as a general objection to my conclusion about the Easy Cases, or to suggest that government should force people to reduce statistical risks at an expense that they deem excessive.

4.5 Third-party effects

A final objection would point to effects on third parties. If outsiders would be adversely affected by the undervaluing of a particular risk, and if their welfare is

not being considered, then the WTP calculus is seriously incomplete. This point demonstrates a general and neglected problem for WTP as it is currently used: agencies consider people's WTP to eliminate statistical risks, without taking account of the fact that others – especially family members and close friends – would also be willing to pay something to eliminate those risks.

John might be willing to pay \$25 to eliminate his own risk of 1/100,000, but his wife, Jane, might be willing to pay \$25 to eliminate John's risk also. When John is hurt or killed, John is not the only person who pays the price. If regulators add the WTP, on John's behalf, of John's friends and relatives, the total WTP might soon exceed \$50. This is a real problem for existing uses of WTP. In principle, regulators should consider the full range of people who are adversely affected, not only the person directly at risk (for discussion, see Jones-Lee, 1991, 1992). A great deal of work remains to be done on this topic (Posner and Sunstein, 2005).

This is a legitimate point, but thus far the discussion has been assuming that there are no third-party effects. In the Easy Cases, the argument for using WTP, on the stated assumptions, is that government should not force people to buy goods that are not worthwhile for them. At least at first glance, this argument seems sound with respect to statistical risks of the kind on which I am focusing.

5 Harder cases

There is an obvious artificiality in the assumptions thus far. Most important, people do not always bear the full social costs of the regulatory benefits that they receive. Sometimes they pay only a fraction of those costs – or possibly close to nothing. When this is so, the analysis is much more complicated.

5.1 Welfare and distribution

We have seen that in the context of air pollution regulation, there is a complex set of distributional effects, and on balance, poor people, and members of minority communities, may well be net gainers. Suppose that the result of an air pollution regulation is to improve public health in poor communities, and that those who benefit pay only a small part of the cost. Suppose too that strictly in terms of welfare, they benefit a great deal, perhaps because they are less likely to get sick, perhaps because they live longer lives. Suppose that most of the cost is paid by people who can easily bear it (and hence do not much suffer from paying). A cost-benefit analysis, based on WTP, might not produce an adequate account of the

welfare effects of air pollution regulation. The reason is that in terms of welfare, the people who gain may end up gaining more than the people who lose end up losing. Use of WTP, and hence of VSL, may produce a cost-benefit analysis suggesting that the regulation is a net loser – but on welfare grounds, the analysis might be misleading. It might point in the wrong direction.

The case of rich and poor may be the most vivid, but it is merely illustrative. We could imagine many cases in which cost-benefit analysis, based on WTP, leads to conclusions that do not promote welfare. A safety regulation, designed to protect workers, might increase welfare even if the cost-benefit analysis suggests otherwise. This is true not only if and because WTP does not capture informed, rational preferences over a lifetime (Adler, 2011), but also if and because the welfare effects are not sufficiently captured by the monetary figures. Of course it may also be true that a regulation that has positive net benefits is also bad from the standpoint of promoting welfare. Indeed, we could distinguish among four kinds of cases: (1) net monetary benefits and net welfare benefits; (2) net monetary benefits but net welfare costs; (3) net monetary costs and net welfare costs; and (4) net monetary costs but net welfare benefits. For present purposes, cases (2) and (4) are the interesting ones.

In any case, the welfare effects might not resolve the question about what to do, because distributional effects are important to consider. Recall that Executive Order 13563 explicitly makes “distributional impacts” relevant. If poor people are gaining a great deal, and wealthy people are losing slightly more, the regulation might be justified on distributional grounds (Livermore and Rosenberg, 2013). Consider the idea of prioritarianism, which suggests that the social goal should be to increase overall welfare, but with priority given to the most disadvantaged (Arneson, 2000).

Here is what distinguishes the Easy Cases from the harder ones. We have seen that in the Easy Cases, it does not make much sense to require people to pay more than they are willing to pay. But in the harder cases, people are not paying all of the cost of the benefits that they receive. If so, it is quite possible that they will gain on balance from the relevant regulation, even if their WTP is significantly lower than the cost of the regulation. The more relevant possibility is that on net, society will gain in terms of welfare as well. The point suggests a potentially serious limitation to the use of WTP in regulatory policy when those who benefit from regulations do not pay for them. In such cases, even a unitary VSL may produce misleading results if welfare is our guide. And a more disaggregated VSL, suggesting a lower figure for poor people, may be especially misleading if poor people stand to gain a great deal in terms of welfare.

Indeed, the use of a unitary rather than disaggregated VSL, in which poor people are given more than they are willing to pay for, might be justified in such

circumstances. Thus Viscusi suggests that by “using a uniform VSL across different populations, agencies engage in an implicit form of income redistribution, as benefits to the poor receive a greater weight than is justified by their VSL and benefits to the rich are undervalued.” (Viscusi, 2013).

Some of the most intuitively plausible defenses of cost-benefit analysis speak in terms of the Kaldor-Hicks criterion (sometimes called potential Pareto superiority), which asks whether the winners win more than the losers lose (Adler and Posner, 2006). The central idea is that if the winners could compensate the losers, and there would be a surplus, satisfaction of the Kaldor-Hicks criterion shows a net welfare gain. The criterion raises many doubts and puzzles, but let us simply stipulate that regulation is ordinarily justified if it produces such a welfare gain. The problem is that under certain circumstances, a net loss, in terms of cost-benefit analysis, may coexist with a net gain, in terms of welfare. Because welfare is the master concept, and because monetized numbers are mere proxies, it would seem clear that the proxies would have to yield in favor of the master concept.

It is sometimes argued that if agencies use cost-benefit analysis, everyone, or almost everyone, will benefit in the long-run. Hicks so argued about the Kaldor-Hicks criterion, suggesting that “although we could not say that all the inhabitants of that community would be necessarily better off than they would have been if the community had been organized on some different principle, nevertheless there would be a strong probability that almost all of them would be better off after the lapse of a sufficient amount of time.” (Hicks, 1941, p. 111). Note that this argument is not strictly about cost-benefit analysis – and it would take a great deal of work to show that if agencies use such analysis as a criterion of decision, based on VSL, almost everyone “would be better off after the lapse of a sufficient length of time.” And there is good reason to question Hicks’ argument (Adler, 2011). If government could measure welfare directly, and make individual decisions so as to promote it, it would unquestionably do better on welfare grounds.

With respect to harder cases, related points can be said with respect to autonomy. If poor people do not bear all of the costs of programs that benefit them, the autonomy argument for use of WTP is greatly weakened. Poor people are enjoying a benefit (in whole or in part) for free. It does not insult people’s autonomy to give them a good on terms that they find acceptable.

It is natural to respond that if redistribution is the goal, then it should be produced not through regulation but through the tax system, which is a much more efficient way of transferring resources to people who need help (Kaplow and Shavell, 1994; Shavell, 1981; Weisbach, 2003). In general, the point is correct. But redistribution might not be feasible through the tax system. If not, then regulation in the harder cases cannot be ruled off-limits (despite its inefficiency). To be

sure, the fact that a regulation is helpful to the most disadvantaged is not decisive in its favor. We need an account by which to measure benefits to the most disadvantaged against costs to others. If a regulation is trivially helpful to the most disadvantaged, and if it inflicts huge costs on everyone else, little can be said for it. But everything depends on the magnitude of the relevant effects. A program that produces large gains for the least well-off might well be justified even if it imposes, in terms of WTP, slightly higher costs than benefits on balance.

5.2 Administrability

Notwithstanding these considerations, regulators do not make direct inquiries into welfare. And while Executive Order 13563 allows them to take account of distributional effects, they do not often do so. Why?

One answer involves administrability, or more particularly, the best way to minimize decision costs and error costs. In the real-world cases, regulators might well think that a direct inquiry into welfare, bypassing WTP, would be extremely difficult or perhaps even impossible to operationalize [even bracketing, as I have throughout, the problem of interpersonal comparisons of well-being (Adler, 2011; Elster and Roemer, 1993)]. Regulators might rely on WTP not because it is perfect as a proxy for welfare, or even close to it, but because any more direct welfare inquiry is not tractable. Regulators lack welfare meters, and for that reason alone, they might use standard cost-benefit analysis instead.

If regulators decide that distributional considerations are relevant, they might fear that interest-group warfare would be the consequence, rather than distribution to those who particularly need and deserve help. The larger point is that it is not easy to identify many regulations for which poor people are the clear beneficiaries while healthy people foot the bill. It is far more usual for the costs and benefits of regulations to be widely distributed, so that a variety of demographic groups both enjoy the benefits and pay the costs. To the extent that regulation increases the price of goods and services, it may even be regressive, and there is good reason to think that is a frequent pattern. A great deal of additional work needs to be done on this topic, in order to specify distributional effects with far more precision than has been done to date.

Under current circumstances, a reasonable approach would be for regulators to use WTP as the foundation for decisions, and generally to follow the results of cost-benefit balancing, but to inquire into welfare or distribution in cases in which there is compelling reason to do so. (Consider cases in which the monetized benefits are only mildly higher than the monetized costs, but in which the costs are borne by those who are well-off, and the benefits are

enjoyed mostly by those who are struggling. Regulators might decide to proceed in such circumstances.) In fact this is generally the correct approach, because it is right in principle, and because it does not impose undue information-gathering burdens on regulators.

6 Conclusion

In many contexts, the use of the WTP criterion is controversial, and for legitimate reasons. But for valuation of statistical mortality risks, that criterion makes a great deal of sense, at least as the place to start. In the Easy Cases, people should not be asked to spend more than they are willing to pay for the elimination of risks, even mortality risks. To be sure, it is possible to doubt whether we have accurately measured people's informed choices. Emphasizing the possibility of behavioral market failures, I have raised several relevant questions here. But to the extent that the measurement is accurate, the argument for use of VSL, based on WTP, is generally secure.

In the hard cases, the analysis is more complicated. Use of VSL may lead to cost-benefit analyses that do not capture the welfare effects of regulations. In some cases, those effects may be positive even if a regulation has net monetary costs. In addition, distributional effects deserve consideration, and they may support regulation even in the face of net costs (consider prioritarianism). At the same time, regulators do not have good tools for measuring welfare effects directly, and consideration of distributional effects may also create serious challenges. In these circumstances, the best approach seems to be to proceed as suggested by cost-benefit analysis, but to allow departures when there are compelling reasons to believe that rules nonetheless increase welfare or are strongly supported by distributional considerations.

Acknowledgements: This Article has benefited from the author's experience as Administrator of the White House Office of Information and Regulatory Affairs from 2009 to 2012. Parts of the Essay draw on, but significantly revise, portions of Cass R. Sunstein, *Valuing Life: A Plea for Disaggregation*, 54 *Duke LJ* 385 (2004). A version will appear in Cass R. Sunstein, *Valuing Life: Humanizing the Regulatory State* (University of Chicago Press, 2013). I am grateful to Matthew Adler, Joseph Aldy, Martha Nussbaum, Eric Posner, and Lisa Robinson for valuable comments on a previous draft. Cassie Chambers provided superb research assistance.

Previously published online July 19, 2013

References

- Adler, M. D. (2008). Risk equity: a new proposal. *Harvard Environmental Law Review* 32, 1. Available at: Available at SSRN: <http://ssrn.com/abstract=981003> or <http://dx.doi.org/10.2139/ssrn.981003>.
- Adler, M. D. (2011). *Well-being and fair distribution*. New York: Oxford University Press.
- Adler, M. D. (2013). Happiness surveys and public policy: what's the use? *Duke Law Journal* 62, 1509–1601.
- Adler, M., & Posner, E. (1999). Implementing cost-benefit analysis when preferences are distorted. *U Chicago Law & Economics, Olin Working Paper*, (88).
- Adler, M. D., & Posner, E. A. (2006). *New foundations of cost-benefit analysis*. Cambridge: Harvard University Press.
- Akerlof, G. A., & Dickens, W. T. (1984). The economic consequences of cognitive dissonance. In: G. A. Akerlof (Ed.) *An economic theorist's book of tales*. New York: Cambridge University Press.
- Aldy, J., & Viscusi, K. W. (2007). Age differences in the value of statistical life. *Review of Environmental Economics and Policy* 1(2), 241–260.
- Arneson, R. J. (2000). Luck egalitarianism and prioritarianism. *Ethics* 110(2), 339–349.
- Ashenfelter, O., & Greenstone, M. (2002). Using mandated speed limits to measure the value of a statistical life. Available at <http://www.nber.org/papers/w9094>.
- Benjamin, D. J., Kimball, M. S., Heffetz, O., & Rees-Jones, A. (2012). What do you think would make you happier? What do you think you would choose? *The American Economic Review* 102(5), 2083.
- Bronsteen, J., Buccafusco, C., & Masur, Jonathan S. (2013). Well-being analysis versus cost-benefit analysis. *Duke Law Journal* 62, 1603–1689.
- Conly, S. (2012). *Against autonomy*. New York: Cambridge University Press.
- Cropper, M., Hammitt, J., & Robinson, L. (2011). Valuing mortality risk reductions: progress and challenges. *Annual Review of Resource Economics* 3, 313–336.
- Department of Transportation. (DOT). (2013). Available at: <http://www.dot.gov/sites/dot.dev/files/docs/DOT%202013%20Signed%20VSL%20Memo.pdf>.
- de Tocqueville, A. D. (1848). *1969, Democracy in America* Trans. by George Lawrence. Garden City, NY: Anchor Books.
- Diamond, P. A., & Hausman, J. A. (1994). Contingent valuation: is some number better than no number? *The Journal of Economic Perspectives* 8(4), 45–64.
- Dorman, P. (1996). *Markets and mortality*. Great Britain: Cambridge University Press.
- Dorman, P., & Hagstrom, P. (1998). Wage compensation for dangerous work revisited. *Industrial & Labor Relations Review* 52, 116–135.
- Dunn, E., & Norton, M. (2013). *Happy money*. New York: Simon and Schuster.
- Dworkin, R. (2002). *Sovereign virtue*. Cambridge: Harvard University Press.
- Elster, J. (1985). *Sour grapes: studies in the subversion of rationality*. New York: Cambridge University Press.
- Elster, J., & Roemer, J. E. (Eds.). (1993). *Interpersonal Comparisons of Well-Being*. Cambridge University Press.
- Environmental Protection Agency (EPA) (2000). Agency Guidelines for Preparing Economic Analyses. Available at: [http://yosemite.epa.gov/ee/epa/erm.nsf/vwAN/EE-0228C-07.pdf/\\$file/EE-0228C-07.pdf](http://yosemite.epa.gov/ee/epa/erm.nsf/vwAN/EE-0228C-07.pdf/$file/EE-0228C-07.pdf).

- Environmental Protection Agency (EPA) (2010). Valuing Mortality Risk Reductions for Environmental Policy: A White Paper (Review Draft). Prepared by the National Center for Environmental Economics for consultation with the Science Advisory Board – Environmental Economics Advisory Committee.
- Gilbert, D. T., & Wilson, T. D. (2000). Miswanting. In: J. P. Forgas (Ed.), *Feeling and thinking*. New York: Cambridge University Press.
- Graham, C. (2013). An economist's perspective on well-being analysis and cost-benefit analysis. *Duke Law Journal* 62, 1691–1700.
- Hedbergian, R. J., Gray, W., & Morgan, C. (2007). Benefits and costs from sulfur dioxide trading: a distributional analysis. In: G. R. Visgilio and D. M. Whitelaw (Eds.), *Acid in the environment: lessons learned and future prospects*. New York, New York: Springer.
- Hicks, J. R. (1941). The rehabilitation of consumers' surplus. *The Review of Economic Studies* 8(2), 108–116.
- Jones-Lee, M. W. (1991). Altruism and the value of other people's safety. *Journal of Risk and Uncertainty* 4(2), 213–219.
- Jones-Lee, M. W. (1992). Paternalistic altruism and the value of statistical life. *The Economic Journal* 102(410), 80–90.
- Kahn, M. (2001). The beneficiaries of clean air act regulation. *Regulation*, 24, 34–38.
- Kahneman, D. (2011). *Thinking, fast and slow*. New York: Farrar, Straus and Giroux.
- Kahneman, D. (2003). Maps of bounded rationality: psychology for behavioral economics. *The American Economic Review* 93(5), 1449–1475.
- Kahneman, D., & Varey, C. (1991). Notes on the psychology of utility. In: J. Elster and J. Roemer (Eds.), *Interpersonal comparisons of utility*. New York: Cambridge University Press.
- Kahneman, D., Wakker, P. P., & Sarin, R. (1997). Back to bentham? Explorations of experienced utility. *The Quarterly Journal of Economics*, 112(2), 375–406.
- Kaplow, L., & Shavell, S. (1994). Why the legal system is less efficient than the income tax in redistributing wealth. *Journal of Legal Studies* 23, 667–681.
- Karlsson, N., Loewenstein, G., & McCafferty, J. (2004). The economics of meaning. *Nordic Journal of Political Economy* 30(1), 61–75.
- Khader, S. J. (2011). *Adaptive preferences and women's empowerment*. OUP USA.
- Kling, C. L. & Swackhamer, D. L. (2011). Review of 'Valuing Mortality Risk Reductions for Environmental Policy: A White Paper' (December 10, 2010). Memorandum to Lisa P. Jackson, EPA Administrator, from the EPA Science Advisory Board and Environmental Economics Advisory Committee. EPA-SAB-11-011.
- Kniesner, T. J., Viscusi, K. W. & Ziliak, J. P. (2012). Willingness to accept equals willingness to pay for labor market estimates of the value of statistical life. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2221038.
- Korobkin, R. (2003). The endowment effect and legal analysis. *Northwestern University Law Review* 97(3), 1227–1293.
- Leung, J., & Guna, J. (2006). Value of statistical life: adults versus children, *Accid Anal. Prev.* 38(6), 1208–1217.
- Lindhjem, H., S. Navrud, N. A. Braathen, & V. Biousque. (2011). Valuing mortality risk reductions from environmental, transport, and health policies: a global meta-analysis of stated preference studies. *Risk Analysis* 31(9), 1381–1407.
- Livermore, M. A., & Rosenberg, J. S. (2013). The shape of distributional analysis. *The Globalization of Cost-Benefit Analysis in Environmental Policy* 5, 69–84.
- Mill, J. S. (1859). *On liberty*. (Kathy Casey ed., 2002). New York: Dover.

- Miller, T. R. (2000). Variations between countries in value of a statistical life. *Journal of Transport Economics and Policy* 34, 169–188.
- Posner, E. A., & Sunstein, C. R. (2005). Dollars and death. *The University of Chicago Law Review* 537–598.
- Robinson, L. A., & Hammitt, J. K. (2011). Valuing health and longevity in regulatory analysis: current issues and challenges, In: *The Handbook on the Politics of Regulation* D. Levi-Faur (Ed.), Cheltenham and Northampton: Edward Elgar. Available at www.regulatory-analysis.com.
- Robinson, L. A., & J. K. Hammitt. (2013). Skills of the trade: valuing health risk reductions in benefit-cost analysis. *Journal of Benefit-Cost Analysis* 4(1), 107–130.
- Sen, A. (2002). *Rationality and freedom*. Cambridge: Harvard University Press.
- Sharot, T. (2010). *The optimism bias*. New York: Random House.
- Shavell, S. (1981). A note of efficiency vs. distributional equity in legal rulemaking: should distributional equity matter given optimal income taxation? *The American Economic Review* 71(2), 414–418.
- Sunstein, C. R. (2002a). *Risk and reason*. New York: Cambridge University Press.
- Sunstein, C. R. (2002b). Probability neglect. *The Yale Law Journal* 112(1), 61–107.
- Sunstein, C. R. (2013a). The storrs lectures: behavioral economics and paternalism. *Yale Law Journal*, 113, Available at: <http://ssrn.com/abstract=2182619> or <http://dx.doi.org/10.2139/ssrn.2182619>.
- Sunstein, C. R. (2013b). *Simpler*. New York: Simon and Schuster.
- Sunstein, C. R. (2014). The real world of cost-benefit analysis. *Colum L Rev* (forthcoming 2014).
- Thaler, R., & Rosen, S. (1976). The value of saving a life. Available at <http://www.nber.org/chapters/c3964.pdf>.
- Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: improving decisions about health, wealth, and happiness*. New Haven: Yale University Press.
- Ubel, P. A., & Loewenstein, G. (2008). Pain and suffering awards: they shouldn't be (just) about pain and suffering. *The Journal of Legal Studies* 37(S2), S195–S216.
- Viscusi, K. W. (1992). *Fatal tradeoffs*. New York: Oxford University Press.
- Viscusi, W. (2010). The heterogeneity of the value of statistical life. *Journal of Risk and Uncertainty* 40(1), 1–13.
- Viscusi, K. W., & Aldy, J. (2003). The value of a statistical life. *Journal of Risk and Uncertainty* 27(1), 5–76.
- Viscusi, K. W. (2013). The benefits of mortality risk reduction. *The Duke Law Journal* 62, 1735–1745.
- Weisbach, D. A. (2003). Should legal rules be used to redistribute income? *The University of Chicago Law Review* 70(1), 439–453.
- Williams, S. H. (2013). Statistical children. *Yale Journal of Regulation* 30, 63–125.
- Wilson, T. D., & Gilbert, D. T. (2003). Affective forecasting. In: M. P. Zanna (Ed.), *Advances in experimental social psychology*. New York: Academic Press.