

January 2018

Moral Context and Risks of Death

Dov Waisman
Southwestern Law School

Follow this and additional works at: <https://scholarworks.uark.edu/alr>



Part of the [Labor and Employment Law Commons](#)

Recommended Citation

Dov Waisman, *Moral Context and Risks of Death*, 71 Ark. L. Rev. 215 (2018).

Available at: <https://scholarworks.uark.edu/alr/vol71/iss1/4>

This Article is brought to you for free and open access by ScholarWorks@UARK. It has been accepted for inclusion in Arkansas Law Review by an authorized editor of ScholarWorks@UARK. For more information, please contact scholar@uark.edu, ccmiddle@uark.edu.

MORAL CONTEXT AND RISKS OF DEATH

Dov Waisman*

When an industry poses a risk of premature death to consumers, workers, or others, regulatory agencies employ a figure known as the “value of a statistical life” (VSL) to monetize the life-saving benefit of regulations designed to reduce that risk. Use of the VSL, which currently hovers around \$9 million, has been highly controversial. While a number of prominent scholars have vigorously endorsed the VSL as necessary to the cost-benefit analysis of mortality risk regulations, other prominent scholars have vehemently rejected the very idea of attaching a monetary value to a statistical human life.

This article stakes out a novel and more nuanced position based on a largely neglected aspect of mortality risk regulation: moral context. Consumption risks—risks of death associated with using or consuming a particular product—typically fall on consumers who not only benefit meaningfully from the industry but who also bear all or substantially all of the costs of risk-reducing regulations. Using a VSL to guide risk regulation in this moral context is defensible on the basis of the norm of personal autonomy. By contrast, workplace risks—risks of death associated with employment in a particular industry—typically fall on workers who benefit from the industry but who do not bear the costs of risk-reducing regulations. In this moral context, using a VSL to guide risk regulation is not normatively defensible. However, using the underlying economic concept of willingness-to-pay to guide the regulation of workplace mortality risks is defensible on the basis of the norm of equity.

*Professor of Law, Southwestern Law School. A.B. Harvard College. M.A., Philosophy, U.C. San Diego. J.D., U.C.L.A. I am indebted to Al Calnan, Michael Dorff, Bryant Garth, Warren Grimes, Danni Hart, Jeff Helmreich, Aaron James, Greg Keating, Hila Keren, Gowri Ramachandran, Ken Simons, Byron Stier, John Tehranian, and Tracy Turner for their feedback and support. Of course, any errors are mine.

TABLE OF CONTENTS

Introduction.....	217
I. Factual and Theoretical Background.....	224
A. How the Value of a Statistical Life is Calculated	224
B. Normative Pluralism	226
II. Consumption Risks and Personal Autonomy.....	229
A. Example: Arsenic in Drinking Water.....	230
B. Welfare Maximization Fails to Justify Use of a VSL....	233
C. The Norm of Personal Autonomy.....	243
D. Personal Autonomy Justifies Use of a VSL	245
1. <i>A Stylized Consumption Risk Scenario: Arsenic (1)</i>	245
2. <i>A More Realistic Scenario: Arsenic (100,000)</i>	249
III. Workplace Risks and Equity.....	255
A. Example: Chromium in the Workplace	257
B. The Norm of Equity and (Ex Ante) Contractualism	260
C. Equity Justifies Use of Willingness-to-Pay, But Not a VSL.....	268
1. <i>A Stylized Workplace Risk Scenario: Chromium (2)</i>	268
2. <i>A More Realistic Scenario: Chromium (200,000)</i> ..	278
3. <i>Another Realistic Scenario: Chromium (1.1 million)</i>	289
Conclusion	293

INTRODUCTION

Operating power plants, refining oil into gasoline, constructing large buildings, manufacturing pharmaceuticals, pesticide-based agriculture—these are all prominent examples of industrial activities that, despite their important social benefits, pose significant risks of premature death to human beings. Indeed, one of the defining features of the modern world has been the proliferation of industries that deliver profound, life-improving benefits to millions at the cost of imposing low, but nontrivial, risks of death on consumers, workers, or bystanders. Not surprisingly, the twentieth century saw the creation of an array of federal agencies tasked with regulating industrial mortality risks in their respective jurisdictions.¹

Since the first Reagan administration, the dominant theoretical framework for federal risk regulation has been cost-benefit analysis (CBA).² According to CBA, a regulation should be issued only if its total benefits exceed its total costs,³ and regulatory levels (e.g., permissible exposure levels to toxic substances) should be set with the goal of maximizing net benefits.⁴ Because standard CBA is founded on a quantitative comparison of costs to benefits, an obvious problem arises when the costs and benefits of a regulation take different forms. When the costs of a regulation are monetary, but the benefits consist in

1. The Environmental Protection Agency (EPA), National Transportation and Highway Safety Administration (NHTSA), Occupational Health and Safety Administration (OSHA), and Food and Drug Administration (FDA) are the most prominent ones.

2. See Exec. Order No. 12291, 46 Fed. Reg. 13193 (Feb. 19, 1981) (requiring federal agencies to perform CBA for all major regulations). Similar orders have been issued by subsequent administrations. See, e.g., Exec. Order No. 12866, 58 Fed. Reg. 51735 (Oct. 4, 1993) (President Clinton); Exec. Order No. 13563, 76 Fed. Reg. 3821 (Jan. 21, 2011) (President Obama).

3. See MATTHEW D. ADLER & ERIC A. POSNER, *NEW FOUNDATIONS OF COST-BENEFIT ANALYSIS 2* (2006) (“To simplify greatly, cost-benefit analysis (CBA) requires the regulatory agency to sum up the costs and benefits of a proposed regulation, and issue the regulation if the benefits exceed the costs.”). For a more technical statement of this point, see MATTHEW D. ADLER, *WELL-BEING AND FAIR DISTRIBUTION: BEYOND COST-BENEFIT ANALYSIS 88-91* (2012).

4. See Michael A. Livermore & Richard L. Revesz, *Rethinking Health-Based Environmental Standards*, 89 N.Y.U. L. REV. 1184, 1190 (2014) (“Cost-benefit analysis, in its most general form, places both costs and benefits along a common metric and supports the standard that maximizes net benefits (the difference between benefits and costs).”); cf. ADLER & POSNER, *supra* note 3, at 68 (noting CBA is “a plausible candidate to be the welfare-maximizing procedure as against currently available alternative procedures”).

the prevention of some number of premature deaths, how do regulators decide whether the regulation is cost-justified? The answer is that regulators convert the life-saving benefits of the regulation into the same currency as its costs by monetizing the benefit of avoiding premature death.⁵ When the benefits of a regulation include the reduction of a mortality risk—and, therefore, the avoidance of some number of premature deaths that would otherwise have been statistically expected to occur—regulators have conducted CBA with reference to what has come to be known as the *value of a statistical life* (VSL).⁶ The VSL—which currently hovers around \$9 million⁷—represents the primary vehicle by which federal regulatory agencies have monetized the benefits of regulations that reduce risks of death.

The use of the VSL to guide mortality risk regulation has been highly controversial. A number of prominent scholars have vehemently rejected the very idea of attaching a monetary value

5. See FRANK ACKERMAN & LISA HEINZERLING, PRICELESS: ON KNOWING THE PRICE OF EVERYTHING AND THE VALUE OF NOTHING 61 (2004) (“The most significant benefits of environmental protection are often the deaths prevented by regulation. To decide whether the benefits of regulation are larger or smaller than the costs, it is essential to assign a dollar value to lives saved.”); CASS R. SUNSTEIN, VALUING LIFE: HUMANIZING THE REGULATORY STATE 85 (2014) (“We have seen that in order to conduct cost-benefit analysis, agencies must assign monetary values to the human lives that would be saved by a proposed regulation.”); Thomas O. McGarity, *A Cost-Benefit State*, 50 ADMIN. L. REV. 7, 63 (1998) (“[P]erhaps the most significant objection to cost-benefit analysis is the inability of economic analysis to reduce the benefits of regulation to dollar equivalents to compare with regulatory costs.”).

6. See, e.g., ACKERMAN & HEINZERLING, *supra* note 5, at 67 (“The standard economic response is that a value like \$6.1 million is not a price on an individual’s life or death. Rather, it is a way of expressing the value of small risks of death, which, when aggregated to produce one death, can be called a ‘statistical life’; for example, the value of one statistical life is one million times the value of a one in a million risk.”); W. Kip Viscusi, *The Value of Life in Legal Contexts: Survey and Critique*, 2 AM. L. & ECON. REV. 195, 196-97 (2000) (“Economic discussions of the value of life almost invariably focus on the value of a statistical life, considering an individual facing a very small probability of death. . . . [T]he value of a statistical life is a prospective measure that in effect establishes the appropriate price society is willing to pay for small risk reductions.”).

7. See W. Kip Viscusi & Joseph E. Aldy, *The Value of a Statistical Life: A Critical Review of Market Estimates Throughout the World*, 27 J. RISK & UNCERTAINTY 5, 18 (2003) (concluding in a meta-study that most studies determine the VSL to lie in a range between \$3.8 and \$9 million); see also SUNSTEIN, *supra* note 5, at 94 (“The most detailed meta-study, [Viscusi & Aldy 2003,] far more comprehensive than the EPA’s own analysis, identifies a central value in the general vicinity of \$8 million and finds that most studies produce VSLs ranging from \$3.8 million to \$9 million. . . . (Note that on the basis of their reading of the technical literature, agencies actually use a narrow range of \$6 million to \$9 million, with increasing consensus in the vicinity of \$9 million.)”).

to a statistical life.⁸ An equally prominent group of scholars has vigorously defended the use of the VSL to guide mortality risk CBA,⁹ often while acknowledging the need for further work to develop and refine the VSL calculation.¹⁰

This article stakes out a novel and more nuanced position based on a largely neglected aspect of mortality risk regulation: *moral context*. Whether and how the VSL is used to guide

8. See, e.g., ACKERMAN & HEINZERLING, *supra* note 5, at 61, 67 (“Putting a price on human life makes most people uncomfortable. It is clearly unacceptable to virtually all religions and moral philosophies. Nonetheless, the quantitative valuation of life has become central to recent analyses of public policies. . . . A more careful restatement of the ethical objection is that there is no ‘price’ for life because its value is immeasurable.”); DOUGLAS A. KYSAR, REGULATING FROM NOWHERE: ENVIRONMENTAL LAW AND THE SEARCH FOR OBJECTIVITY 111-13 (2010) (critiquing the method by which CBA monetizes human life as “driven by the procedure’s purely individualistic conception of value” and “tolerat[ing] treatment of statistical victims that would be considered well nigh criminal in ordinary contexts”); John Broome, *Trying to Value a Life*, 9 J. PUB. ECON. 91, 95 (1978); David M. Driesen, *Distributing the Costs of Environmental, Health, and Safety Protection: The Feasibility Principle, Cost-Benefit Analysis, and Regulatory Reform*, 32 B.C. ENVTL. AFF. L. REV. 1, 53 (2005) (“CBA proponents contemplate comparison of costs to benefits. To facilitate comparison, analysts attempt to place a dollar value on the averted harms, such as deaths and illnesses. The methodologies for doing this involve numerous highly questionable value assumptions.”); see also W. Kip Viscusi, *How to Value a Life*, 32 J. ECON. & FIN. 311, 322 (2008) (noting “many non-economists continue to attack the entire concept of monetizing risks to life”).

9. See, e.g., ADLER & POSNER, *supra* note 3, at 178-82 (“Pricing life does not ‘cheapen’ life. There are numerous contexts, other than CBA, in which premature death or the risk of premature death is priced”); STEPHEN BREYER, BREAKING THE VICIOUS CIRCLE: TOWARD EFFECTIVE RISK REGULATION 16 (1993) (“[E]very day, each of us implicitly evaluates risks to life. . . . We find it worth spending money on an ordinary fire alarm system, but not worth installing state-of-the-art automatic-phone-dialing fire protection. We believe it worth installing guard rails on bridges, but not worth coating the Grand Canyon in soft plastic to catch those who might fall over the edge.”); SUNSTEIN, *supra* note 5, at 7 (“[F]ar from being preposterous, efforts to value human life (more accurately, statistical mortality risks) are rooted in exceedingly appealing ideas about welfare and autonomy—ideas that deserve a prominent place in a free society.”); Viscusi, *supra* note 8, at 311 (“[T]he economic approach to valuing life, or more specifically, valuing risks to life, is quite different and more legitimate than an accounting procedure based on one’s income.”).

10. See, e.g., ADLER & POSNER, *supra* note 3, at 181 (“A different problem with VSL, as currently practiced, involves disaggregation. Agencies almost always employ a single VSL figure, regardless of the number of life-years saved or lost or other individual characteristics. Perfect welfarist measurement would be more disaggregated.”); SUNSTEIN, *supra* note 5, at 95-106 (arguing VSL should be individuated across different types of mortality risks and differently-situated persons); W. Kip Viscusi, *Policy Challenges of the Heterogeneity of the Value of Statistical Life*, 6 FOUND. & TRENDS IN MICROECONOMICS 99, 117-120 (2010) (discussing variability of VSL based on the magnitude of the mortality risk and age of those exposed to the risk); see also Cass R. Sunstein, *Valuing Life: A Plea for Disaggregation*, 54 DUKE L.J. 385, 389 (2004).

mortality risk regulation should depend on two morally significant features of the risk imposition in question: (1) the extent to which those exposed to the risk benefit from the industrial activity that gives rise to the risk; and (2) the extent to which those exposed to the risk bear the costs of compliance with the risk-reducing regulation. These two features vary in characteristic ways depending on the type of risk imposition at issue. This article distinguishes broadly between two types of mortality risk impositions—*consumption risks* and *workplace risks*—and makes a claim about the justifiability of using the VSL to guide risk regulation in each case.¹¹

As explained in Part I, the VSL used in federal risk regulation is based on the general economic concept of willingness-to-pay and, more specifically, on empirical studies of the amount of money individuals are willing to pay to avoid (or the amount of money they are willing to accept to assume) very small mortality risks. The basic calculation goes as follows. Suppose labor market and consumer studies show that, on the average, individuals are willing to pay \$90 to eliminate a 1 in 100,000 risk of death to which they would otherwise be subject. This implies that, were the risk imposed on 100,000 individuals, the collective willingness-to-pay to eliminate the risk, and thereby avoid one statistical death, would be \$9 million (\$90 multiplied by 100,000). On this basis, the value of a statistical life is determined to be \$9 million.¹²

Armed with this information, CBA can be applied to mortality risk regulation. If a regulation that would cost \$20

11. For a brief discussion of the moral context created by a third type of mortality risk imposition—environmental risks—see *infra* text accompanying note 274. In future work, I plan to explore the justifiability of using the VSL to guide the regulation of such risks. The present article focuses on consumption and workplace risks.

12. See, e.g., SUNSTEIN, *supra* note 5, at 51 (“With these values, the government is not actually ‘valuing life.’ It is valuing the reduction of mortality risks—typically by eliminating low-level risks, for example, risks of 1 in 100,000. When it is said that a life is ‘worth’ \$9 million in such cases, what is really meant is that people are willing to pay, or ask to be paid, \$90, on average, to eliminate a risk of 1 in 100,000.” (footnote omitted)); W. Kip Viscusi, *The Heterogeneity of the Value of Statistical Life: Evidence and Policy Implications*, in *BENEFIT-COST ANALYSES FOR SECURITY POLICIES: DOES INCREASED SAFETY HAVE TO REDUCE EFFICIENCY* 78, 81 (Carol Mansfield & V. Kerry Smith eds., 2015) (“The VSL pertains to the trade-off between money and very small risks of death. Suppose that there is a risk of one chance in 10,000 to 10,000 people, so that this group will experience one expected death. If each person would be willing to pay \$800 to eliminate the risk, the VSL in this instance would be \$8 million, or 10,000 people x \$800 per person.”).

million is expected to save a single statistical life by eliminating a 1 in 100,000 risk of death to 100,000 people, CBA would reject the regulation on the grounds that it is not cost-justified, since the monetized value of the regulation's life-saving benefit is \$9 million, less than half of its cost. CBA would reach this same conclusion regardless of the moral context of the risk imposition. By the lights of CBA, the conclusion would follow with equal force regardless of whether those exposed to the risk benefit from the activity that gives rise to the risk, and regardless of whether those exposed to the risk bear the cost of complying with risk-reducing regulations.

This one-size-fits-all approach to using the VSL in mortality risk regulation is a fundamental error. Whether and how the VSL is deployed should vary with the moral context.

The VSL is most defensible in the context of consumption risks—risks associated with using or consuming a particular good or service.¹³ Such risks typically fall on consumers who not only benefit meaningfully from the industrial activity giving rise to the risk, but who also bear all or substantially all of the costs of risk-reducing regulations.¹⁴ For example, the city residents exposed to a small risk of death from drinking city water containing trace amounts of arsenic both benefit from the activity that puts them at risk (public provision of drinking water) and, by paying higher water bills reflecting passed through regulatory costs, bear the costs of complying with EPA regulations limiting the amount of arsenic in the water.¹⁵

As explained in Part II, using the VSL to guide mortality risk CBA in such circumstances is defensible on the basis of the norm of *personal autonomy* (though not, as is commonly asserted, on the basis of the norm of welfare maximization).¹⁶ Where the cost of complying with a risk-reducing regulation is borne by, and divided equally among, the same group of people that is exposed

13. See Thomas J. Kniesner & W. Kip Viscusi, *Value of a Statistical Life: Relative Position vs. Relative Age*, 95 AM. ECON. REV. 142, 142 (2005).

14. See *id.* at 145.

15. See generally EPA, No. 815-R-00-026, *Arsenic in Drinking Water Rule: Economic Analysis* (2000) [hereinafter EPA Analysis].

16. Cass Sunstein has argued that using a VSL as a guide for mortality risk regulation in such circumstances (which Sunstein calls “[e]asy [c]ases”) is defensible on the basis of both autonomy and welfare. See SUNSTEIN, *supra* note 5, at 113-15. For reasons explained in Part II.B, the welfare-based rationale appears to be open to serious objection.

to the risk, honoring the risk-bearers' autonomy requires rejecting any regulation whose total cost exceeds the sum the risk-bearers would collectively be willing to pay to eliminate the risk. Enacting such a regulation would amount to forcing each risk-bearer to invest more in their own safety than they would rationally wish to, thereby violating each risk-bearer's right to be sovereign in matters affecting their interests alone.

By contrast, the VSL is more problematic in the context of workplace risks—risks associated with employment in producing a particular good or service.¹⁷ Such risks typically fall on workers who benefit meaningfully from the industrial activity giving rise to the risk (because it provides them with a livelihood), but who bear no appreciable share of the costs of risk-reducing regulations.¹⁸ These costs are typically passed through to consumers or deducted from shareholder profits.¹⁹ For example, the electroplating workers exposed to a small risk of death from exposure to hexavalent chromium (a carcinogenic chemical used in a number of industries) benefit significantly from the existence of the electroplating industry, but are unlikely to bear a substantial share of the costs of compliance with OSHA regulations limiting their exposure to hexavalent chromium.²⁰

This changes the moral context in an important way. As explained in Part III, using the general concept of willingness-to-pay to guide risk regulation in such circumstances is morally defensible on the basis of the norm of *equity*, rather than personal autonomy.²¹ Suppose the workplace chromium regulation discussed above would eliminate a 1 in 100,000 risk of death to 100,000 workers. Suppose further that its \$20 million cost would be evenly distributed among 1 million consumers in the form of a \$20-per-unit price increase. Standard CBA is insensitive to whether the \$20 million cost of saving one statistical life is

17. Kniesner & Viscusi, *supra* note 13, at 142.

18. See W. Kip Viscusi, *Monetizing the Benefits of Risk and Environmental Regulation*, 33 FORDHAM URB. L.J. 1003, 1022-23 (2006).

19. Dov Waisman, *Equity and Feasibility Regulation*, 50 U. RICH. L. REV. 1263, 1266 (2016).

20. See generally Occupational Exposure to Hexavalent Chromium, 71 Fed. Reg. 10100 (Feb. 28, 2006) (to be codified at 29 C.F.R. pts. 1910, 1915, 1917, 1918, & 1926).

21. Cf. Waisman, *supra* note 19, at 1265-66, 1269-73 (arguing that the norm of equity provides the basis for feasibility-based regulation, which requires mortality risks to be reduced to the maximum extent technologically and economically feasible).

distributed among the same 100,000 people subject to the 1 in 100,000 mortality risk or among 1 million people who are not subject to that risk. In either case, the regulation's \$9 million benefit cannot justify its \$20 million cost.

From the standpoint of equity, however, the manner in which costs and risks are distributed is significant. In the case of the chromium regulation, a safety benefit worth \$90 to each worker could be provided at a cost of just \$20 to each cost-bearing consumer. Equity—a norm which attends to and compares the burdens and benefits experienced by each affected individual as the result of the action or policy under consideration—would weigh strongly in favor of the regulation,²² because it would arguably be unreasonable for each consumer to reject having to pay a \$20 cost so that each worker can avoid a risk imposition that is equivalent to a \$90 cost. In these circumstances, equity supports enacting the regulation even though, by the lights of standard CBA, it is not cost-justified.

This equity-based line of reasoning embodies a modified, individualized form of CBA, one that attends to the risks and costs experienced by *each affected risk-bearer and cost-bearer*, rather than to risks and costs considered in the aggregate.²³ This individualized form of CBA relies directly on the individual willingness-to-pay data upon which the VSL is based, rather than on the collective willingness-to-pay calculation which the VSL itself reflects.²⁴ When it comes to workplace mortality risks, there does not seem to be a normative basis for conducting CBA with reference to the VSL. Instead, standard VSL-based CBA should be subordinate to the individualized form of CBA described in this article.

22. *See id.* at 1267.

23. *See* Dov Waisman, *Reasonable Precaution for the Individual*, 88 ST. JOHN'S L. REV. 653, 676-79 (2014) (introducing and defending an "individual risk principle" that determines regulatory levels based on a comparison between each cost-bearer's share of regulatory costs and each risk-bearer's resulting reduction in mortality risk).

24. *Id.* at 697-98.

I. FACTUAL AND THEORETICAL BACKGROUND

A. How the Value of a Statistical Life is Calculated

Suppose a given regulation is expected to save 100 lives at a cost of \$1 billion. Do the life-saving benefits of the regulation justify its monetary cost? At least since President Reagan's 1981 Executive Order directing federal regulatory agencies to engage in cost-benefit analysis of all major regulations,²⁵ regulators have been faced with the difficult task of answering such questions.

The approach regulatory agencies have taken is to convert the health benefits of such a regulation—primarily, deaths avoided—into the same currency as the regulation's cost—dollars.²⁶ In the early 1980s, the notion of a “value of statistical life” was born, and has since been used pervasively by federal agencies charged with regulating risks of death or serious bodily harm.²⁷

How exactly do regulatory agencies calculate the VSL? This question presents a dilemma. On the one hand, the view that a statistical death should be prevented at *any* monetary cost strikes many as implausible.²⁸ On the other hand, using willingness-to-pay or willingness-to-accept as a method for attaching a monetary value to life itself is problematic because, except in very unusual circumstances, no sane person would be willing to forfeit his or her life in exchange for any amount of money.²⁹ How have economists resolved this dilemma? “The answer,” explains Cass Sunstein in a 2005 book, “involves real-world markets, producing evidence of compensation levels for actual risks.”³⁰ Moreover, Sunstein previously suggested,

25. See Exec. Order No. 12291, 46 Fed. Reg. 13193 (Feb. 19, 1981).

26. See Eric A. Posner & Cass R. Sunstein, *Dollars and Death*, 72 U. CHI. L. REV. 537, 538 (2005).

27. See *id.* at 538 n.9; see also Waisman, *supra* note 19, at 1299.

28. See, e.g., BREYER, *supra* note 9, at 16 (“[E]very day each of us implicitly evaluates risks to life. We begin to run risks to achieve our daily objectives the instant we get out of bed. . . . We believe it worth installing guard rails on bridges, but not worth coating the Grand Canyon in soft plastic to catch those who might fall over the edge.”).

29. See, e.g., ACKERMAN & HEINZERLING, *supra* note 5, at 68-69; Broome, *supra* note 8, at 92, 95; David M. Driesen, *Two Cheers for Feasible Regulation: A Modest Response to Masur and Posner*, 35 HARV. ENVTL. L. REV. 313, 340 (2011).

30. CASS R. SUNSTEIN, LAWS OF FEAR: BEYOND THE PRECAUTIONARY PRINCIPLE 132 (2005); see also Viscusi, *supra* note 10, at 104 (“Suppose that there is a risk of one

In the workplace and for consumer goods, additional safety has a ‘price’; market evidence, involving the compensation people actually receive, is investigated to identify that price. Agency valuations are largely a product of studies of workplace risks, attempting to determine how much workers are paid to assume mortality hazards. The relevant risks usually are in the general range of 1/10,000 to 1/100,000. The calculation of the value of a statistical life (VSL) is a product of simple arithmetic. Suppose . . . that workers must be paid \$[9]00, on average, to eliminate a risk of 1/10,000. If so, the value of a statistical life would be said to be \$[9] million.³¹

It is important to bear in mind that this method of attaching a monetary value to human life is, in one very clear sense, based on a fiction.³² There is a huge inferential leap from the empirical proposition that an average worker would require a payment of \$900 to be exposed to 1 in 10,000 risk of death to the proposition that a statistical worker’s *life* has a value of \$9 million.³³ Sunstein, in fact, explicitly rejects the latter proposition:

With these values, the government is not actually ‘valuing life.’ It is valuing the reduction of mortality risks—typically by eliminating low-level risks, for example, risks of 1 in 100,000. When it is said that a life is ‘worth’ \$9 million in such cases, what is really meant is that people are willing to pay, or ask to be paid, \$90, on average, to eliminate a risk of 1 in 100,000.³⁴

chance in 10,000 to 10,000 people so that this group will experience one expected death. If each person would be willing to pay \$800 to eliminate the risk, the VSL in this instance would be \$8 million, or 10,000 people x \$800 per person. This is the amount that could be raised to prevent one expected death.”).

31. Posner & Sunstein, *supra* note 26, at 551 (footnotes omitted) (in the last two sentences of the quote, the original text uses figures of \$600 and \$6 million). Professor W. Kip Viscusi is generally credited with introducing the practice of calculating VSL based on willingness-to-pay figures derived from actual labor market behavior. *See, e.g.*, W. Kip Viscusi, *Risk Equity*, 29 J. LEGAL STUD. 843, 845-46 (2000).

32. Posner & Sunstein, *supra* note 26, at 544.

33. *See* Viscusi, *supra* note 10, at 104.

34. SUNSTEIN, *supra* note 5, at 51 (footnote omitted). Sunstein further notes, “it would be grossly misleading to offer the following suggestion: *The value of a statistical life is \$9 million*. It would be much more accurate to say that for risks of 1/10,000, the median WTP in the relevant population is \$900—or that for risks of 1/100,000, the median WTP is \$90.” *Id.* at 95.

Obviously, it cannot be inferred that an average worker would forfeit their life in exchange for a monetary payment of *any* amount of money, much less a payment of \$9 million.³⁵ To generalize this point, it cannot be inferred that the amount a worker would be willing to pay to eliminate a risk of death would increase in a linear fashion as the risk itself increases. Rather, a person's willingness-to-pay to avoid a risk of death increases in a nonlinear manner as the magnitude of the risk increases.³⁶ For example, although a worker might be willing to pay only \$900 for the elimination of a 1 in 10,000 risk of death, she might well be willing to pay more than ten times as much to eliminate a risk ten times as great.³⁷ Judge Richard Posner has put this point as follows: “[C]ost-benefit analysis values risks, and the so-called value of life that cost-benefit analysts refer to is just a mathematical transformation.”³⁸

B. Normative Pluralism

The theoretical basis of this article's argument is a moral theory known as *normative pluralism*. Normative pluralists hold that moral judgments must be understood to rest on multiple norms and deny that any one norm—e.g., maximizing overall well-being—can adequately explain the moral rightness or wrongness of all actions.³⁹ Normative pluralism allows a number

35. See, e.g., ACKERMAN & HEINZERLING, *supra* note 5, at 68-69; JOHN BROOME, *ETHICS OUT OF ECONOMICS* 180-81 (Cambridge Univ. Press, 2004); Driesen, *supra* note 29, at 340.

36. See SUNSTEIN, *supra* note 5, at 95 (“It is plausible to think that people’s WTP to reduce statistical risks is nonlinear. As the probability approaches 100 percent, people become willing to pay an amount for risk reduction that rises nonlinearly to 100 percent of their wealth; as the risk approaches zero, WTP nonlinearly approaches nothing.”); Richard A. Posner, *CBA: Definition, Justification, and Comment on Conference Papers*, in *COST-BENEFIT ANALYSIS: LEGAL, ECONOMIC, AND PHILOSOPHICAL PERSPECTIVES* 324-25 (Matthew D. Adler & Eric A. Posner, eds.) (2001) (noting “people are much more reluctant (plausibly more than a thousand times as reluctant) to be subjected to” a 1 in 1,000 mortality risk than to a 1 in 1 million mortality risk). Cf. W. Kip Viscusi, *The Devaluation of Life*, 3 *REG. & GOVERNANCE* 103, 106 (2009) (observing that “as the magnitude of the risk change increases, willingness-to-pay amounts increase less than proportionally and willingness-to-accept amounts increase more than proportionally”); Viscusi, *supra* note 10, at 140, fig. 12.1 (showing nonlinear function for relation of worker wages to workplace fatality risks).

37. See Posner, *supra* note 36, at 324-25.

38. *Id.* at 324.

39. See William Twining, *Normative and Legal Pluralism: A Global Perspective*, 20 *DUKE J. COMP. & INT’L L.* 473, 477 (2010).

of normative factors—including overall well-being, deontological constraints and options, equity, fairness, personal autonomy, rights, and possibly others—to ground moral determinations.⁴⁰ It further allows multiple factors to jointly determine the moral rightness or wrongness of a single action, as different factors interact with one another in determining whether the action is right or wrong, all things considered.⁴¹

Though a general defense of normative pluralism lies beyond the scope of this article, suffice it to say that normative pluralism is a popular position among both contemporary moral philosophers⁴² and philosophically-inclined legal theorists interested in CBA and risk regulation.⁴³ This article argues in Parts II and III that, when it comes to determining whether it is morally justifiable to conduct mortality risk CBA with reference to a VSL, different normative factors are relevant and decisive in

40. Normative pluralism is often thought to go hand-in-hand with intuitionism. *See, e.g.*, JOHN RAWLS, *A THEORY OF JUSTICE* 34 (1971) (“Intuitionist theories, then, have two features: first, they consist of a plurality of first principles which may conflict to give contrary directives in particular types of cases; and second, they include no explicit method, no priority rules, for weighing these principles against one another: we are simply to strike a balance by intuition, by what seems to us most nearly right.”).

41. *See* SHELLY KAGAN, *NORMATIVE ETHICS* 177 (1998).

42. *See, e.g., id.* at 72, 80-81 (noting possibility of coherent normative pluralism as to wrong-making properties when discussing deontological constraints and moderate (“threshold”) deontology); Johann Frick, *Contractualism and Social Risk*, 43 *PHIL. & PUB. AFF.* 175, 219, 221-22 (2015) (proposing a “pluralist account of moral rightness” that embraces considerations of both equity and well-being); *see also* 3 TERENCE IRWIN, *THE DEVELOPMENT OF ETHICS: A HISTORICAL AND CRITICAL STUDY* 690, 906 (2009) (noting that both W.D. Ross and the 18th century moral philosopher Richard Price accept a “pluralist view” that posits a “pluralism about ultimate moral principles,” “none of which is always prior to the others”); *Value Pluralism*, *STAN. ENCYCLOPEDIA PHIL.* (2018), <http://plato.stanford.edu/entries/value-pluralism/> [<https://perma.cc/MK52-AX5R>] (noting that, in addition to the prominent early 20th-century moral philosopher W.D. Ross, numerous contemporary moral philosophers are committed to foundational normative pluralism, including Judith Jarvis Thomson, Bernard Williams, Charles Taylor, and Michael Stocker).

43. *See, e.g.*, ADLER & POSNER, *supra* note 3, at 53, (“Our position is not utilitarianism, but ‘weak welfarism.’ ‘Weak welfarism’ claims that overall welfare is morally relevant, not that it is morally decisive. Morality may encompass a plurality of moral factors. It certainly includes overall welfare; but it may also include such factors as moral rights, the fair distribution of welfare, and even moral considerations wholly detached from welfare, such as intrinsic environmental values.” (footnote omitted); SUNSTEIN, *supra* note 5, at 7, 92-93, 113-15, 129-30 (discussing relevance of two norms—welfare and autonomy—to analysis of normative basis of mortality risk regulation); *see also* Matthew D. Adler & Eric A. Posner, *Implementing Cost-Benefit Analysis When Preferences are Distorted*, 29 *J. LEGAL STUD.* 1105, 1106-08 (2000).

different moral contexts. The most salient norms are overall well-being, personal autonomy, and equity.⁴⁴

Because normative pluralism holds that normative judgments should be informed by a number of independent norms, none of which can be fully expressed in terms of the others, it makes room for the possibility that a certain norm or group of norms should prevail in some situations, while a different norm or group of norms should prevail in others.⁴⁵ For example, a normative pluralist might hold that considerations of fairness or equity should prevail over considerations of overall well-being when an act or policy will deliver tiny benefits to a vast number of people, but inflict devastating losses on a very small number of people. Such a policy could be rejected on equitable grounds notwithstanding that it might increase aggregate well-being significantly on net. On the other hand, a pluralist might hold that considerations of overall well-being should prevail over equitable considerations when, for example, choosing between a policy that poses a 1 in 1,000 risk of death to 1 million people (1,000 expected deaths) and a policy that poses a 1 in 10 risk of death to 100 people (10 expected deaths). The latter policy could be justified on grounds of overall well-being notwithstanding that it imposes a far greater *ex ante* burden on each of the 100 persons it would affect than the alternative policy would impose on each of the 1 million persons *it* would affect.

When normative considerations conflict, a normative pluralist would reach a considered judgment about the morally correct course of action by gauging the *strength* of those competing considerations.⁴⁶ If, for example, equitable considerations weigh strongly in favor of a particular course of action, and considerations of overall well-being weigh only weakly against it, the normative pluralist could sensibly hold the course of action to be morally permissible.⁴⁷ The more difficult cases are those in which the conflicting normative considerations are of relatively equal strength.

44. See *infra* Parts II & III.

45. See Twining, *supra* note 39, at 480-81; see also Dietmar von der Pfordten, *Five Elements of Normative Ethics – A General Theory of Normative Individualism*, 15 ETHICAL THEORY & MORAL PRAC. 449, 449-50 (2012).

46. See KAGAN, *supra* note 41, at 183.

47. Waisman, *supra* note 19, at 1296-97.

As a normative pluralist, it is tempting to look for a decision rule or algorithm with which to resolve such cases. That temptation should be resisted, in my view. The task of moral theory is to identify the relevant normative factors that should inform our moral judgments and to explain why those factors matter,⁴⁸ not to provide a quasi-mathematical rule for making all moral judgments. Moral judgment in particular cases is inherently multifactorial, and, therefore, deeply fact-dependent, making a rigid algorithm an inappropriate substitute for considered reflection and deliberation when the normative factors one has identified stand in equipoise.⁴⁹

II. CONSUMPTION RISKS AND PERSONAL AUTONOMY

Commonly, an industrial activity poses a risk of death to a group of persons that both benefits meaningfully from the activity and shoulders all or substantially all of the costs of complying with regulations that reduce the risk.⁵⁰ Consumption risks are the best example of this type of risk imposition.⁵¹ Many industries produce goods that pose a small but nontrivial risk of death to the consumers who purchase and use them. Automobiles, heavy machinery, pharmaceuticals, publically-provided drinking water, and pesticide-treated food are some of the most obvious examples. All or substantially all of the costs of complying with government regulations designed to reduce such risks are

48. See KAGAN, *supra* note 41, at 17-22, 189-203 (distinguishing between “factorial” views that identify which normative factors or properties make acts wrong and “foundational” views that explain why the identified normative factors are wrong-making).

49. See Frick, *supra* note 42, at 223 (“The aim and ambition of moral philosophy should be to inform our judgment, by making us alive to the relevant ethical considerations, not to abolish the need for judgment altogether”).

50. See *supra* text accompanying notes 14-15.

51. However, workplace risks could conceivably exhibit the same structure if all or substantially all of the costs of regulatory compliance were met by reducing wages or laying off workers. See SUNSTEIN, *supra* note 5, at 117 (noting that, with respect to workers’ compensation regulation, “nonunionized workers faced a dollar-for-dollar wage reduction, corresponding almost perfectly to the expected value of the benefits that they received.”). By the same token, some consumption risks might *not* exhibit this structure. For example, if the costs of complying with regulations reducing the risk of using or consuming a product were deducted from shareholder profits rather than passed on to consumer, the risk would exhibit the structure of the typical workplace risk.

typically passed on to consumers in the form of price increases.⁵² When this occurs, the group of people exposed to the risk in question is roughly the same group responsible for paying the costs of reducing that risk.⁵³ Moreover, these risk-bearing, cost-bearing consumers typically benefit meaningfully from the industrial activity that gives rise to the risk, as they willingly pay to consume its product or service.⁵⁴ These features of the typical consumption risk create a particular moral context, one that is relevant to the justifiability of using a VSL to guide mortality risk regulation. In this Part, I argue that, when it comes to consumption risks exhibiting this structure, using a VSL to guide mortality risk CBA is morally defensible on the basis of the norm of personal autonomy.⁵⁵ It is not, however, defensible on the basis of welfare maximization, the norm on which mortality risk CBA is often defended.⁵⁶

A. Example: Arsenic in Drinking Water

In 1996, Congress directed EPA to newly regulate the presence of arsenic in publically provided drinking water.⁵⁷ Arsenic is a naturally occurring carcinogen.⁵⁸ When consumed in drinking water, arsenic has been found to cause cancer in a number of organs, including the lungs, bladder, kidneys, and liver.⁵⁹ A 1999 report from the National Research Council concluded that, at the then-prevailing maximum contaminant

52. *See id.* at 113, 117 (noting that “Easy Cases,” in which “the cost of eliminating any risk is borne entirely by those who benefit from eliminating that risk . . .”).

53. HOWARD BEALES ET AL., REGULATORY TRANSPARENCY PROJECT, GOVERNMENT REGULATION: THE GOOD, THE BAD, & THE UGLY 4 (2017), <http://regproject.org/wp-content/uploads/RTP-Regulatory-Process-Working-Group-Paper.pdf> [<https://perma.cc/535W-WNFE>].

54. Arno J. Rethans, *Towards Determinants of Acceptable Risk: The Case of Product Risks*, 8 ADVANCES CONSUMER RES. 506, 506-07 (1981), <http://www.acrweb.site.org/volumes/5846/volumes/v08/NA-08> [<https://perma.cc/242F-4478>].

55. *See* SUNSTEIN, *supra* note 5, at 113-15 (arguing that using a VSL to guide CBA in this type of case is justifiable on the basis of autonomy).

56. *But see id.* at 114-15 (arguing that, in addition to autonomy, welfare provides an independent normative basis for using a VSL to guide CBA in this context); ADLER & POSNER, *supra* note 3, at 178-82 (arguing that “the VSL method . . . allows CBA to determine with reasonable accuracy whether lifesaving projects really do increase overall welfare, given their additional welfare effects (for example, compliance costs)”).

57. *See* EPA Analysis, *supra* note 15, § 1.1, § 2.3, at 2-6 to 2-7.

58. *Id.* § 2.2.1.

59. *Id.* §§ 1.2, 2.2.1.

level of 50 µg/L (micrograms per liter) of drinking water, arsenic could pose an aggregate cancer risk “on the order of 1 in 100.”⁶⁰

The Safe Drinking Water Act directs the EPA to undertake CBA when setting regulatory levels.⁶¹ Specifically, the Act provided that,

After determining an MCL [maximum contaminant level] based on affordable technology for large systems, EPA must complete an economic analysis to determine whether the benefits of the standard justify the costs. If not, EPA may adjust the MCL to a level that ‘maximizes health risk reduction benefits at a cost that is justified by the benefits.’⁶²

In its Economic Analysis, EPA noted that section 1412(b)(6)(A) of the Act allowed it to set the MCL for arsenic at a less stringent level than would be economically feasible so long as that level “maximizes health risk reduction at a level where costs and benefits are also considered.”⁶³

After conducting a thorough cost-benefit analysis, EPA set the new MCL for arsenic at 10 µg/L.⁶⁴ The annual monetized health benefit of reducing the arsenic MCL from 50 µg/L to 10 µg/L was found to be in the range of \$139.6 million to \$197.7 million (median of \$168.65 million).⁶⁵ These benefits included the avoidance of both fatal and non-fatal lung and bladder cancer.⁶⁶ Reducing the arsenic MCL to 10 µg/L would result in the avoidance of 21.3 to 29.8 premature cancer deaths annually (a median of 25.55 deaths avoided annually).⁶⁷ In monetizing the benefit of avoiding premature deaths caused by lung or bladder cancer, the EPA used a VSL of \$6.1 million.⁶⁸ To monetize the

60. *Id.* § 1.2.

61. *Id.* § 2.3, at 2-6.

62. Safe Drinking Water Act of 1996, Pub. L. No. 104-182, § 1412(b)(6); *see also* EPA Analysis, *supra* note 15, § 3.1.3.

63. EPA Analysis, *supra* note 15, § 1.3.

64. *Id.* The EPA determined that the mean cancer risk posed by arsenic at 10 µg/L was between .63 in 10,000 and 2.99 in 10,000. *Id.* § 5.3.2, at 5-14 exhibit 5-4(c).

65. *Id.* § 5.4, at 5-26 exhibit 5-11.

66. EPA’s quantitative benefit analysis was confined to the avoidance of lung and bladder cancer because these were the only types of cancer as to which EPA had sufficient data to perform a quantitative analysis. *See id.* § 5.3.1. EPA determined that lung cancer carries a twenty-year mortality rate of 88%, whereas the twenty-year mortality rate for bladder cancer is just 26%. *See id.* § B.1.7, at B-12.

67. *Id.* § 5.4.1, at 5-22 exhibit 5-9(c).

68. *See* EPA Analysis, *supra* note 15, § 5.4.2, at 5-23.

benefit of avoiding non-fatal lung or bladder cancer, the EPA used a figure of \$607,162.⁶⁹

The total annual cost of reducing the MCL from 50 µg/L to 10 µg/L was found to be \$205.6 million, exceeding the median monetized health benefits by \$7.9 million to \$66 million.⁷⁰ Although the costs of reducing the level to 10 µg/L were, therefore, found to marginally exceed the monetized benefits, EPA's decision to nevertheless set the level at 10 µg/L was informed by its conclusion that doing so carried substantial non-quantifiable benefits.⁷¹

The EPA further noted that the costs of complying with the new arsenic regulation were likely to be passed on to consumers in the form of price increases.⁷² The EPA determined that reducing the arsenic MCL to 10 µg/L would increase the average household's annual water bill by \$31.85.⁷³

Thus, as Cass Sunstein has observed, the EPA's regulation of arsenic in drinking water exhibits the structure of the typical consumption risk.⁷⁴ The risky activity at issue—the public provision of drinking water—obviously provides a significant benefit to each risk-bearing water consumer. Moreover, as the EPA noted, the costs of complying with arsenic regulations would be passed on to water consumers in the form of rate increases.⁷⁵ Thus, the people exposed to cancer risks from arsenic both benefit significantly from the industrial activity that gives rise to those risks, *and* bear all or substantially all of the costs of compliance

69. *See id.* at 5-24

70. *Id.* § 5.4.3, at 5-26 exhibit 5-11. The \$205.6 million total annual cost figure is based on a discount rate of 7%. *Id.* § 7.3.1, at 7-2 exhibit 7-1. If a lower discount rate of 3% is used, the total annual cost is \$180.4 million. *Id.* § 1.4.

71. *See id.* § 5.5.3 (“Were EPA able to quantify some of the currently non-quantifiable health effects and other benefits associated with arsenic regulation, monetized benefits estimates could be significantly higher than what are shown”); *see also id.* § 1.4, at 1-7, § 5.4.3, at 5-26 exhibit 5-11.

72. *See id.* § 6.3.3, at 6-34 (“Household level costs are considered a good proxy for the affordability of rule compliance with regard to CWSs [community water systems], since water systems recover costs at the household level through increased water rates.”); *see also* SUNSTEIN, *supra* note 5, at 117 (“For drinking water regulation, . . . [t]he entire cost of regulation is passed onto consumers in the form of higher water bills.”).

73. EPA Analysis, *supra* note 15, § 6.6.3, at 6-35 exhibit 6-17.

74. *See* SUNSTEIN, *supra* note 5, at 113, 117 (observing that regulation of arsenic in drinking water is one in which “the cost of eliminating any risk is *borne entirely by those who benefit from eliminating that risk*”).

75. *Id.* at 113.

with regulations that reduce those risks. This makes the cancer risk posed by arsenic in drinking water a good real-world example of a consumption risk.

B. Welfare Maximization Fails to Justify Use of a VSL

Suppose the mortality risk from arsenic in drinking water is 1 in 100,000 and that this risk falls on 100,000 people. Suppose the risk can be eliminated by a regulation that will cost the industry \$20 million, and suppose this cost will be entirely passed on to and equally distributed among the 100,000 people subject to the risk, for a per capita cost of \$200 (call this example *Arsenic (100,000)*). If the government decides not to regulate under these circumstances, can its decision be defended on the basis of maximizing overall well-being? Although a number of commentators have argued or suggested that it can,⁷⁶ a close analysis suggests otherwise.

Suppose that the most that any of the 100,000 people subject to the arsenic risk would be willing to pay to avoid a 1 in 100,000 mortality risk is \$90. Most people will agree that this fact shows that each person would *prefer* bearing the 1 in 100,000 risk to bearing the \$200 cost, which in turn arguably shows that imposing the risk rather than the cost would leave each person *better off*. One might posit that these facts have a clear implication for overall well-being, namely, that imposing the \$20 million cost on the group of 100,000 risk-bearers would leave the group worse off than would imposing a 1 in 100,000 risk on the group. One might then infer, on this basis, that enacting the regulation—and thereby imposing the cost rather than the risk—would be normatively indefensible. Here is the argument laid out in steps:

Individual Willingness-to-Pay: The most that any of these 100,000 people would be *willing to pay* to eliminate the 1 in 100,000 risk of death facing them is \$90.

Individual Preference: Therefore, each person would *prefer* bearing a 1 in 100,000 risk to bearing a \$200 cost.

76. See, e.g., ADLER & POSNER, *supra* note 3, at 178-82; SUNSTEIN, *supra* note 5, at 92, 113-15, 127-30; cf. Adler & Posner, *supra* note 43, at 1110-11.

Individual Well-being: Therefore, imposing a 1 in 100,000 risk, rather than a \$200 cost on each, person would leave each person *better off*.⁷⁷

Overall Well-being: Therefore, imposing the risk rather than the cost would leave the entire group of 100,000 affected persons *better off*. Thus, overall well-being—the aggregate well-being of all affected persons—would be greater in a world in which the risk is imposed than it would be in a world in which the cost is imposed.

Normative Conclusion: Therefore, it is normatively indefensible to impose the cost.

There might be reason to question the move from step 2 to step 3. People's preferences might be irrational or ill-informed, such that the course of action they prefer is not the course of action that maximizes their welfare.⁷⁸ There also might be reason to question the move from step 4 to step 5. Another norm—such as equity or autonomy—might make it normatively defensible to impose the cost rather than the risk notwithstanding that overall well-being would be greater if the risk were imposed rather than the cost.⁷⁹

But the move I want to focus on is the move from step 3 to step 4. From an *ex ante* point of view, each person is better off bearing the risk than bearing the cost. For each person, bearing the risk is equivalent to bearing a \$90 cost, whereas eliminating the risk requires each person to bear a \$200 cost. If the risk is imposed, society bears a total cost of \$9 million, whereas if the cost is imposed and the risk eliminated, society bears a total cost of \$20 million. Thus, from an *ex ante* point of view, overall well-being is greater if the risk, rather than the cost, is imposed.

However, even if all of that is true, it does not seem to follow that, *all things considered*, overall well-being is greater if the risk is imposed instead of the cost. That is because the analysis thus

77. See SUNSTEIN, *supra* note 5, at 92 (“Suppose that people are willing to pay \$60, but no more, to eliminate a risk of 1/100,000. If so, then it might be assumed that their welfare is increased by asking them to pay that amount—and that their welfare is decreased by asking them to pay more.”).

78. See, e.g., *id.* at 118-22 (discussing phenomena of “miswanting” and informational and behavioral market failures); Adler & Posner, *supra* note 43, at 1116-22 (discussing phenomena of uninformed preferences, adaptive preferences, and objectively bad preferences).

79. See *supra* Part I.B (discussing normative pluralism).

far has focused only on ex ante effects on welfare, but has ignored ex post effects.

When a 1 in 100,000 risk of death is imposed on 100,000 people, two welfare-related losses would be expected to occur, one ex ante and one ex post.⁸⁰ Ex ante, each of the 100,000 people exposed to the risk experiences a welfare-related loss⁸¹ just by virtue of the imposition of the risk. By hypothesis, this loss is equivalent to the loss associated with paying a \$90 monetary cost, as each person would be willing to pay that much to avoid the risk at issue. But in addition, one unlucky person is expected to die if the risk is imposed. The welfare loss associated with this *ex post* effect of the risk imposition—one expected death—seems to get ignored in the move from step 3 to step 4 in the five-step argument laid out above. The argument sensibly deploys willingness-to-pay to capture the welfare effect of the risk's *imposition* on 100,000 people, but fails to take account of the welfare effect of the risk's expected *materialization* in the death of a single person.

I am certainly not the first to lodge this critique of the welfare maximization justification for the VSL. Matthew Adler has observed that the critique originated with a 1978 article by the philosopher John Broome.⁸² Adler describes Broome's critique as follows:

Imagine the government must choose between two outcomes: the status quo, and an outcome in which (1) exactly one person will die; (2) one million persons are at risk of suffering this death; and (3) other benefits, valued at \$B, are realized. . . . Each person P_i , out of the million persons at risk of dying, would be willing to accept some amount V_i as compensation for his one-in-one-million risk.

80. See SUNSTEIN, *supra* note 5, at 92.

81. *Id.*; Frank Ackerman & Lisa Heinzerling, *Pricing the Priceless: Cost-Benefit Analysis of Environmental Protection*, 150 U. PA. L. REV. 1553, 1565 (2002) (“In acknowledging the monetary value of reducing risk, economic analysts have contributed to our growing awareness that life-threatening risk itself—and not just the end result of such risk, death—is an injury.”); see also Posner, *supra* note 36, at 324-25 (noting the “ex ante” costs” of projects that impose risks of death and suggesting that ex ante costs can differ even when ex post costs are identical). It may be more appropriate, technically, to think of the imposition of a small risk of death as involving a *setback to a welfare-related interest* of the risk-bearer rather than an actual *welfare loss* to the risk-bearer. (I thank Johann Frick for impressing this distinction on me during a conversation in early 2014.) For the sake of simplicity, I will continue to refer to this as a “welfare loss.”

82. See generally Broome, *supra* note 8.

These (small) monetary amounts can be aggregated to produce an overall cost figure $\$C = \sum V_i$ *On an ex post view there seem to be two kinds of harm here, not one.* First, each of the million suffers a risk of death; and that harm, cumulatively, is captured by $\sum V_i$. Second, and quite separately, the person who will die is very seriously harmed: his life is shortened, and (in the typical case) the balance of welfare goods realized over the course of his life will be much lower. *But this second harm is, apparently, ignored when the total cost of the policy is set equal to $\sum V_i$.* If $\sum V_i$ is, in some way, a measure of this second harm, then shouldn't the total cost of the policy be set equal to $2 * \sum V_i$?⁸³

Now, it is of course true that, due to the stochastic nature of risk, it is *expected*, but not certain, that one person will die. There is a significant chance that the imposition of a 1 in 100,000 mortality risk on 100,000 people will not result in a death.⁸⁴ However, it is more likely than not that at least one person will die, and there is a significant possibility that more than one person will die.⁸⁵ In light of these probabilities, any determination of the likely effect of the risk imposition on overall well-being needs to take account of, among other things, the effect that an actual death would have on overall well-being.

83. Matthew D. Adler, *Risk, Death, and Time: A Comment on Judge Williams' Defense of Cost-Benefit Analysis*, 53 ADMIN. L. REV. 271, 282-83 (2001) (emphasis added); see also Lisa Heinzerling, *The Rights of Statistical People*, 24 HARV. ENVTL. L. REV. 189, 203-206 (2000) ("One way in which analysts treat the valuation of risk as equivalent to a valuation of life is that they do not calculate the value of both statistical life and life itself. They calculate only the value of statistical life The value of a discrete risk, however, remains the same regardless of whether anyone actually ends up dying as a result of that risk. Risk and death are two separate injuries." (footnote omitted)); see also Ackerman & Heinzerling, *supra* note 81, at 1564-65 ("In practice, however, analysts often ignore the distinction between valuing risk and valuing life. Many regulations reduce risk for a large number of people and avoid actual death for a much smaller number. A complete cost-benefit analysis should, therefore, include valuation of both of these benefits. However, the standard practice is to calculate a value only for 'statistical' life and to ignore life itself." (footnote omitted)).

84. As to each risk-bearer, there is a 99.999% (.99999/1) chance that the imposition will *not* result in death. Thus, the chance that the risk imposition will not result in *any* risk-bearer's death is therefore $.99999^{100,000}$ or roughly 36.8%.

85. Imposing a 1 in 100,000 risk of death on 100,000 people creates roughly a 63.2% chance of at least one death ($1 - .99999^{100,000}$). The chance that *exactly* one death will occur is roughly 36.8% ($.99999^{99,999} \times (.00001) \times (100,000)$), while the chance that two or more deaths will occur is roughly 26.4%.

When the welfare setback associated with the one expected death is added to the welfare setback associated with the imposition of a 1 in 100,000 risk of death on 100,000 people, it is not at all clear that, all things considered, imposing the risk would leave *everyone better off* than imposing the \$20 million cost would. But this is precisely what is claimed in step 4 of the above argument. In the unlikely (but possible) event that the risk imposition did not result in anyone's death, then it is plausible that imposing the risk rather than the cost turned out to be the welfare-maximizing course of action.⁸⁶ This follows straightforwardly from step 3 of the argument: if each person is better off being exposed to the risk than paying the cost, then imposing the risk maximizes overall well-being *assuming the only welfare-related loss that occurs is that associated with the mere imposition of the risk.*⁸⁷ However, in the considerably more likely event that the risk imposition materializes in at least one death, the truth of the claim made in step 3 of the argument does not entail the truth of the claim made in step 4.

Against this conclusion, Matthew Adler and Eric Posner have argued that the welfare loss associated with death itself *does* get captured in the move from step 3 to step 4. Claiming that "the overall-welfare loss in premature death is not infinite, but the *finite* difference between a longer and shorter life history," they argue that the VSL represents a plausible monetization of this finite welfare loss based on a person's willingness-to-pay to avoid very small mortality risks.⁸⁸ Adler and Posner concede that using willingness-to-pay or willingness-to-accept as a measure of well-being breaks down when it comes to avoiding *certain* death, since a rational person would be willing to pay all of their wealth to avoid premature death and would be willing to accept only an infinite amount of money to succumb to premature death,

86. ADLER & POSNER, *supra* note 3, at 179.

87. *See id.* at 180.

88. *See id.* at 179-180; *see also* John Broome, *Cost-Benefit Analysis and Population*, in *COST-BENEFIT ANALYSIS: LEGAL, ECONOMIC, AND PHILOSOPHICAL PERSPECTIVES* 117, 122 (Matthew D. Adler & Eric A. Posner, eds., 2001) ("When a person's life is saved, she lives a longer life than she would have lived. The benefit to her is the difference between the goodness, or value, of her longer life and the goodness, or value, of the shorter life she would have lived. Conversely, if an event kills a person, the harm done her is the difference between the value of the longer life she would have lived and the value of the shorter life she actually lives.").

notwithstanding that the welfare loss associated with premature death is finite.⁸⁹ In their view, “[t]he VSL method avoids this problem, by translating” the welfare loss associated with premature death “into dollars at an ordinary, ‘premortem’ rate rather than the postmortem rate.”⁹⁰ This “premortem” conversion rate between welfare and dollars is the rate that would determine a person’s willingness-to-pay for “marginal nonlongevity changes to his welfare . . . (such as small changes in consumption, in pleasurable or painful experiences, or the provision of some public good)”⁹¹

Although the premortem conversion rate is supposed to apply to “marginal nonlongevity” welfare changes—that is, welfare changes *not* related to premature death—Adler and Posner nevertheless attempt to derive this rate from a person’s willingness-to-pay to avoid a small risk of premature death.⁹² Adler and Posner reason that if a person is willing to pay \$90 to avoid a 1 in 100,000 risk of premature death, this implies that:

$$k \times \Delta U \times 1/r = \$90$$

where k is the premortem conversion rate between welfare and dollars, ΔU is the welfare loss associated with premature death, and $1/r$ is the probability of premature death occurring (here, $1/100,000$).⁹³ Multiplying both sides of this equation by r , Adler and Posner obtain the following result:

$$k \times \Delta U = \$9 \text{ million.}^{94}$$

According to Adler and Posner, $k \times \Delta U$ represents “the utility difference between the shorter life . . . and the longer life . . . converted into dollars at rate k .”⁹⁵ If this is correct, the \$9 million VSL *does* represent an accurate monetization of the welfare loss associated with premature death. In that case, a regulation that would cost \$20 million but prevent only one premature death

89. See ADLER & POSNER, *supra* note 3, at 180 (“Absent a bequest function, the dying individual asked about WTP/WTA at the moment before death would pay his entire wealth (above a subsistence level) to avoid death, and would have an infinite WTA . . .”).

90. *Id.*

91. *Id.*

92. *See id.*

93. *See id.*

94. See ADLER & POSNER, *supra* note 3, at 180.

95. *See id.*

would decrease overall well-being and could therefore be rejected on welfare maximization grounds.⁹⁶

There are two problems with this argument, however. First, because willingness-to-pay to avoid mortality risks increases in a nonlinear fashion as the magnitude of the mortality risk increases,⁹⁷ Adler and Posner's approach leads to the arbitrary result that the value of $k \times \Delta U$ varies with the magnitude of the mortality risk ($1/r$) upon which the willingness-to-pay calculation is based. For example, a person who would be willing to pay \$90 to avoid a mortality risk of 1 in 100,000 might well be willing to pay significantly more than 100 times as much money (\$9,000) to avoid a mortality risk 100 times as great (1 in 1,000). If someone were willing to pay \$20,000 to avoid a mortality risk of 1 in 1,000, this would yield a value of \$20 million for $k \times \Delta U$.⁹⁸ Why should the \$9 million figure generated by Adler and Posner's analysis when r is set at 100,000 be preferred to the \$20 million figure that is generated when r is set at 1,000? Both risks are low enough to deliver only the sort of "small" or "marginal" changes to welfare Adler and Posner are concerned with, yet they would likely yield significantly different results for the VSL.⁹⁹

96. *See id.* ("So the VSL method is not just an arbitrary device to avoid infinite dollar valuations. Because it translates the utility loss that occurs with an individual's premature death into dollars by using a utility-to-dollar conversion factor proximate to the conversion factor implicit in WTP/WTA for welfare impacts other than premature death, *this technique allows CBA to determine with reasonable accuracy whether lifesaving projects really do increase overall welfare*, given their additional welfare effects (for example, compliance costs).") (emphasis added).

97. *See* sources cited *supra* note 35-38.

98. The converse nonlinear phenomenon occurs as the magnitude of the risk decreases to a level close to zero. For example, although a person might be willing to pay \$90 to avoid a 1 in 100,000 mortality risk, she might well be willing to pay *less* than a tenth as much (\$9) to avoid a risk a tenth as great (1 in 1 million).

99. In earlier work, Adler himself seemed to recognize the force of this objection. *See* Adler, *supra* note 83, at 284 ("Where one person out of a very large population is certain to die, each person's V_i might equal zero, and then $\sum V_i$ will equal zero—which can hardly be an accurate measure of the welfare loss that will befall the one person who ends up dying. In general, if we imagine one person out of a population dying, with the size of the population varying, the $\sum V_i$ will presumably vary as well—and yet the (ex post) harm from the death remains the same. Why take this invariant harm to be accurately measured by the variable $\sum V_i$?"); *see also* Broome, *supra* note 8, at 93 ("There may be a very small degree of probability which people do not distinguish from no chance at all, so they will accept the chance of being killed without requiring compensation. Suppose the probability is one in ten million. Then a country with a population of fifty million could kill five people at random, assessing their deaths at no cost at all.").

There would seem to be two ways to explain the finding that the value of $k \times \Delta U$ increases as the value of $1/r$ increases. One could either posit that the value of k increases when a higher magnitude mortality risk is involved or posit that the value of ΔU increases when a higher magnitude mortality risk is involved. The latter possibility is plainly absurd. The welfare loss associated with premature death cannot plausibly be thought to vary depending on the magnitude of the risk of premature death occurring.¹⁰⁰ The former possibility, however, contradicts Adler and Posner's assertion that k "is precisely the rate that would determine [a person's] WTP/WTA for marginal nonlongevity changes to his welfare"¹⁰¹ Adler and Posner say nothing to suggest the possibility that k can vary. They certainly make it sound as though k is a constant, fixed conversion rate for all "marginal nonlongevity changes to welfare." However, Adler and Posner's approach suggests the value of k would change *significantly* depending on the magnitude of the mortality risk at issue. As the magnitude of the risk increases, k increases significantly.¹⁰² The question then becomes: why should the k value derived from a person's willingness to pay to avoid a 1 in 100 or 1 in 1,000 mortality risk be any less valid than the k value derived from a person's willingness to pay to avoid a 1 in 100,000 or 1 in 1 million mortality risk? At the end of the day, there does not seem to be any non-arbitrary way to set the value of k , nor do Adler and Posner suggest one. But if that is true, then the monetization of the welfare loss associated with premature death itself ($k \times \Delta U$) is arbitrary as well.

The other problem with Adler and Posner's analysis is that, even if it is correct as far as it goes, it ignores the *ex ante* welfare setback associated with imposition of a risk of premature death. According to Adler and Posner, the VSL represents a reasonably accurate monetization of the welfare loss associated with premature death. However, as Adler himself noted when summarizing Broome's critique in the passage quoted above,

100. See Adler, *supra* note 83, at 284.

101. See ADLER & POSNER, *supra* note 3, at 180.

102. For example, assuming a person were willing to pay \$90 to avoid a 1 in 100,000 mortality risk, but \$20,000 to avoid a 1 in 1,000 mortality risk, the value of k derived from an analysis based on a 1 in 1,000 mortality risk would be more than double the value derived from an analysis based on a 1 in 100,000 mortality risk.

when a premature death is expected to result from the imposition of a small mortality risk on a large number of people, *two* welfare-related losses are involved. In addition to the welfare loss associated with the premature death itself, there is also the welfare loss associated with the mere *imposition* of the mortality risk on the risk-bearers.¹⁰³ If Adler and Posner are correct that the VSL represents a monetization of the *ex post* welfare loss associated with an actual premature death, then how does the *ex ante* welfare loss associated with the risk imposition get accounted for and monetized so that it can be incorporated into CBA?

The standard response to this objection is that the mere imposition of a risk of death need does not involve an actual welfare loss (unless the risk-bearer is aware of the risk and experiences fear or apprehension as a result) and therefore need not be incorporated into CBA.¹⁰⁴ There are, however, two problems with this response. To begin with, it does not square well with statements by Viscusi, Sunstein, and the EPA to the effect that the VSL represents the tradeoff between money and small risks of death, not the tradeoff between money and death itself.¹⁰⁵ Such statements suggest—more or less explicitly—that

103. Adler and Posner seem to implicitly acknowledge this point when they posit that “small changes to [a person’s] utility . . . are converted to dollars at rate k ” and then use k in the equation $k \times \Delta U \times 1/r$ to represent a person’s “WTP to avoid the 1-in- r risk.” *Id.*; see also sources cited *supra* notes 77, 81, & 83.

104. See, e.g., Adler, *supra* note 83, at 286 (“[R]isk, per se, is not a harm. Fear and other affective states associated with risk can surely constitute welfare setbacks—and should be separately measured by the cost-benefit analyst—but the mere fact that a person is at risk of death does not diminish his well-being. Most plausibly, this is true where the person is unaware of the risk; it also may be true where he is aware of the risk, but does not fear it.”); Stephen R. Perry, *Risk, Harm, and Responsibility*, in *PHILOSOPHICAL FOUNDATIONS OF TORT LAW* 321, 336 (David G. Owen ed., 1995) (“[I]f the processes that caused or might in the future cause physical harm are deterministic, then there is no basis for saying that a person who has been put at risk by another of suffering such harm has, just by reason of being put at risk, sustained damage distinct in kind from the physical harm.”).

105. See, e.g., Viscusi, *supra* note 10, at 101 (“The value of a statistical life (VSL) is the individual’s money-risk tradeoff for small risks of death.”); SUNSTEIN, *supra* note 5, at 51 (“With these [VSL] values, the government is not actually ‘valuing life.’ It is valuing the reduction of mortality risks—typically by eliminating low-level risks, for example, risks of 1 in 100,000.”); *Mortality Risk Valuation: What Does It Mean To Place A Value On Life?*, EPA, <https://www.epa.gov/environmental-economics/mortality-risk-valuation#whatisvsl> [<https://perma.cc/97P3-5HNC>] (“The EPA does not place a dollar value on individual lives. Rather, when conducting a benefit-cost analysis of new environmental policies, the Agency uses estimates of how much people are willing to pay for small reductions in their risks of dying from adverse health conditions that may be caused by

a mortality risk imposition does, in fact, involve an actual welfare loss. Indeed, a 2010 EPA guidance document directly implies as much:

Some EPA policies are designed to reduce the risk of contracting a potentially fatal health effect such as cancer. *Reducing these risks of premature death provides welfare increases to those individuals affected by the policy.* These policies generally provide marginal changes in relatively small risks. . . . For BCA, analysts generally aggregate these small risks over the affected population to derive the number of statistical lives saved (or the number of statistical deaths avoided) and then use a ‘value of statistical life’ (VSL) to express these benefits in monetary terms.¹⁰⁶

The most natural way of reading such statements by the EPA, Viscusi, Sunstein, and others is that the mere imposition of a mortality risk involves a welfare loss to the risk-bearer and it is this welfare loss that drives individuals’ reported willingness to pay to be free of small mortality risks.

But even assuming for the sake of argument that, absent awareness or fear, the mere imposition of a mortality risk does not involve an actual welfare loss to the risk-bearer, the question is whether this means that such a risk imposition lacks moral significance. If I point a partially loaded gun at your head while you are asleep or unconscious and play Russian roulette, but no bullet is fired, have I violated an interest of yours? Have I wronged you? If you were later to view a film of the incident, would you be justified in feeling moral outrage? The intuitive answer to these questions seems to me clearly to be “Yes,” even if the source of the wrong were not an actual welfare setback.¹⁰⁷ If this intuition can be trusted, then it follows that, when a 1 in 100,000 risk of death imposed on 100,000 people results in one

environmental pollution. In the scientific literature, these estimates of willingness to pay for small reductions in mortality risks are often referred to as the ‘value of a statistical life.’”)

106. NAT’L CTR. FOR ENVTL. ECON., EPA, *Appendix B: Mortality Risk Valuation Estimates*, in GUIDELINES FOR PREPARING ECONOMIC ANALYSES B-1, B-1 (2014) (emphasis added).

107. See, e.g., Aaron James, *The Distinctive Significance of Systemic Risk*, 30 RATIO JURIS 239, 244 (2017) (“The interest we have in our exposure to risk is an objective interest in its own right, which is not to be confused with or reduced to our separate interests in not suffering the *experience* of risk exposure.”). I borrow the Russian roulette example from James’s paper.

person's death, 99,999 people suffered a setback to a welfare-related interest (even if not an actual welfare loss) and one person suffered the welfare loss associated with death itself. If Adler and Posner are correct that the \$9 million VSL is intended to capture the latter only, the former remains unaccounted for. Leaving risk imposition out of the calculation in this way is hard to understand. It is common sense that eliminating a mortality risk benefits every person who would have been exposed to it, not just those unlucky individuals who would have died were the risk to have materialized.

In sum, the VSL does not persuasively monetize the welfare loss associated with premature death itself. What it does persuasively monetize is the welfare loss (or the setback to welfare-related interests) associated with the imposition of a small *risk* of premature death on a large number of people. And even if the VSL does accurately capture the welfare loss associated with the *ex post* harm of death, it then fails to capture the *ex ante* interest setback resulting from the imposition of a small risk of death on a large number of people. Either way, the norm of maximizing overall well-being fails to persuasively justify the use of the VSL in standard cost-benefit analysis of mortality risk regulations, whether the type of mortality risk at issue is a consumption risk (as in *Arsenic (100,00)*), a workplace risk, or an environmental risk. The question then becomes whether other norms might be able to do the job.

C. The Norm of Personal Autonomy

The notion of autonomy has played a central role in Western moral philosophy, at least since Kant.¹⁰⁸ Where the autonomy of individual persons—rather than nations or other collective entities—is concerned, it is commonplace to refer to “personal” autonomy.¹⁰⁹ An individual possesses personal autonomy to the extent the individual is self-governing; that is, to the extent the

108. See IMMANUEL KANT, METAPHYSICAL PRINCIPLES OF VIRTUE, in ETHICAL PHILOSOPHY 44-54 (James W. Ellington trans., 1983) (1797).

109. See, e.g., *Personal Autonomy*, STAN. ENCYCLOPEDIA PHIL. (2018), <http://plato.stanford.edu/entries/personal-autonomy> [<https://perma.cc/Y7MX-8RUV>] (distinguishing group autonomy from autonomy of an individual person).

individual acts and lives in accordance with her own will.¹¹⁰ In the words of the moral philosopher John Christman,

[T]o be autonomous is to be one's own person, to be directed by considerations, desires, conditions, and characteristics that are not simply imposed externally upon one, but are part of what can somehow be considered one's authentic self. Autonomy in this sense seems an irrefutable value, especially since its opposite—being guided by forces external to the self and which one cannot authentically embrace—seems to mark the height of oppression.¹¹¹

Another contemporary moral philosopher, Shelly Kagan, offers the following definition:

To have autonomy is to have the various aspects of one's life under one's control. Typically, if I have autonomy over some aspect of my life (whether my career, my hair color, or how I spend this Thursday afternoon), then I can deliberate concerning how I want that aspect of my life to go, choose among the various alternatives open to me, and act so as to make my life the way I want it to be in that regard.¹¹²

From the standpoint of morality, respect for a person's autonomy clearly seems to be one of the normative factors that can determine the rightness or wrongness of a particular action. Kagan puts this point in terms of a moral prohibition on interfering with another's autonomy:

110. *See id.* (“Autonomous agents are self-governing agents. . . . In short, every agent has an authority over herself that is grounded, not in her political or social role, nor in any law or custom, but in the simple fact that she alone can initiate her actions.”).

111. *Autonomy in Moral and Political Philosophy*, STAN. ENCYCLOPEDIA PHIL. (2015), <http://plato.stanford.edu/entries/autonomy-moral> [<https://perma.cc/SW-E9-Z6W2>]; *see also* JOHN CHRISTMAN & JOEL ANDERSON, *AUTONOMY AND THE CHALLENGES TO LIBERALISM* 3 (2005). This definition of autonomy differs from that offered by Gerald Dworkin in his influential study. Dworkin understands autonomy as “a second-order capacity of persons to reflect critically upon their first-order preferences, desires, wishes, and so forth and the capacity to accept or attempt to change these in light of higher-order preferences.” GERALD DWORKIN, *THE THEORY AND PRACTICE OF AUTONOMY* 20 (1988). The more traditional understanding of autonomy adopted by Christman and Kagan—and the one I adopt in this article—seems to correspond roughly to what Dworkin calls “liberty.” *See id.* at 14 (“Suppose we think of liberty as being, roughly, the ability of a person to do what she wants, to have (significant) options that are not closed or made less eligible by the actions of other agents.”).

112. KAGAN, *supra* note 41, at 111.

One natural suggestion is that there should be a prohibition against interfering with the autonomy of another. If I act in such a way as to reduce, or to undermine, the control that someone else has over their own life, then I harm them in a profound way. My act fails to reflect the fact that my victim has his own conception of how he wants his life to go and that—but for my action—he would have been better able to live as he has chosen.¹¹³

Alternatively, one can put the idea in terms of rights: an individual has the right to determine for herself how her own life will go and how she will or will not act in the world.¹¹⁴ Although not absolute, this right of autonomy would seem to be entitled to substantial respect, particularly where an action or decision implicates the interests of the agent *only*, and not the interests of any other person.

D. Personal Autonomy Justifies Use of a VSL

1. A Stylized Consumption Risk Scenario: Arsenic (1)

Suppose the government is deciding whether to enact an arsenic regulation that would affect a single person, A (call this *Arsenic (1)*). Suppose the regulation would eliminate a 1 in 100,000 risk of death to A at a cost of \$200, which cost would be entirely passed on to A in the form of an increase in A's water rates. Suppose further that, after closely studying consumption- and employment-related choices involving mortality risks in the general population, the government concludes that very few people would be willing to pay more than \$90 to avoid a 1 in 100,000 risk of death. From the standpoint of autonomy, what is the normative significance of that conclusion for the government's decision whether to enact the regulation at issue in *Arsenic (1)*?

One possible claim is that, if very few people would be willing to pay \$200 to avoid a 1 in 100,000 risk of death, then enacting the arsenic regulation—which effectively forces A to make that payment—would show a lack of respect for A's

113. *Id.* at 291.

114. See *Autonomy in Moral and Political Philosophy*, *supra* note 111.

personal autonomy.¹¹⁵ This claim rests on the implicit inference that, because few people in the general population would be willing to pay \$200 to avoid a 1 in 100,000 mortality risk, A would not be willing to make that payment. This is, of course, a *non sequitur*.¹¹⁶ A might be one of the few people who *would* be willing to pay \$200 to avoid a 1 in 100,000 mortality risk. However, the government's data arguably shows that, as to A or any other individual in the general population, it is very *unlikely* that the individual would be willing to make such a payment. If, for some reason, information about A's degree of mortality risk aversion and risk-related preferences were unavailable, the government might reasonably rely on such data to make a determination about A.

Suppose one accepts the claim that, based on the government's data, it is highly unlikely that A would be willing to pay \$200 to avoid a 1 in 100,000 mortality risk. It would then be reasonable to ask: assuming A would be unwilling to make such a payment, what are the implications for whether it is morally permissible for the government to enact the regulation in *Arsenic (1)*?

To enact the regulation in such circumstances would seem to be an unjustifiable violation of A's personal autonomy. Because A is the only person who would be affected by the regulation—we are assuming that the associated cost and risk fall entirely on her—the government could not base a decision to enact the regulation on a need to safeguard the interests of any *other* person. Because all of the costs and benefits associated with enacting or not enacting the contemplated regulation fall on A and no one else, the norm of personal autonomy is very strongly implicated,

115. See SUNSTEIN, *supra* note 5, at 93 (“Perhaps regulatory policy should not be based on welfare. . . . Even if so, WTP might be defended on the ground of personal autonomy. On this view, people should be sovereign over their own lives. Government should respect people's choices about how to use limited resources (again so long as those choices are informed). When people decline to devote more than \$90 to the elimination of a 1/100,000 risk, it is because they would prefer to spend the money in a way that seems to them more desirable. If regulators do not use people's actual judgments, then they are insulting their dignity.”); see also Viscusi, *supra* note 8, at 322 (“Although many non-economists continue to attack the entire concept of monetizing risks to life, these implicit tradeoffs are reflective of how people themselves value the risks and *respect consumer sovereignty* in much the same way as do prices in other economic markets.” (emphasis added)).

116. See SUNSTEIN, *supra* note 5, at 92-93.

whereas norms like fairness and equity—which have to do with the distribution of burdens and benefits among *multiple* persons or parties—do not seem to be relevant. If, after informed deliberation, A freely and voluntarily prefers a 1 in 100,000 mortality risk to a \$200 cost, A's right to be sovereign in matters which affect her, and only her, interests would seem to have a clear moral implication: to enact the regulation in *Arsenic (I)* would be an unjustifiable violation of A's personal autonomy.

It is important to note that, as presented above, the autonomy-based case against regulating in *Arsenic (I)* is *not* founded on a claim that A voluntarily consented to the mortality risk associated with drinking water containing arsenic.¹¹⁷ The argument is based on an assumption—made on the basis of government data on the general population—about A's general preferences and behavior in the area of mortality risks. The argument is *not* based on any specific decision A made after being informed of the arsenic risk and asked to choose between taking on the 1 in 100,000 mortality risk or paying the \$200 cost. Instead, the autonomy-based case against the regulation rests implicitly on a claim that, had this choice been given to A, A *would have* chosen to assume the risk rather than pay the cost. It is, in other words, based on A's hypothetical consent, not her actual consent.¹¹⁸

In *Arsenic (I)*, the autonomy-based case against regulating seems relatively strong. It is not airtight, however. As noted above, the case depends on the assumption that A is *not* one of the very few people in the general population who would be willing to pay \$200 to avoid a 1 in 100,000 mortality risk. But, what if it turned out that A *were* a member of that tiny minority? What if, given all the relevant information, A would freely and deliberately choose to pay \$200 rather than be exposed to a 1 in 100,000 risk of premature death? If that were true, it would seem difficult to defend the choice not to regulate in *Arsenic (I)* on the basis of autonomy. Choosing not to regulate—and thereby imposing the risk rather than the cost on A—would run contrary to A's free and informed preferences.¹¹⁹ Because the regulation

117. *See id.*

118. For further discussion of the moral significance of hypothetical (as opposed to actual) consent to a risk imposition, see *infra* Part II.D.2 & text accompanying notes 142-44.

119. *See Autonomy in Moral and Political Philosophy, supra* note 111.

affects A and A alone, this would, arguably, violate A's right to sovereignty over her own life.¹²⁰

In such a case, the government might offer an alternative justification for choosing not to regulate: welfare maximization. The government could claim that A's preference for paying \$200 rather than accepting a 1 in 100,000 mortality risk, despite being informed, voluntary, and free, is *irrational* because, as a matter of fact, paying the cost would leave A worse off than being exposed to the risk. The government might defend that conclusion on the basis of its general population data by claiming that the fact that the *vast majority* of people prefer a 1 in 100,000 mortality risk to a \$200 cost shows that *all* people would be better off bearing the risk than bearing the cost. On this view, the few people who prefer the cost to the risk are guilty of what Cass Sunstein has termed "miswanting": freely and voluntarily preferring something contrary to their own best interests.¹²¹ If this is so, the government might defend a decision not to regulate in *Arsenic (I)* as an instance of justifiable paternalism,¹²² one similar to the government's requirement that occupants of moving automobiles wear seatbelts. Some people might make a free and informed decision not to wear a seatbelt when driving. But, by requiring all drivers and passengers to wear a seatbelt, the government effectively deems such a decision to be irrational. It deems the preference for not wearing a seatbelt to be an instance of miswanting, one that can therefore be justifiably overridden in order to maximize the well-being of the miswanter.

120. *Id.*

121. See SUNSTEIN, *supra* note 5, at 118-22 (discussing phenomena of "miswanting" and informational and behavioral market failures); see also Adler and Posner, *supra* note 43, at 1116-22 (discussing phenomena of uninformed preferences, adaptive preferences, and objectively bad preferences).

122. See KAGAN, *supra* note 41, at 112 ("Prohibitions against *paternalism* rule out interfering with another's autonomy . . . [by] forcing someone to do what's best for himself."); DEREK PARFIT, REASONS AND PERSONS 321 (1984) ("We are paternalists when we force someone act in his own interests. It provides some justification for paternalism, when this involves coercion or the infringement of someone's autonomy, if we are stopping this person from acting irrationally.").

2. *A More Realistic Scenario: Arsenic (100,000)*

The next question is whether the autonomy-based case against regulation loses any of its force when we move from *Arsenic (1)* to the more realistic case of *Arsenic (100,000)*, in which an arsenic mortality risk of identical magnitude—1 in 100,000—is imposed on 100,000 consumers. As above, suppose this risk could be eliminated at a total cost of \$20 million, which would be entirely passed on to the risk-bearing consumers in the form of a \$200-per-capita cost. As with *Arsenic (1)*, it should be assumed that the government's regulatory decision is based on general population data showing that very few people would be willing to pay more than \$90 to avoid a 1 in 100,000 mortality risk.

From an autonomy standpoint, enacting the regulation in *Arsenic (100,000)* features 100,000 instances of the autonomy violation involved in enacting the regulation in *Arsenic (1)*. Arguably, if enacting the regulation in *Arsenic (1)* is wrong because it would unjustifiably violate the autonomy of one person, then enacting the regulation in *Arsenic (100,000)* is *egregiously* wrong because it would unjustifiably violate the autonomy of 100,000 people.

Unlike in *Arsenic (1)*, however, the norm of maximizing overall well-being arguably favors a choice to regulate in *Arsenic (100,000)*. Where the risk in question would fall on 100,000 people, rather than one person, not enacting the regulation would be expected to result not just in the imposition of a very small risk of death on one person (as in *Arsenic (1)*), but rather in a single premature death, one which could be prevented by distributing a \$20 million cost among 100,000 people. For reasons discussed in the previous section, it is not at all clear that overall well-being would be greater in a world featuring *both* the imposition of a 1 in 100,000 risk on 100,000 people *and* the risk's materialization in one expected premature death than it would be in a world featuring the imposition of a \$200 cost on each of 100,000 people. There may, in fact, be good reason to believe that the welfare-maximizing course of action would be to impose the \$20 million

cost, thereby not only sparing 100,000 people from the risk imposition but also sparing one person from premature death.¹²³

A decision to regulate in *Arsenic (100,000)* might also be defended on the basis of fairness. In *Arsenic (100,000)*, it is clear from an ex ante point of view that *someone* would probably be killed were the regulation not enacted.¹²⁴ It is not as though this is a single-person case (like *Arsenic (1)*), in which it is highly unlikely that the risk will materialize in anyone's death. In *Arsenic (100,000)*, it is probable that at least one person will die. Therefore, it might be argued, it is not just a matter of respecting the autonomy interests of a group of 100,000 individuals, all of whom are identically situated with respect to the risk imposition in question. It is a matter of fairly balancing the interests of the 100,000 people affected ex ante against those of the one person likely to be fatally affected ex post. Thus, this situation is, arguably, not relevantly different from one in which 100,000 town residents are each asked to pay \$200 to subsidize the \$20 million rescue of a miner trapped at the bottom of the town mine (call this *Trapped Miner*). In both cases, it would arguably be unfair to allow the death of one person to occur simply to spare each of 100,000 people from having to pay a cost of \$200. In both cases, each of the 100,000 cost-bearers can, on the basis of the norm of fairness, justifiably be required to invest \$200 to save a life.¹²⁵ Thus, even if the welfare maximization case for

123. Of course, to say that the imposition of a 1 in 100,000 risk of death on 100,000 people would be *expected* to result in the death of one person is not to say that there is a 100% chance of it doing so. As to each risk-bearer, there is a 99.999% (.99999/1) chance that the imposition will *not* result in death. The chance that the risk imposition will not result in *any* risk-bearer's death is therefore .99999^{100,000} or roughly 36.8%. Thus, imposing a 1 in 100,000 risk of death on 100,000 people creates roughly a 63.2% chance of at least one death. The chance that *exactly* one death will occur is roughly 36.8% (.99999^{99,999} x (.00001) x (100,000)), while the chance that two or more deaths will occur is roughly 26.4%.

124. See Sophia Reibetanz, *Contractualism and Aggregation*, 108 ETHICS 296, 303-304 (1998).

125. The analogy between *Trapped Miner* and *Arsenic (100,000)* assumes that it makes no difference whether, at the time the decision must be made, the cost-bearers know the identity of the person who will die unless action is taken. *Trapped Miner* involves saving an identified life, whereas *Arsenic (100,000)* involves saving a statistical life. The analogy between the two situations therefore rests on the implicit claim that this difference ought not to be significant for purposes of determining the right course of action in either case. The critique of the analogy in the following paragraph does not take issue with this claim. Instead, it focuses on a different dissimilarity between the two situations: that, in *Arsenic (100,000)*, the \$200-per-capita investment eliminates a mortality risk to each cost-bearer and would be expected to save the life of one cost-bearer; by contrast, in *Trapped Miner*, the

regulating in *Arsenic (100,000)* were deemed uncertain—after all, who can really say whether one premature death would decrease overall well-being by more than imposing a \$200 cost on each of 100,000 people would?—fairness might provide an alternative basis for doing so.

However, there is a way of meeting this fairness-based argument. To begin with, *Trapped Miner* is importantly different from *Arsenic (100,000)*. In *Trapped Miner*, because the 100,000 cost-bearing individuals are not faced with a choice of whether to invest \$200 in their *own* safety, the norm of personal autonomy is not implicated in the same way it is in *Arsenic (100,000)*. One could not say to the trapped miner’s family that saving their loved one’s life would have required forcing their loved one to make an investment in her own safety that, from an *ex ante* standpoint, she could not rationally have wished to make. But, this is precisely the sort of justification that *could* be offered to the family of the deceased water consumer in *Arsenic (100,000)*.¹²⁶ Such a justification would focus on the point that, not only would preventing the death have required unjustifiably violating the autonomy of 100,000 people, but further that the unlucky victim would have been *among* those whose autonomy would have been violated.¹²⁷ “Knowing what she knew at the time the regulatory decision was made,” one might say to the family, “your loved one herself could not rationally have wished to pay her share of the cost of eliminating the risk that materialized in her death. Preventing your loved one’s death would have meant *forcing* her to engage in an irrational transaction: paying \$200 for a safety benefit worth just \$90. Thus, arguing that fairness required enacting the life-saving regulation is tantamount to arguing—rather absurdly—that it would have been unfair to your loved one *ex post* not to have disrespected her autonomy *ex ante*.” Of

\$200-per-capita investment would not have any effect on the mortality risks facing cost-bearers, nor would it save the life of any cost-bearer.

126. I borrow this analytical device from Sophia Reibetanz Moreau. See Reibetanz, *supra* note 124, at 303-04.

127. See SUNSTEIN, *supra* note 5, at 115 (“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 judgments, then they are insulting their autonomy. Suppose that people in a free society are entitled to have a kind of mastery over the conduct of their own lives. If so, then they should be permitted to make such allocations as they choose.”).

course, one might still insist that, for welfare maximization reasons, the regulation should be enacted, notwithstanding the autonomy-based reasons for opening it. However, given the apparent strength of the autonomy-based case against regulation, it is at least highly questionable whether this welfare maximization argument should prevail.

In *Arsenic (100,000)*, we saw that the government's regulatory decision was based not on risk-bearers' actual choice to assume a mortality risk rather than the pay the cost of eliminating the risk, but rather on studies of risk aversion and other risk-related preferences in the general population.¹²⁸ This reflects the way in which decisions are actually made by federal agencies when regulating mortality risks: they are based on studies of the amount persons in the general population are willing to pay to avoid small mortality risks.¹²⁹ Based on this sort of general information, we saw that an inference could be made about what the persons actually affected by a particular regulation (i.e., the 100,000 risk-bearers in *Arsenic (100,000)*) would have preferred with respect to the mortality risk in question. More specifically, we saw that, based on the general population data, it was reasonable to infer that each risk-bearer in those cases would very likely have preferred being exposed to the 1 in 100,000 mortality risk to paying the \$200 cost of eliminating that risk. Thus, the autonomy-based argument against regulating was founded on hypothetical consent to the mortality risk in question, not actual consent.

The question, then, is whether and to what extent the autonomy-based case against regulation in *Arsenic (100,000)* suffers by virtue of being based on hypothetical rather than actual consent. On one hand, the norm of autonomy seems to be more strongly implicated when an informed, free, and voluntary decision to assume (or pay to avoid) a particular mortality risk has *actually been made* by each potential risk-bearer¹³⁰ than it is

128. See *supra* Part II.D.2.

129. See, e.g., EPA Analysis, *supra* note 15, § 5.4.2, at 5-24 (“VSL does not refer to the value of an identifiable life, but instead to the value of small reductions in mortality risks in a population. . . . VSL estimates are appropriate only for valuing small changes in risk; they are not values for saving a particular individual’s life.”); see also *supra* Part I.A.

130. The connection between autonomy and actual consent is a strong one. See KAGAN, *supra* note 41, at 292 (“[I]f a person *consents* to a given form of treatment, then typically, at least, being treated in that way will not constitute a violation of autonomy. For

when, on the basis of data about the general population's risk aversion and preferences, a determination is made about what each risk-bearer *would have* decided in that regard.¹³¹ On the other hand, I believe it would be a mistake to think that the absence of an actual decision by risk-bearers about whether to assume the risk in question means that the norm of autonomy is not implicated at all.¹³² It would be a mistake, in other words, to think that a plausible determination as to what risk-bearers *would have* preferred under the circumstances should have *no* implications for whether the course of action under consideration would respect the risk-bearers' autonomy.¹³³

To see why, suppose the government were trying to decide whether to enact a regulation that would eliminate a 1 in 100,000 risk of death to 100,000 people at a cost of \$20,000 to each risk-bearer (for a total cost of \$2 billion). In such circumstances, should the government be required to actually put the question before each risk-bearer—and receive a negative answer from each one—in order to be able to defend a decision not to regulate on the basis of autonomy? It seems not. The fact that it is highly plausible to suppose that no risk-bearer *would be* willing to pay \$20,000 to avoid a 1 in 100,000 mortality risk supplies ample reason to believe that imposing the cost, rather than the risk, would violate the autonomy of each risk-bearer, even if it is not known for certain whether any of the risk-bearers *actually is* willing to do so. Enacting a regulation that the risk-bearers would almost certainly have objected to is little different from enacting a regulation that they actually objected to. If the latter would violate their autonomy, then the former arguably would too.

example, if I harm you with your permission, or if I fail to act on a promise from which you have released me, this doesn't actually interfere with your autonomy: you are still in control of your life, deciding how it is to go.”).

131. See P.J. Thomas, *Measuring Risk-Aversion: The Challenge*, 79 MEASUREMENT 285, 285-86 (2016).

132. See DWORKIN, *supra* note 111, at 88-89 (noting that, while hypothetical consent “is not actual consent” and “is simply a judgment about what the agent would have agreed to under certain circumstances,” hypothetical consent “appeal[s] to consent or something like consent in explaining the obligations and rights [it] create[s]”).

133. See SUNSTEIN, *supra* note 5, at 115 (“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 judgments, then they are insulting their autonomy.”).

In *Arsenic (100,000)*, by hypothesis, the general population data on which the government's willingness-to-pay determinations were based established that *very few* people would be willing to pay more than \$90 to avoid a mortality risk of 1 in 100,000. If that were true, it would follow that enacting the regulation would mean requiring the *vast majority* of risk-bearers to pay \$200 to avoid a mortality risk that, if given a choice, they would have instead chosen to assume. Arguably, this would disrespect the personal autonomy of the vast majority of risk-bearers. The fact that a decision not to regulate in *Arsenic (100,000)* is based on hypothetical rather than actual consent does not seem to undermine, in any significant way, an argument that this decision can be defended on the basis of autonomy.

Of course, autonomy works as a justification for using a VSL to guide the regulation of consumption risks only to the extent the government's willingness-to-pay data provides a secure basis for determining what the risk-bearers would or would not have consented to in a given set of circumstances. The urgent questions therefore become epistemic ones. In a given instance, how sure can the government be of the accuracy of its determinations about what risk-bearers would or would not have consented to under the circumstances? How can the government make its willingness-to-pay determinations sensitive to the fact that risk preferences will inevitably vary across different types of mortality risks,¹³⁴ as well as across different individuals depending on, among other factors, their wealth, socioeconomic status, age, and cultural preferences?¹³⁵ Even assuming the government were able to achieve a high degree of individuation regarding the willingness-to-pay preferences among a given group of risk-bearers, what figure should the government rely on in regulating the risk at issue: the average willingness-to-pay, the median willingness-to-pay, or some other statistical figure? Proponents of mortality risk CBA often speak in terms of *average* or *median* WTP of the risk-bearing population,¹³⁶ though other

134. See SUNSTEIN, *supra* note 5, at 96-101; Viscusi, *supra* note 10, at 152-69.

135. See SUNSTEIN, *supra* note 5, at 101-06; Viscusi, *supra* note 10, at 117-20, 133-34.

136. See, e.g., SUNSTEIN, *supra* note 5, at 95 (“[I]t would be grossly misleading to offer the following suggestion: *The value of a statistical life is \$9 million*. It would be much more accurate to say that for risks of 1/10,000, the median WTP in the relevant population is \$900—or that for risks of 1/100,000, the median WTP is \$90.”); Posner & Sunstein, *supra*

like-minded commentators appear to suggest that the relevant WTP is uniform.¹³⁷ If the basis of mortality risk CBA is autonomy, rather than welfare, it is difficult to see how the average or median WTP of the risk-bearing population could be a legitimate basis on which to set regulatory levels, because regulatory decisions would presumably contravene the preferences of the substantial proportion of risk-bearers with above-average or above-median risk aversion preferences.

These epistemic inquiries, although critically important, lie beyond the scope of this article. The discussion above establishes that, assuming the government can say with a reasonable degree of certainty how much money the persons exposed to a given consumption risk would have been willing to pay to avoid that risk, personal autonomy provides a sound normative basis—and seemingly the *only* sound normative basis—for using a VSL to guide the regulation of such a risk. This is a significant conclusion, one which challenges both the claim that using a VSL in mortality risk regulation has no normative basis in any moral context¹³⁸ and the claim that use of a VSL in the moral context of consumption risks can be defended on the basis of maximizing overall well-being.¹³⁹

III. WORKPLACE RISKS AND EQUITY

Commonly, an industrial activity poses a risk of death to a group of persons that benefits meaningfully from the activity, but does not bear any significant share of the costs of complying with regulations that reduce the risk.¹⁴⁰ Workplace risks are perhaps

note 26, at 551 (“Suppose, for example, that workers must be paid \$[9]00, on average, to eliminate a risk of 1/10,000. If so, the value of a statistical life would be said to be \$[9] million.”); Posner, *supra* note 36, at 324 (“Suppose that it is discovered by studies of people’s behavior that the average person would be willing to incur a maximum cost of \$1 to avoid the one-in-a-million chance of being killed by some hazard that a proposed project would eliminate.”).

137. See, e.g., Viscusi, *supra* note 36, at 105 (“[A] group of 10,000 people facing an individual risk of death of 1/10,000 will incur one expected fatality in the group. If *each person* is willing to pay \$500 to eliminate the risk, then there is a collective willingness to pay \$5 million to eliminate the one statistical death.” (emphasis added)).

138. See, e.g., sources cited *supra* note 8.

139. See, e.g., sources cited *supra* note 9.

140. Sunstein calls risk impositions of this type “harder cases,” since such cases tend to be more difficult to analyze than cases in which the persons exposed to a risk also bear the costs of reducing it. See SUNSTEIN, *supra* note 5, at 127 (“There is an obvious artificiality

the best example of this type of risk imposition.¹⁴¹ Many industries employ production processes that pose serious health risks to workers in those industries. For example, many industries use toxic chemicals that pose a risk of death or serious illness to the workers exposed to the chemicals in the course of performing their job duties.¹⁴² As with consumption risks, all or substantially all of the costs of complying with government regulations designed to reduce workplace risks tend to be passed on to consumers in the form of price increases or deducted from shareholder profits.¹⁴³ When this occurs, the group of people exposed to the risk in question (workers) is *not* the same group responsible for paying the costs of reducing that risk (consumers or shareholders). This is what makes workplace risks fundamentally different from consumption risks from a normative standpoint. Although workers obviously benefit from the industrial activity that places them at risk (because it provides them with gainful employment), the burden of complying with risk-reducing regulations tends to fall on others. These features

in the assumptions thus far. Most important, people do not always bear the full social costs of the regulatory benefits that they receive. . . . When this is so, the analysis is much more complicated.”).

141. However, environmental risks might also exhibit this structure. For example, the persons living in proximity to a plant that emitted carcinogenic fumes into the air might significantly benefit from the risk-generating industry, whether by consuming its products or in some other way. *See, e.g., id.* at 127-28 (offering air pollution regulation as an example of a case in which those exposed to the risk at issue do not bear the costs of regulations that reduce the risk).

142. A prominent example is the use of benzene in the oil refining industry. Benzene poses a cancer risk to workers exposed in the course of their duties. *See Indus. Union Dep’t. v. Am. Petroleum Inst.*, 448 U.S. 607, 613 (1980) (discussing OSHA’s regulation of workplace exposure to benzene).

143. *See Identification, Classification, and Regulation of Potential Occupational Carcinogens*, 45 Fed. Reg. 5002, 5237 (Jan. 22, 1980) (to be codified at 29 C.F.R. pt. 1990) (“The public, by and large, equally benefitted from and paid the costs for cleaner water. . . . However, in the case of occupational safety and health, different groups enjoy the economic savings of not regulating and take the risks. Consumers save through lower prices and employers benefit through higher profits from not regulating, while risks are borne by workers, often in the lower economic groups. *Therefore, occupational safety and health posed an ‘equity’ question.*” (emphasis added)) (summarizing testimony of Dr. Nicholas Ashford of MIT); *see also Occupational Exposure to Hexavalent Chromium*, 71 Fed. Reg. 10100, at 10280-281 (Feb. 28, 2006) (to be codified at 29 C.F.R. pts. 1910, 1915, 1917, 1918, & 1926) (“Price elasticity refers to the relationship between the price charged for a service and the demand for that service; that is, the more elastic the relationship, the less able is an establishment to pass the costs of compliance through to its customers in the form of a price increase and the more it will have to absorb the costs of compliance from its profits.”).

of the typical workplace risk create a particular moral context, one that is relevant to the justifiability of using a VSL to guide mortality risk regulation. In this Part, I argue that, when it comes to risks exhibiting this structure, using the basic concept of individual willingness-to-pay to guide mortality risk regulation is morally defensible on the basis of the norm of equity.¹⁴⁴ However, using a VSL to guide mortality risk CBA is not defensible in this context because it relies on a collective, rather than individual, willingness-to-pay determination.

A. Example: Chromium in the Workplace

Hexavalent chromium (Cr(VI)) is a toxic substance used in a number of industries.¹⁴⁵ Compounds containing chromium are used intentionally to perform metal electroplating and to produce chemical catalysts and pigments for textile dyes, paints, inks, glass, and plastics.¹⁴⁶ Chromium compounds are also formed incidentally as a byproduct of certain welding processes and as an impurity found in portland cement.¹⁴⁷ According to OSHA, there are over 30 industry sectors in which workers may be exposed to chromium.¹⁴⁸

Compounds containing chromium can exist in mist, dust, or fume form, and have long been known to pose health risks to workers when inhaled or upon contact with skin.¹⁴⁹ Most significantly, exposure to chromium can cause lung cancer, which

144. Cf. *Identification, Classification, and Regulation of Potential Occupational Carcinogens*, 45 Fed. Reg. at 5237, 5239 (“[O]ccupational safety and health pose[s] an ‘equity’ question. . . . [OSHA believes that] as a matter of policy, efficiency criteria alone are not appropriate because they ignore equity considerations. The economic savings from less protective regulation accrue to industry in the form of higher profits and consumers in the form of lower prices. But the costs are borne by workers through increased industrial illness and death rates.”); Waisman, *supra* note 19, at 1265-67 (arguing feasibility-based regulation of workplace mortality risks is defensible on the basis of the norm of equity).

145. See Pub. Citizen Health Research Grp. v. U.S. Dep’t. of Labor, 557 F.3d 165, 169 (3d Cir. 2009) (denying petition challenging OSHA’s regulation of hexavalent chromium except with respect to the employee notification aspects of the standard).

146. *Id.*

147. *Id.*

148. *Id.*

149. *Id.*

is often fatal.¹⁵⁰ Chromium exposure can also cause nonfatal ailments such as asthma, dermatitis, nasal irritation, and gastrointestinal ulcers.¹⁵¹

In 1971, OSHA set a permissible exposure level for chromium of 52 micrograms per cubic meter ($\mu\text{g}/\text{m}^3$).¹⁵² By 2004, chromium's carcinogenic properties had become clear, and OSHA consequently proposed reducing the chromium exposure standard to $1 \mu\text{g}/\text{m}^3$.¹⁵³ In 2006, after extensive comments and hearings, OSHA issued its final rule, which set a somewhat higher chromium PEL of $5 \mu\text{g}/\text{m}^3$.¹⁵⁴ This was the exposure level that OSHA ultimately found to satisfy the feasibility standard applicable under the Occupational Safety and Health Act.¹⁵⁵

As part of its analysis, OSHA calculated the annual costs and monetized benefits of setting the chromium PEL at the various alternative levels it was considering.¹⁵⁶ The benefits included the prevention of both fatal and nonfatal cancers.¹⁵⁷ To monetize the avoidance of fatal cancers, OSHA adopted the EPA's then-prevailing \$6.8 million VSL; to monetize the avoidance of nonfatal cancers, OSHA employed an analysis based on both the VSL and a cost-of-illness approach.¹⁵⁸ Discounting costs and

150. *Pub. Citizen Health Research Grp.*, 557 F.3d at 169-70; *see also Occupational Exposure to Hexavalent Chromium*, 71 Fed. Reg. 10100, 10224 (Feb. 28, 2006) (to be codified at 29 C.F.R. pts. 1910, 1915, 1917, 1918, & 1926).

151. *Occupational Exposure to Hexavalent Chromium*, 71 Fed. Reg. at 10108, 10166, 10174.

152. *Pub. Citizen Health Research Grp.*, 557 F.3d at 169.

153. *Id.*

154. *Id.*

155. *See id.*; *see also* 29 U.S.C. § 655(b)(5) (2012) (providing that, in promulgating regulatory standards for toxic materials or other hazardous substances, OSHA "shall set the standard which most adequately assures, to the extent feasible, on the basis of the best available evidence, that no employee will suffer material impairment of health or functional capacity" (emphasis added)).

156. *See Occupational Exposure to Hexavalent Chromium*, 71 Fed. Reg. at 10305, 10307.

157. *See id.* at 10305. Based on various scientific studies OSHA considered, it estimated ranges of the annual cancer-prevention benefits of setting the chromium PEL at the various alternative levels it was considering. For annual avoided lung cancer deaths, the reported ranges for each PEL were as follows: $.25 \mu\text{g}/\text{m}^3$: 66-258; $.5 \mu\text{g}/\text{m}^3$: 62-243; $1 \mu\text{g}/\text{m}^3$: 58-224; $5 \mu\text{g}/\text{m}^3$: 40-145; $10 \mu\text{g}/\text{m}^3$: 27-95; $20 \mu\text{g}/\text{m}^3$: 15-47. *Id.* at 10304, tbl.VIII-10. For annual avoided non-fatal cancers, the reported ranges for each PEL were as follows: $.25 \mu\text{g}/\text{m}^3$: 9-35; $.5 \mu\text{g}/\text{m}^3$: 8-33; $1 \mu\text{g}/\text{m}^3$: 8-31; $5 \mu\text{g}/\text{m}^3$: 5-20; $10 \mu\text{g}/\text{m}^3$: 4-13; $20 \mu\text{g}/\text{m}^3$: 2-6. *Id.*

158. *See id.* at 10305, 10307.

median monetized benefits by 7%,¹⁵⁹ OSHA determined that the greatest net benefit (\$6 million) would be realized at a chromium PEL of 10 $\mu\text{g}/\text{m}^3$.¹⁶⁰ Discounting costs and median monetized benefits by 3%, the greatest net benefit (\$231 million) would be realized at a PEL of 5 $\mu\text{g}/\text{m}^3$.¹⁶¹ Although the 5 $\mu\text{g}/\text{m}^3$ level may not have been cost-justified, OSHA fixed the chromium PEL at that level because, in its feasibility analysis, OSHA determined the 5 $\mu\text{g}/\text{m}^3$ level to be the lowest PEL economically and technologically feasible.¹⁶² As noted above, OSHA was mandated by federal statute to set the chromium PEL at the lowest feasible level.¹⁶³

OSHA determined that regulatory costs will generally either be passed on to consumers in the form of price increases, or absorbed by firms in the form of reductions to profits.¹⁶⁴ From OSHA's chromium report, it is difficult to determine either the proportion of regulatory costs that would be passed through to consumers or the amount of any per-unit price increases by which such pass-throughs would be accomplished. OSHA report does provide some helpful information, however. OSHA determined that firms in over 30 different industry sectors would be affected by the new chromium regulation.¹⁶⁵ Many of the firms in these industry sectors—e.g., those in the welding and construction sectors—may primarily serve other businesses, rather than individual consumers, so there may be multiple cost pass-

159. OSHA relied on an assumption of a 7% discount rate in its analysis of costs, adding an alternative 3% discount rate calculation as part of what it called a "sensitivity analysis" when calculating net monetized benefits. See *id.* at 10263 ("[A]ll costs are annualized at a discount rate of 7 percent. (A sensitivity analysis using a discount rate of 3 percent is presented in the discussion of net benefits.)").

160. See *id.* at 10306, tbl.VIII-11, 10308, tbl.VIII-12.

161. See *Occupational Exposure to Hexavalent Chromium*, 71 Fed. Reg. at 10306, tbl.VIII-11, 10308, tbl.VIII-12.

162. Elsewhere, I have argued that the 5 $\mu\text{g}/\text{m}^3$ chromium PEL on which OSHA settled was normatively defensible on the basis of equity. See Waisman, *supra* note 19, at 1306-12.

163. See 29 U.S.C. § 655(b)(5) (2012).

164. See, e.g., *Occupational Exposure to Hexavalent Chromium*, 71 Fed. Reg. at 10280-81 ("Price elasticity refers to the relationship between the price charged for a service and the demand for that service; that is, the more elastic the relationship, the less able is an establishment to pass the costs of compliance through to its customers in the form of a price increase and the more it will have to absorb the costs of compliance from its profits.")

165. See *id.* at 10272-82, tbl.VIII-7.

throughs to take into account before a regulatory cost reaches an individual “end user.”¹⁶⁶

Thus, OSHA’s regulation of chromium exposure exhibits the structure of the typical workplace risk.¹⁶⁷ The risky activity at issue—industrial use of chromium—provides a benefit to each risk-bearing worker (in the form of gainful employment) and to each cost-bearing consumer (in the form of a desirable product). Moreover, as OSHA noted, the costs of complying with chromium regulation would likely be passed on either to consumers of the regulated industries in the form of price increases or to shareholders of firms in those industries in the form of reductions to profits.¹⁶⁸ Thus, the workers exposed to cancer risks from chromium exposure benefit significantly from the industrial activity that gives rise to those risks, but do not bear the costs of compliance with regulations that reduce those risks. This makes the cancer risk posed by chromium exposure a good real-world example of a workplace risk.

B. The Norm of Equity and (Ex Ante) Contractualism

The term *equity* is defined in Webster’s Third New International Dictionary as “a free and reasonable conformity to accepted standards of natural right, law, and justice without prejudice, favoritism, or fraud and without rigor entailing undue hardship.”¹⁶⁹ Although often taken to be synonymous with fairness,¹⁷⁰ equity, as its etymology suggests, specifically connotes the aspect of fairness that has to do with equality.¹⁷¹ “[T]he essence of Justice or Equity,” observed renowned moral philosopher Henry Sidgwick, “is that different individuals are not to be treated differently, except on grounds of universal

166. See sources cited *supra* note 143.

167. See *supra* text accompanying notes 152-66.

168. See *Occupational Exposure to Hexavalent Chromium*, 71 Fed. Reg. at 10280-82.

169. *Equity*, WEBSTER’S THIRD NEW INTERNATIONAL DICTIONARY (2002). Throughout this Article, I use the term in this more popular sense, rather than in the technical legal sense that refers to the legal system and body of principles originating in the English Court of Chancery.

170. See, e.g., *Equity*, BLACK’S LAW DICTIONARY (10th ed. 2014) (defining equity as “[f]airness; impartiality; evenhanded dealing”).

171. The term derives from the Latin words *aequitas* (“equality”) and *aequus* (“equal”). See *Equity*, *supra* note 169.

application.”¹⁷² Equity is therefore a somewhat narrower concept than fairness, as it is specifically concerned with achieving fairness in the distribution of burdens and benefits.¹⁷³ Equity, as I understand the term here, is broadly concerned with equalizing the burdens borne by differently-situated individuals as the result of some socially desirable act, practice, or policy.¹⁷⁴

Like the norm of fairness, the norm of equity presupposes the evaluation of an action or policy from the point of view of each affected individual, rather than an impersonal evaluation from a “view from nowhere” or “God’s eye” perspective.¹⁷⁵ When determining whether an action or policy is equitable, one is implicitly and necessarily asking whether it is equitable *to* a particular affected party or *to* each and every affected party.¹⁷⁶ It is impossible to determine whether an action, policy, or rule is equitable, in a general sense, without assessing and comparing its impact on each affected individual.¹⁷⁷ This is to be contrasted with norms like overall well-being or efficiency, which do not require that the perspective of each affected individual be consulted and considered.¹⁷⁸ It would not make sense, for example, to ask whether an action or policy maximizes overall well-being *to* or *from the standpoint of* a particular party.

172. HENRY SIDGWICK, *THE METHODS OF ETHICS* 496 (1930).

173. By contrast, one can intelligibly speak of fairness in a non-distributive sense. For example, consider the notions of “fair play,” *see, e.g.*, *International Shoe Co. v. Washington*, 326 U.S. 310, 316 (1945) (observing that an assertion of personal jurisdiction must not offend “traditional notions of fair play and substantial justice” (citation omitted)), and “unfair surprise,” *see, e.g.*, 8 SAMUEL WILLISTON & RICHARD A. LORD, *A TREATISE ON THE LAW OF CONTRACTS* § 18:7 (4th ed. 2010) (discussing notion of “unfair surprise” in the context of the unconscionability doctrine in contract law).

174. *See* SIDGWICK, *supra* note 172, at 496-97.

175. *See, e.g.*, John Broome, *Fairness*, 91 *PROC. ARISTOTELIAN SOC’Y* 87, 87-96 (1991) (“Fairness is concerned only with how well each person’s claim is satisfied compared with how well other people’s are satisfied. It is concerned with relative satisfaction, not absolute satisfaction.” (emphasis omitted)).

176. This is reflected in the famous definition of justice offered by the Roman jurist Ulpian: “Justice is the constant and perpetual will to allot to every man his due.” A.B. Neil, *Justice and Natural Law*, 22 *TENN. L. REV.* 1025, 1025 (1951) (translating *DIG. 1.1.10* (Ulpian, *Regularum* 1) (“*Iustitia est constans et perpetua voluntas ius suum cuique tribuendi.*”)).

177. *See* Frick, *supra* note 42, at 221.

178. *See id.* (“[E]quity is an ‘individualistic’ moral notion. Being equitable is a property that attaches to actions, not in virtue of their overall or aggregate effects, but in virtue of how they treat each person individually.”).

However, what exactly equity requires in a particular case may not be deducible from these rather abstract observations. For purposes of my analysis in this article, I therefore adopt a particular interpretation of the notion of equity. Specifically, I will hold an action, practice, policy, or rule to be equitable if and only if it would be permitted by *contractualism*.¹⁷⁹ Introduced by the philosopher T.M. Scanlon in 1982,¹⁸⁰ contractualism is a relatively new theory of normative ethics founded on the social contract tradition in political philosophy (embodied in the work of Locke,¹⁸¹ Rousseau,¹⁸² and Rawls¹⁸³) and on the Kantian tradition in moral philosophy.¹⁸⁴

Contractualism is not a general theory of morality, but rather a theory of *interpersonal* morality or, to use Scanlon's famous phrase, of "what we owe to each other."¹⁸⁵ In this area of morality, Scanlon intended contractualism to offer an account of what it means for an action to be wrong.¹⁸⁶ In this article, I follow an alternative interpretation of contractualism, one originally offered by Derek Parfit and since endorsed by other philosophers.¹⁸⁷ The alternative interpretation holds that contractualism specifies a wrong-making property of an action, rather than explaining what it means for an action to be wrong.¹⁸⁸ On this view, an action might be wrong for non-contractualist reasons, or it might be wrong from a contractualist point of view but right all things considered.¹⁸⁹ This interpretation of

179. See *id.* at 220-21 (arguing that the notion of equity best captures the "wrong-making property" that contractualism specifies).

180. See generally T.M. SCANLON, WHAT WE OWE TO EACH OTHER (1998) (especially Chapter 5, pp. 189-247) [hereinafter SCANLON, WHAT WE OWE]; T.M. Scanlon, *Contractualism and Utilitarianism*, in UTILITARIANISM AND BEYOND 103, 103-14 (Amartya Sen & Bernard Williams eds., 1982) [hereinafter Scanlon, *Contractualism*].

181. See generally JOHN LOCKE, SECOND TREATISE OF GOVERNMENT (Richard H. Cox ed., Harlan Davidson 1982) (1690).

182. See generally JEAN-JACQUES ROUSSEAU, *On the Social Contract*, in BASIC POLITICAL WRITINGS 153-252 (Donald A. Cress ed. & trans., 2011).

183. See generally RAWLS, *supra* note 40.

184. See generally KANT, *supra* note 108, at 44-54.

185. Thus, contractualism does not directly address questions of *political* morality, the moral strictures applicable to the actions and policies of coercive institutions, nor does it address our moral obligations to animals or future persons.

186. See Scanlon, *Contractualism*, *supra* note 180, at 110.

187. See 1 DEREK PARFIT, ON WHAT MATTERS 368-70 (Samuel Scheffler ed. 2011); Frick, *supra* note 42, at 220.

188. Frick, *supra* note 42, at 220.

189. See *id.* at 222.

contractualism obviously squares well with normative pluralism. For the pluralist, contractualism specifies the requirements of equity, which is one of the normative factors that can determine the rightness or wrongness of an action.

According to contractualism, “an act is wrong if its performance under the circumstances would be disallowed by any set of principles for the general regulation of behavior that no one could reasonably reject as a basis for informed, unforced general agreement.”¹⁹⁰ What exactly does that mean? An example will be a useful place to start in answering this question.

Suppose the government is deciding between two different policies with which to accomplish a particular goal, Policy One and Policy Two. (Call this *Policy Choice*.) Both policies will affect a group of 1 million people, Group A, and a separate group of 100 people, Group B. At the time of the government’s decision, all members of both groups enjoy roughly equal levels of well-being and all members of both groups enjoy a decent, intermediate standard of living. Policy One would leave members of Group A much better off by increasing their standard of living from decent to high, but would leave members of Group B much worse off by lowering their standard of living from decent to low. Policy Two would leave members of Group A slightly better off by making a very small improvement in their standard of living, but would leave members of Group B slightly worse off by making a very small reduction in their standard of living. From a normative point of view, which policy should the government adopt?

The norm of welfare maximization provides a normative basis for choosing Policy One. The welfare-maximizing policy is the policy under which the aggregate (overall) well-being of all affected persons is greater than it would be under any alternative policy.¹⁹¹ In the above example, following this norm would mean

190. SCANLON, WHAT WE OWE, *supra* note 180, at 153.

191. The norm of welfare maximization, as I understand it here, is essentially equivalent to classical utilitarianism, except that it holds that the goodness of the outcome of an action or policy is to be evaluated on the basis of how it affects the well-being of individuals, rather than on how much pleasure or pain it produces. See *Consequentialism*, STAN. ENCYCLOPEDIA PHIL. (Oct. 22, 2015), <https://plato.stanford.edu/entries/consequentialism/> [https://perma.cc/68Z7-9F6E] (“When a welfarist theory of value is combined with the other elements of classic utilitarianism, the resulting theory can be called *welfarist consequentialism*.”).

calculating the net aggregate gain (or loss) in welfare associated with each policy, and choosing the policy that delivers the greater net gain (or smaller net loss) to aggregate welfare. It seems clear that both policies would increase aggregate welfare on net. This follows from the fact that, while each policy would result in a reduction in the well-being of 100 people (Group B), each policy would also deliver an equally-sized increase to the well-being of *1 million* people (Group A). For this reason, both policies would, from a welfare maximization point of view, plausibly result in a net improvement over the status quo.¹⁹² The question is which policy delivers the bigger improvement, i.e., the greater increase to overall well-being.¹⁹³ On that question, it seems clear that Policy One wins out. While Policy One would make each Group A member *much* better off, Policy Two would make each Group A member only *slightly* better off. Were Group A the only group of people affected, it is clear that Policy One would result in a significantly greater increase in aggregate well-being than Policy Two would. Does this change when the policies' respective effects on Group B are taken into consideration? Almost certainly not. Although Policy One would leave members of Group B worse off than Policy Two would, Group B is only a tiny fraction (1/10,000) of the size of Group A. The policies' respective effects on 100-person Group B therefore play a negligible role in determining how they would each affect the aggregate well-being of all 1,000,100 affected persons. For these reasons, Policy One would deliver a greater net increase to aggregate well-being than Policy Two would (or so I am assuming). The norm of welfare maximization therefore favors Policy One.

Contractualism provides a normative basis for choosing Policy Two. Contractualists believe that the normatively appropriate policy is the one that could be justified not to the entire group of affected persons considered as a whole, but to *each* affected person in light of how the policy, and the available alternatives, would affect him or her.¹⁹⁴ To use contractualist

192. To put this point in the language of welfare economics: although neither policy would represent a Pareto improvement over the status quo, both policies would result in a Kaldor-Hicks improvement.

193. See *Consequentialism*, *supra* note 191.

194. See *Contractualism*, STAN. ENCYCLOPEDIA PHIL. (Aug. 2, 2012), <https://plato.stanford.edu/entries/contractualism/> [<https://perma.cc/SH9C-YUBJ>].

language, the morally permissible policy is the one permitted by a principle that *no person could reasonably reject*.¹⁹⁵ Policy One is not justifiable to members of Group B because there is an alternative policy—Policy Two—under which no person is left as bad off as members of Group B would be left under Policy One. A contractualist would maintain that members of Group B could, on that basis, reasonably reject a principle licensing Policy One, making Policy One morally impermissible. Members of Group A might *wish* to reject a principle permitting Policy Two in favor of Policy One, but they could not *reasonably* do so, as they would implicitly be asking each member of Group B to put up with being left substantially worse off (rather than only slightly worse off) so that each member of Group A can avoid being left only slightly better off (rather than substantially better off).

The contractualist criterion of moral rightness involves three core ideas: (i) the idea that interpersonal morality presupposes the requirement of justifiability to *each* affected person considered as an individual, rather than the requirement of justifiability to *all* affected persons considered in the aggregate;¹⁹⁶ (ii) the idea that the moral status of a particular act (its rightness or wrongness) is a function of the moral validity of the *general principle* licensing the act;¹⁹⁷ and (iii) tying together the first two tenets, the idea that a principle is justifiable to each person if and only if it would command the free assent of all persons, i.e., if and only if no person could *reasonably reject* it as a principle for the general regulation of behavior.¹⁹⁸ Thus, under contractualism, an act is morally right if and only if no one could reasonably reject a general principle permitting the act.

When, according to contractualism, can a principle be reasonably rejected? The concept underlying the notion of

195. *Id.*

196. SCANLON, WHAT WE OWE, *supra* note 180, at 390 n.8 (“What is basic to contractualism as I understand it is the idea of justifiability to each person (on grounds that he or she could not reasonably reject).”).

197. *Id.* at 197 (“To justify an action to others is to offer reasons supporting it and to claim that they are sufficient to defeat any objections that others may have. To do this, however, is also to defend a principle, namely one claiming that such reasons are sufficient grounds for so acting under the prevailing conditions.”).

198. *Id.* at 197, 390 n.8.

reasonable rejectability is the *minimax criterion*.¹⁹⁹ Among a set of candidate principles, the non-rejectable principle P is the one of which the following is true: the strongest complaint any person could make against P, were P generally accepted, is weaker than the strongest complaint that could be made against every other alternative principle.²⁰⁰ As Scanlon puts it, “[S]omeone can reasonably reject a principle if there is some alternative to which no other person has a complaint that is as strong.”²⁰¹ The principle no one could reasonably reject is the principle that, among a set of candidate principles, *minimizes* the strength of the complaint that could be lodged by the *maximally* burdened person.²⁰² Supposing I am the person that would be most burdened by general acceptance of a particular principle, I still cannot *reasonably* reject that principle if every alternative principle would, if generally accepted, impose a greater burden on someone else.

Contractualism thus contemplates a *rejectability inquiry* the goal of which is to identify the principle or policy satisfying the minimax criterion. This is a fundamentally comparative inquiry that takes into account not just the extent to which general acceptance of each candidate principle burdens each affected person in an absolute sense, but also the differential each person experiences in the burdens they would bear under the respective principles.²⁰³ That is, supposing A is the most burdened party under Principle P-1 and B the most burdened party under Principle P-2, we ask not only if A’s burden under Principle P-1 is weightier than B’s burden under Principle P-2, but also if A’s gain in moving from P-1 to P-2 is more significant than B’s gain in moving from P-2 to P-1. The question to be asked, in other

199. I follow Sophia Reibetanz Moreau in using the minimax criterion to explain the notion of reasonable rejectability. See Reibetanz, *supra* note 124, at 300 (describing a “Minimax Complaint Model” of reasonable rejectability) (emphasis omitted).

200. See *id.*; see also Scanlon, *Contractualism*, *supra* note 180, at 111-12.

201. SCANLON, WHAT WE OWE, *supra* note 180, at 229; see also Scanlon, *Contractualism*, *supra* note 180, at 111 (“[I]t would be unreasonable . . . to reject a principle because it imposed a burden on you when every alternative principle would impose much greater burdens on others.”).

202. Applying the minimax rule to complaints or burdens is roughly (though not perfectly—see *infra* note 204 and accompanying text) equivalent to applying the more familiar “maximin” rule to outcomes for individual well-being.

203. See Scanlon, *Contractualism*, *supra* note 180, at 113; Reibetanz, *supra* note 124, at 299.

words, is whether it would be unreasonable for A to refuse to accept the loss she would experience in moving from P-2 to P-1 in order that B can avoid the loss she would experience in moving from P-1 to P-2.²⁰⁴

A critically important feature of contractualism—one that makes it particularly appropriate as an interpretation of equity—is what has been termed its *individualist restriction*, i.e., “the insistence that the justifiability of a moral principle depends only on various *individuals’* reasons for objecting to that principle and alternatives to it.”²⁰⁵ In other words, according to the individualist restriction, the strength of a complaint lodged against a particular principle can never be a function of the *sum* of different individuals’ gain (or loss) in well-being under that principle as compared with some alternative principle.²⁰⁶ Contractualism, in other words, does not allow for the interpersonal aggregation of complaints.²⁰⁷ Contractualism instead contemplates a series of “pairwise comparisons” in which *one* representative individual’s burden under a particular principle is compared to *one other* representative individual’s burden under an alternative principle.²⁰⁸ Thus, in *Policy Choice*, the fact that there are many more people in Group A than in Group B plays no role in determining which policy is justifiable to each person. That determination is made entirely on the basis of a pairwise comparison between representative members of each group.²⁰⁹

204. See Scanlon, *Contractualism*, *supra* note 180, at 123. For example, suppose that under P-1, A’s well-being is 102 and B’s 103 and that under P-2, A’s well-being is 150 and B’s 101. The mere fact that the loser under P-1 (A) is slightly better off than the loser under P-2 (B) does not necessarily mean that P-1 is the non-rejectable principle satisfying the minimax criterion. This is because A might plausibly be taken to have a stronger complaint with P-1 being chosen over P-2 than B would have with P-2 being chosen over P-1. Why? Because it would arguably be unreasonable for B to refuse to accept a 2-unit (roughly 2%) reduction in well-being in order that A can avoid a 48-unit (roughly 32%) reduction in well-being. Thus, B could not reasonably reject P-2.

205. See SCANLON, *WHAT WE OWE*, *supra* note 180, at 229.

206. See Reibetanz, *supra* note 124, at 300.

207. *Id.*

208. *See id.*

209. In cases involving the imposition of a risk of harm, as opposed to the imposition of harm itself, contractualism requires a further specification. Are the possible policies or courses of action to be evaluated on the basis of the *risk of harm* they pose to each person they *might* impact (so-called “ex ante contractualism”) or on the basis of the *actual harm* they cause to the person or persons they *do* impact (so-called “ex post contractualism”). For

C. Equity Justifies Use of Willingness-to-Pay, But Not a VSL

1. A Stylized Workplace Risk Scenario: Chromium (2)

Suppose the government is deciding whether to enact a chromium regulation that will affect two and only two people, W and C. W, the only worker employed in the industry in question, is exposed to chromium in the course of her work. C is the industry's only consumer. Suppose that, prior to regulation, chromium exposure poses a 1 in 50,000 mortality risk to W. The risk can be reduced in half, to 1 in 100,000, at a cost of \$40 to the industry, which cost will be entirely passed on to C in the form of a price increase. The risk can be entirely eliminated at a cost of \$240 to the industry, which cost will also be entirely passed on to C. (Call this *Chromium (2)*.) Suppose, as above, that the government has determined that, although very few people would be willing to pay *more* than \$90 to avoid a 1 in 100,000 mortality risk, the vast majority of people would be willing to pay \$90 to avoid a mortality risk of that magnitude. From the standpoint of morality, which of the following three actions is appropriate: not to regulate at all, to regulate so that the risk is reduced to 1 in 100,000 (moderate regulation), or to regulate so that the risk is eliminated (aggressive regulation)? The following chart shows the risk and cost burdens associated with each regulatory option.

reasons I have described elsewhere, I find ex ante contractualism to be the more plausible and defensible form of contractualism in cases involving imposition of a risk of harm. See Waisman, *supra* note 19, at 1291-95. A number of commentators have reached a similar conclusion. See, e.g., Frick, *supra* note 42, at 180, 219; Aaron James, *Contractualism's (Not So) Slippery Slope*, 18 LEGAL THEORY 263, 274, 292, & n.7 (2012); Rahul Kumar, *Risking and Wringing*, 43 PHIL. & PUB. AFF. 27, 48 (2015).

Chromium (2)

Regulatory Action	Worker's Burden	Consumer's Burden
No regulation	1 in 50,000 risk (WTP \geq \$180 ²¹⁰)	\$0
Moderate regulation	1 in 100,000 risk (WTP = \$90)	\$40
Aggressive regulation	No risk (WTP = \$0)	\$240

Notice first that the norm of personal autonomy is not implicated here in the same way it is in the case of consumption risks.²¹¹ Recall that in *Arsenic (1)*, the risk-bearer's autonomy-based objection to the risk-reducing regulation rested on the fact that the regulation's costs and benefits would fall entirely on her and her alone.²¹² For that reason, the risk-bearer's rational preference for assuming the risk, rather than paying the cost of avoiding it, could not be overridden without disrespecting her right to be sovereign in matters bearing exclusively on her own interests.²¹³ Absent a determination that the risk-bearer's preference was an irrational one, we saw that there did not seem to be any normative basis for forcing A to invest more in her own safety than she would wish to.²¹⁴

By contrast, if the government decides to regulate in *Chromium (2)*, C may be forced to pay a cost she does not wish to pay, but this would not involve forcing her to invest in her *own* safety. In *Chromium (2)*, by hypothesis, the costs of the regulation fall on C, but the associated risk-reduction benefit accrues to W. For this reason, were C to object to the regulation on autonomy grounds, the objection could not be based on C's

210. The assumption here is that, if W is willing to pay \$90 to avoid a 1 in 100,000 mortality risk, she would be willing to pay *at least* twice as much to avoid a 1 in 50,000 mortality risk. See *supra* note 36 (noting that willingness-to-pay increases nonlinearly with the magnitude of the mortality risk avoided).

211. See SUNSTEIN, *supra* note 5, at 130 (noting that, in cases where risk-bearers do not bear the costs of risk-reducing regulations, "it is much harder to argue that the use of cost-benefit analysis promotes autonomy").

212. See *supra* Part II.D.1.

213. *Supra* notes 114-19 and accompanying text.

214. See *supra* Part II.D.1.

right to be sovereign in matters affecting her own and only her own interests. The regulatory decision here affects C's interests, but it also clearly affects W's.

Of course, this does not mean that the norm of autonomy is not implicated *at all*. C might argue that forcing her to make any sort of investment in W's safety would violate her autonomy, because C might rationally prefer not to invest *anything* in reducing the risk chromium exposure poses to W. However, unlike the autonomy-based case against regulation in *Arsenic (1)*, the autonomy-based case against regulation in *Chromium (2)* faces a powerful counterargument based on the norm of equity, a norm which attends to the distribution of burdens and benefits among distinct persons.²¹⁵

This counterargument would begin by observing that C and W both benefit in a direct and meaningful way from the existence of the industry that, through its use of chromium, poses a risk of premature death to W. C benefits from consuming the industry's product. W benefits from gainful employment in the industry's production process. Because C and W both derive significant benefits from the industry's existence, it is reasonable to ask them both to bear a share of the burdens necessary to provide those benefits. C's burden takes the form of the monetary cost she must pay to purchase and consume the industry's product. W's burden takes the form of the risk she must bear to participate in the industry's production process. Notice that these burdens are inversely proportional: because the industry will typically pass on the costs of regulatory compliance to consumers, lessening the mortality risk W faces will mean increasing the cost C must pay for the industry's product. Because W and C have a mutual burden-sharing responsibility, the key question then becomes how to equitably *balance* the burdens that C and W respectively bear. Because C and W both benefit significantly from the activity, equity arguably requires that, as far as possible, their respective burdens be equalized.²¹⁶

215. See Waisman, *supra* note 19, at 1265-66.

216. If the net benefit W realizes from the activity is substantially greater than the net benefit C realizes (or vice versa), equity would arguably require that a larger share of the activity's mortality risk burden be borne by the party that realizes the larger net benefit. The analysis in this Part proceeds on the assumption that the net benefits realized by workers and consumers are roughly equal. Where this is not the case and there is a significant disparity in net benefit between worker and consumer, equity may require that a multiplier be used to

With these observations as background, it is not difficult to see why *W* would have a powerful equity-based argument for choosing moderate regulation over no regulation. Assume that *W* and *C* are identical with respect to their preexisting levels of wealth and their willingness to pay for reductions in mortality risk.²¹⁷ Assume further that *W* and *C* would both be willing to pay \$90 for a 1 in 100,000 reduction in mortality risk. Finally, assume *W* and *C* would both suffer the identical reduction in well-being as the result of bearing a given monetary cost or a given mortality risk. A decision not to regulate at all, rather than to regulate moderately, would mean denying *W* a safety benefit worth \$90 (to both *W* and *C*) for the sake of sparing *C* from a \$40 cost. From the standpoint of contractualism, the question is whether *C* could reasonably reject a principle requiring her to pay \$40 to spare *W* from a risk imposition that both *C* and *W* would, by hypothesis, be willing to pay \$90 to avoid. Put in terms of welfare rather than dollars, could *C* reasonably reject a principle requiring her to bear the welfare loss associated with paying a \$40 cost in order to spare *W* from experiencing the presumably *larger* welfare loss associated with a risk imposition that is equivalent to paying a \$90 cost?²¹⁸

It seems that she could not. As between no regulation and moderate regulation, moderate regulation is the course of action that satisfies contractualism's minimax requirement. As the chart

adjust the burden of risk regulation accordingly. For example, if, on average, each worker's net benefit from the activity were twice as large as each consumer's, it would arguably be inequitable to require a consumer to bear any regulatory cost *c* unless the associated reduction in each worker's mortality risk were one for which the worker would be willing to pay $2c$ or more. In *Chromium (2)*, the moderate regulation, which imposes a \$40 cost on the consumer, would satisfy this test because the safety benefit it provides, a 1 in 100,000 reduction in the worker's mortality risk, is one for which the worker would, by hypothesis, be willing to pay \$90, which is more than twice the amount of its \$40 cost.

217. These are, of course, unrealistic assumptions. I make them here just for purposes of stating the equity-based arguments in the simplest and clearest way possible. I later consider how differences in preexisting wealth levels and risk aversion would complicate the argument. See *infra* note 233 and accompanying text.

218. When asking this question, it is critical to bear in mind that *C* is not simply being asked to pay \$40 to reduce a mortality risk to someone else. Rather, *C* is being asked to pay \$40 to mitigate the mortality risk burden that an activity from which *C* benefits places on someone else. Although *C* is not herself imposing a mortality risk on *W*, the risk is imposed as a necessary incident of an industrial activity that provides a significant benefit to *C*. For that reason, as we saw above, it is reasonable to ask *C* to bear a share of the cost of reducing the risk to *W*.

above illustrates, under either no regulation or moderate regulation, W is the most burdened party. If no regulation is enacted, W bears a 1 in 50,000 mortality risk (which she would be willing to pay at least \$180 to avoid) and C experiences no increase in the cost of the product. If the government opts for moderate regulation, W bears a 1 in 100,000 mortality risk (which she would be willing to pay \$90 to avoid), and C experiences a \$40 increase in the product's cost. W is therefore the most burdened party under either regulatory scenario, but her burden is obviously lighter under moderate regulation. From the standpoint of the absolute level of the affected parties' burdens, the minimax criterion therefore favors moderate regulation. The minimax criterion also favors moderate regulation from a comparative standpoint. Compared to no regulation, moderate regulation delivers to W a risk-reduction benefit worth at least \$90, while imposing on C a monetary cost of just \$40. C could not reasonably refuse to put up with the \$40 cost of moderate regulation in order that W can enjoy the \$90 safety benefit that she would miss out on were the government to decide not to regulate at all.²¹⁹ For these reasons, contractualism would hold that W could reasonably reject a principle allowing for no regulation of chromium exposure at all, whereas C could not reasonably reject a principle allowing moderate regulation. Taking contractualism as an interpretation of equity, this means that it would be inequitable for the government to choose no regulation over moderate regulation. In other words, a decision not to regulate chromium exposure at all, rather than to regulate moderately, would mean asking W to bear more than her fair share of the burden of the industrial activity that benefits both C and W in a meaningful way.

If equity favors moderate regulation over no regulation in *Chromium (2)*, it also favors moderate regulation over aggressive regulation. Regulating aggressively would reduce W's mortality risk by 1 in 50,000 at a cost of \$240 to C. Compared with moderate regulation, aggressive regulation would reduce W's mortality risk by an additional 1 in 100,000 at an incremental cost

219. Cf. Scanlon, *Contractualism*, *supra* note 180, at 123 ("The question to be asked is, is it unreasonable for someone to refuse to put up with the Losers' situation under [one principle] in order that someone else should be able to enjoy the benefits which he would have to give up under [an alternative principle]?").

of \$200 to C. Could C reasonably reject a principle requiring her to pay an additional \$200 to provide W with an additional risk-reduction benefit for which W herself would be willing to pay no more than \$90? In other words, could C reasonably refuse to accept the welfare loss associated with paying a \$200 cost so that W can avoid the *smaller* welfare loss associated with a risk imposition that is equivalent to paying a \$90 cost?

It seems that she could. As between moderate regulation and aggressive regulation, moderate regulation is the course of action that satisfies the minimax criterion. Although W is the most burdened party under moderate regulation, C is the most burdened party under aggressive regulation. W's burden under moderate regulation is a 1 in 100,000 mortality risk, while C's burden under aggressive regulation is a cost of \$240. Since, by hypothesis, C and W would be willing to pay no more than \$90 to avoid a 1 in 100,000 mortality risk, C's monetary burden under aggressive regulation is arguably weightier than W's mortality risk burden under moderate regulation. Thus, from an absolute standpoint, moderate regulation would seem to be the course of action that minimizes the extent of the burden borne by the maximally burdened party. The same is true from a comparative standpoint. Compared to moderate regulation, aggressive regulation lowers W's mortality risk by 1 in 100,000 while increasing C's monetary burden by \$200. In other words, aggressive regulation would mean imposing an incremental burden of \$200 on C for the sake of providing an incremental risk reduction benefit of \$90 to W. To do so would be inequitable. C could reasonably refuse to accept a \$200 cost for the sake of delivering a risk reduction benefit worth just \$90 to W. To ask C to accept this burden would be to ask her to invest more in W's safety than W *herself* would have been willing to invest in her *own* safety. If W herself would not have been willing to pay more than \$90 to avoid a 1 in 100,000 mortality risk, it seems unreasonable to ask C to pay \$200—over twice as much—to provide W with a risk reduction of that magnitude.

Thus, adopting ex ante contractualism as an interpretation of equity, the most equitable course of action in *Chromium (2)* is for the government to regulate moderately, reducing W's risk of death from chromium exposure from 1 in 50,000 to 1 in 100,000 at a cost of \$40 to C. Neither W nor C could *reasonably* reject a

principle allowing for moderate regulation because both of the other regulatory alternatives—no regulation and aggressive regulation—involve imposing a greater burden on some party (W and C, respectively) than moderate regulation imposes on the party it burdens most heavily (W). This is just what it means to say that, among all three regulatory alternatives, moderate regulation is the alternative that satisfies the minimax criterion.

Recall that C might object to regulation in *Chromium (2)* on autonomy grounds.²²⁰ A decision to regulate moderately or aggressively would arguably violate C's autonomy by requiring her to pay a cost she may not wish to pay. It is now clear that this objection could be met forcefully on the basis of equity. If C wishes to avail herself of the benefits of the chromium-using industry, she can reasonably be expected to bear an equitable share of the associated burden. She is not being forced to purchase the good or service offered by the chromium-using industry. Once she chooses to do so, she can reasonably be asked to bear her fair share of the burden the industry imposes on society. Contractualism's minimax criterion represents a plausible specification of what each party's equitable share of the industry's mortality risk burden would be. Any violation of C's autonomy resulting from adherence to the minimax criterion in setting the appropriate level of regulation would therefore seem to be justifiable on the basis of equity. Thus, from the standpoint of normative pluralism, moderate regulation would seem to be the right course of action, all things considered.

What general point or principle does this analysis of *Chromium (2)* establish? The most fundamental point is that the norm of equity requires that individuals' willingness-to-pay to avoid small mortality risks—the current basis of the VSL—serve as a guide for the regulation of workplace risks, just as the norm of personal autonomy requires that willingness-to-pay serve as a guide for the regulation of consumption risks.²²¹ This general point can be formulated as a principle that I have elsewhere dubbed the *individual risk principle*: when it comes to workplace risks, equity requires that each cost-bearing consumer make all and only those investments in risk-reduction that each risk-

220. See *supra* notes 211-15 and accompanying text.

221. See Waisman, *supra* note 23, at 676-79.

bearing worker would have been willing to make on her own behalf.²²² As noted above, this principle rests on the premise that, because worker and consumer both benefit meaningfully from the industrial activity that gives rise to the risk, it is fair to ask both to bear a fair share of the activity's mortality risk burden.²²³ According to the individual risk principle, the way to accomplish an equitable distribution of this burden is to require each consumer to invest in reducing each worker's mortality risk up to the point at which the worker herself would cease investing were she to bear each consumer's share of the regulatory costs.²²⁴

The individual risk principle must be qualified in three important ways, however. First, as formulated above, the individual risk principle assumes that (1) the worker and consumer both would suffer the identical reduction in welfare as the result of the imposition of a given monetary cost or a given mortality risk (2) the worker and consumer have the identical willingness to pay to avoid mortality risks of varying magnitudes (e.g., they would both be willing to pay no more than \$90 to avoid a mortality risk of 1 in 100,000).²²⁵ In the real world, these assumptions are unlikely to hold true. Levels of preexisting wealth and income may vary significantly among those individuals likely to be affected by a given workplace risk regulation.²²⁶ Such differences are likely to affect the amount individuals are willing to pay to avoid a given mortality risk. This follows from the general point that willingness-to-pay is a function of ability-to-pay.²²⁷ Moreover, disparities in pre-existing levels of wealth or income may impact the extent of the welfare loss individuals experience as the result of the imposition of a given monetary cost. A \$200 cost will almost certainly have

222. *See id.* (describing the individual risk principle).

223. *Id.* at 677.

224. *See id.* (noting that individual risk principle represents a "single-owner" approach to risk regulation); *cf.* Kenneth W. Simons, *Deontology, Negligence, Tort, and Crime*, 76 B.U. L. REV. 273, 282 (1996) (describing a "'single-owner' conception" of cost-benefit analysis "which asks what the actor . . . would do if he owned all the resources in question and would therefore internalize all the costs and benefits of the decision").

225. *See supra* notes 221-24 and accompanying text; *see also* Waisman, *supra* note 23, at 676.

226. *See* Waisman, *supra* note 23, at 679.

227. *See* Mark S. Thompson, *Willingness to Pay and Accept Risks to Cure Chronic Disease*, 76 AM. J. PUB. HEALTH 392, 392 (1986).

a greater effect on the well-being of a very poor person than on that of a very rich person.

For example, suppose W in *Chromium (2)* is quite poor, while C is quite wealthy. Even assuming W and C have identical degrees of aversion to risks of premature death, W would plausibly be willing to pay significantly less to avoid a mortality risk of 1 in 100,000 than C would. Suppose that, while C would be willing to pay \$90 to avoid a risk of that magnitude, W would only be willing to pay \$30. Should this give C a sound equity-based objection to moderate regulation, which would require her to invest \$10 more in W's safety than W herself would have been willing to invest in her own safety, thereby violating the individual risk principle? Plausibly not. It is very possible that, while the imposition of a 1 in 100,000 mortality risk would reduce both W and C's welfare by the same amount, C would be willing to pay significantly more to avoid a risk of that magnitude due solely to her greater ability-to-pay to avoid mortality risks generally. In other words, it is quite possible that, if W had C's greater ability-to-pay, W would be willing to pay \$90, rather than just \$30, to avoid the mortality risk to which she is subject. If that were true, it would be inequitable to W to allow C to escape the cost of moderate regulation due solely to W's lesser ability-to-pay. Because the setback to W's welfare due to the 1 in 100,000 mortality risk would be greater than the setback to C's welfare due to the \$40 cost of eliminating that risk, the equitable course of action is moderate regulation.

What this shows is that, in applying considerations of equity to mortality risk regulation, the weight of the burdens that differently situated individuals are determined to bear should really be a function of the individuals' welfare or welfare-related interests. Willingness-to-pay is just a proxy for welfare. As Cass Sunstein has noted, "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."²²⁸ The individual risk principle might, then, be reformulated as follows in terms of welfare rather than willingness-to-pay: in the case of workplace risks, equity requires that each cost-bearing consumer invest in risk-reducing

228. SUNSTEIN, *supra* note 5, at 129.

regulation up to the point at which further investment would reduce the well-being of each consumer by more than it would increase the well-being of each risk-bearing worker. Equivalently, equity requires that each cost-bearing consumer make all, and only those, risk-reducing investments in worker safety that a worker would have been willing to make in her own safety, *were the worker in the economic situation of the cost-bearer*.²²⁹

The second qualification of the individual risk principle is that it must be defeasible in circumstances where it would prescribe a level of regulatory investment so great as to threaten to bankrupt the regulated industry, assuming the industry's existence is normatively desirable.²³⁰ For example, suppose the \$40 price increase associated with enacting the moderate regulation in *Chromium (2)* were enough to dissuade C from purchasing the industry's good or service altogether. This reduction in consumer demand could make the industry unprofitable and thereby drive it out of business. If the total social benefit the industry delivers exceeds the total cost it imposes on society,²³¹ bankrupting the industry might be a normatively undesirable result. In such circumstances, capping regulatory investment *below* the level dictated by the individual risk principle—and thereby preserving the industry—could therefore be defended on the basis of the norms of welfare maximization, equity, or both.

The final qualification relates to normative pluralism. The individual risk principle is meant to represent a specification of what the norm of equity (interpreted as *ex ante* contractualism) requires in circumstances of the typical workplace risk, i.e., circumstances in which risk-bearers do not bear the costs of risk-reducing regulations, but both risk-bearers and cost-bearers benefit meaningfully from the industrial activity that gives rise to

229. See Waisman, *supra* note 23, at 677.

230. See *id.* at 681 (noting that regulatory costs should be capped “short of the level required by the individual risk principle if this is necessary to preserve the underlying socially beneficial activity”); Waisman, *supra* note 19, at 1316-31 (same).

231. See Waisman, *supra* note 19, at 1317-20 (describing the notion of all-things-considered beneficialness, i.e., whether the industry in question is socially beneficial taking into account considerations of overall well-being, equity, and other relevant normative factors).

the risk.²³² Equity is just one of the several norms that may be applicable in such circumstances, and normative pluralism makes room for the possibility that, in some cases, the equitable course of action may *not* be the morally right course of action all things considered.²³³ Thus, there may be circumstances in which the individual risk principle would dictate a course of action that turns out, on a complete normative analysis, to be wrong. For example, suppose a 1 in 1,000 mortality risk to 1,000 workers can be eliminated by enacting a regulation costing \$5 billion, the cost of which would be spread equally among 1 million consumers for a per-capita cost of \$5,000. The individual risk principle would require enacting the regulation, since, taking \$9 million as a VSL, each worker would be willing to spend up to \$9,000 to avoid the mortality risk to which they are subject, a figure substantially in excess of the \$5,000 each consumer would have to pay were the regulation enacted. However, in such circumstances, a very powerful argument against regulation could be made based on the norm of welfare maximization. Imposing a \$5 billion cost on 1 million people by requiring each person to pay \$5,000—which represents a significant expense for most people—would arguably lower overall well-being by significantly more than the combination of a single expected death and the imposition of a 1 in 1,000 mortality risk on 1,000 people would. For this reason, welfare maximization considerations might plausibly trump equitable considerations in this instance, making it morally impermissible to enact the regulation. This example is offered to underscore that the individual risk principle merely specifies the requirements of equity as part of a broader, normatively pluralistic theoretical framework. The individual risk principle is not here claimed to be a regulatory principle that ought to be adhered to in every instance.

2. A More Realistic Scenario: Chromium (200,000)

The next question is whether the equity-based case for moderate regulation in *Chromium (2)* loses any of its force when we move to a more realistic workplace risk scenario in which chromium poses a 1 in 50,000 mortality risk to 100,000 workers

232. See Waisman, *supra* note 23, at 677.

233. See Waisman, *supra* note 19, at 1303.

in a given industry, with the costs of regulation spread among 100,000 industry consumers in the form of a \$40-per-capita price increase (for the moderate regulation) and a \$240-per-capita price increase (for the aggressive regulation). (Call this *Chromium (200,000)*.) The chart below illustrates the individual and aggregate burdens associated with each regulatory option. Notice that, because of the much larger number of risk-bearers involved, a choice to regulate in *Chromium (200,000)* carries the expected benefit of saving actual lives (one life in the case of moderate regulation, two in the case of aggressive regulation).

Chromium (200,000)

Regulatory Action	Each Worker's Ex Ante Burden	Each Consumer's Burden	Ex Post Burden on 100,000 Workers	Total Burden on 100,000 Consumers
No regulation	1 in 50,000 risk (WTP \geq \$180)	\$0	2 expected deaths	\$0
Moderate regulation	1 in 100,000 risk (WTP = \$90)	\$40	1 expected death	\$4 million
Aggressive regulation	No risk (WTP = \$0)	\$240	0 expected deaths	\$24 million

In *Chromium (2)*, we saw that equity, interpreted as ex ante contractualism, decisively favored moderate regulation over no regulation.²³⁴ In *Chromium (200,000)*, the equity case for moderate regulation is arguably even stronger, because a decision not to regulate in *Chromium (200,000)* would mean treating not just one but 100,000 workers inequitably.²³⁵ Moreover, in

234. See Part III.C.1.

235. Cf. *supra* Part II.D.2 (making similar point about the autonomy-based case against regulation in *Arsenic (100,000)*, which would involve 100,000 instances of the autonomy violation involved in *Arsenic (1)*).

addition to equity, welfare maximization arguably provides an independent normative basis for moderate regulation in *Chromium (200,000)*.²³⁶ Choosing not to regulate would arguably deliver a greater setback to overall well-being than moderate regulation would.²³⁷ That is because the imposition of a \$40 cost on each of 100,000 people would plausibly decrease overall well-being by *less* than the combination of one (expected) death and the imposition of an additional 1 in 100,000 mortality risk on 100,000 people would. Thus, in *Chromium (200,000)*, both equity and welfare maximization arguably favor moderate regulation over no regulation.

The more interesting question is whether the case for aggressive regulation is stronger in *Chromium (200,000)* than in *Chromium (2)*. After all, regulating aggressively instead of moderately in *Chromium (200,000)* would mean saving one additional life.²³⁸ For this reason, welfare maximization might well favor aggressive over moderate regulation. As we saw in connection with *Arsenic (100,000)*, the imposition of a \$200 cost on each of 100,000 people would arguably decrease overall well-being by less than the combination of one (expected) death and the imposition of an additional 1 in 100,000 mortality risk on 100,000 people would.²³⁹ If this is so, the case for choosing to regulate moderately rather than aggressively in *Chromium (200,000)* would have to rest entirely on equity, just as the case against regulation in *Arsenic (100,000)* rested entirely on personal autonomy. Is the equity-based case for moderate regulation in *Chromium (200,000)* strong enough to overcome the arguable welfare maximization case in favor of life-saving aggressive regulation?

In *Arsenic (100,000)*, the case against regulation rested on the fact that saving the life of the unlucky consumer killed by arsenic would have necessitated violating that consumer's autonomy *ex ante*.²⁴⁰ Because a decision not to regulate was in the *ex ante* interest of each of the 100,000 cost-paying consumers, a choice to regulate would have meant forcing 100,000 people to

236. See *supra* tbl. Chromium (200,000).

237. See *supra* tbl. Chromium (200,000).

238. See *supra* tbl. Chromium (200,000).

239. See *supra* Part II.D.2.

240. See *supra* Part II.D.2.

make an investment in their own safety that none of them would rationally have been willing to make.²⁴¹ To the family of the unlucky consumer killed as the result of a decision not to regulate, that decision could be justified by pointing out that saving their loved one's life would have required violating their loved one's autonomy ex ante.²⁴²

What makes aggressive regulation in *Chromium (200,000)* different is that the risk at issue is a workplace risk, not a consumption risk. The costs of regulation are borne by consumers who are not exposed to the risk at issue, rather than by the risk-exposed workers who benefit from risk-reducing regulations.²⁴³ Thus, unlike the decision not to regulate in *Arsenic (100,000)*, the decision to regulate moderately rather than aggressively in *Chromium (200,000)* cannot be justified on grounds of personal autonomy. To the family of the unlucky worker killed when the mortality risk tolerated under moderate regulation materializes, one could not justify moderate regulation by pointing out that saving their loved one's life would have meant forcing her to invest more in her own safety than she would rationally have wished to. Moreover, unlike the choice not to regulate in *Arsenic (100,000)*, a choice to regulate moderately in *Chromium (200,000)* does not work to the ex ante advantage of the person unlucky enough to lose their life as the result of that choice.²⁴⁴ Whereas opting not to regulate in *Arsenic (100,000)* spares each risk-bearing consumer from a \$200 cost, opting for moderate rather than aggressive regulation in *Chromium (200,000)* provides no such monetary benefit to the risk-bearing workers.²⁴⁵

In arguing that the government should have chosen to regulate aggressively, the unlucky worker's family might therefore draw an analogy between *Chromium (200,000)* and the case of the *Trapped Miner*, discussed in Part II. Arguably, if it would be unreasonable for the 100,000 town residents to refuse to pay \$200 each to save the trapped miner's life, then it would be similarly unreasonable for the 100,000 consumers in

241. See *supra* Part II.D.2.

242. See *supra* Part II.D.2.

243. See *supra* Part III.C.2.

244. See *supra* Part III.C.2.

245. See *supra* Part III.C.2.

Chromium (200,000) to refuse to each accept a \$200 price increase to save the life of a worker. Both scenarios involve a group of 100,000 people being asked to pay \$200 each to save the life of a person who is not among the 100,000 would-be cost-bearers (which is what makes both scenarios relevantly different from *Arsenic (100,000)*). In both cases, it would arguably be inequitable to allow the death of one person to occur simply to spare each of 100,000 people from having to pay a cost of \$200. Thus, the family might argue that the equitable course of action in *Chromium (200,000)* is aggressive regulation.

This analogy between *Trapped Miner* and *Chromium (200,000)* rests on an implicit equivalence between identified and statistical lives. The implicit claim is that there is no *morally relevant* difference between an action that will avoid a death otherwise certain to befall an identifiable person, and an action that will avoid a death otherwise statistically expected to befall an ex ante unidentifiable person.²⁴⁶ Is this claim defensible? Does it make any difference, from a moral point of view, that each of the 100,000 cost-bearers in *Trapped Miner* is being asked to pay \$200 each to avoid certain death for a person who is, at the time of the request, identifiable, whereas each of the 100,000 cost-bearers in *Chromium (200,000)* is being asked to pay \$200 each to avoid the expected death of a person who is unidentifiable at that time?

One possible reason it *might* make a difference is that a death that is statistically expected to occur is not *certain* to occur. As we saw above, when a 1 in 100,000 mortality risk is imposed on 100,000 people, there is a substantial chance—roughly 36.8%—that *no* death will occur.²⁴⁷ It might plausibly be argued that the moral obligation to avoid a certain death is greater than the moral obligation to avoid a death that has only a 63.2% chance of occurring. However, this distinction is rather unsatisfying because many mortality risks *are* practically certain to result in the death of at least one person. For example, when a 1 in 10,000 risk of death is imposed on 100,000 people, 10 deaths are statistically expected, but the chance of at least *one* death occurring is extremely high: 99.99996% ($1 - .9999^{100,000}$). If

246. On the difference between identified and statistical lives, see sources cited *supra* note 42. See also Frick, *supra* note 42, at 212 n.41.

247. See *supra* notes 94-95.

Chromium (200,000) were modified so that aggressive regulation would involve the elimination of a 1 in 10,000 mortality risk to 100,000 workers, it begins to look very similar to *Trapped Miner* with respect to the certainty that the regulatory action will avoid at least one death. Many mortality risks that are the subject of government regulation are similar in this respect—they are statistically expected to result in multiple deaths and practically certain to result in at least one death.²⁴⁸

Suppose then, for the sake of argument, we assume that the aggressive regulation in *Chromium (200,000)* were *certain* to result in the saving of a life that would otherwise have been lost due to the chromium exposure tolerated under moderate regulation. With this assumption in place, does it remain a morally relevant difference that the identity of the person whose life is to be saved in *Trapped Miner* is known at the time of the decision, whereas the identity of the person whose life is to be saved in *Chromium (200,000)* is not?

From one standpoint, it seems irrelevant whether this particular piece of information—the identity of the person whose life will be saved by the action under consideration—is known at the time of the decision. If we know for certain that the action under consideration will save someone’s life, what does it matter whether we know that person’s name, age, gender, ethnicity, religion, or other identifying characteristics?²⁴⁹

However, there is another difference between *Trapped Miner* and *Chromium (200,000)* which *does* seem relevant. The difference relates to the fact that, at the time the payment decision is to be made in *Trapper Miner*, there is at least one way of identifying the individual whose life will be lost absent action being taken: one can say, “It is the person trapped at the bottom of this mine.” One could pick out the person in this way even if

248. See Lisa A. Robinson, *How US Government Agencies Value Mortality Risk Reductions*, 1 REV. ENVTL. ECON. & POL’Y 283, 283-85 (2007) (explaining how government agencies calculate risks and create regulations). These risks are based on deaths statistically expected but, as demonstrated above, the likelihood that at least one death will occur is extremely high, 99.99996%.

249. See Reibetanz, *supra* note 124, at 304 (“[W]hen we know that someone will be harmed or benefited in a certain way, we should assign him a complaint based upon the full magnitude of that harm or benefit. It is irrelevant, for moral purposes, whether we can identify the bearer of this complaint in advance.”); see also Frick, *supra* note 42, at 184 (discussing argument that “the identity of the eventual victims is *irrelevant information*”).

one knew none of his or her identifying characteristics (other than being a live human being). By contrast, at the time the regulatory decision is to be made in *Chromium (200,000)*, there is *absolutely no way* of identifying or picking out the unlucky worker whose life will be lost if the government opts for moderate rather than aggressive regulation. At that time, all it is possible to say is that the unlucky worker, whoever she may turn out to be, is presently among a group of 100,000 workers, each of whom faces a 1 in 100,000 risk of premature death if chromium is regulated moderately.

This difference matters from the standpoint of equity. When one decides whether equity requires me to assume a burden in order to provide a benefit to you, one relevant question may be whether *you* would have been willing to assume that burden in order to provide the identical benefit to yourself.²⁵⁰ Call this the *single-person test*, as it involves imagining that the burden and benefit in question would both be experienced by a single person. The question is, had you been faced with the choice whether to assume a particular burden in order to provide yourself with a corresponding benefit, would you have been willing to assume that burden? If so, then equity *may* require that I assume that burden for your benefit. Contractualism's minimax criterion, which is embodied in the individual risk principle, represents a more-or-less straightforward application of this single person inquiry.²⁵¹

Of course, the answer to the single-person test need not be *dispositive* of the question of whether equity requires me to assume a burden for your benefit. In many circumstances, it seems that I could equitably reject bearing a burden that you would have been willing to assume for your own benefit. Certainly, your willingness to pay a large sum of your own money for the benefit of owning a new car does not mean I would be acting inequitably by refusing to pay the same sum of *my* money to buy the car for you. And one can imagine circumstances in which equity would require me to assume a *greater* burden for your benefit than you yourself would have been willing to

250. See Waisman, *supra* note 19, at 1284-85 (demonstrating equity requires an examination of the benefits and burdens to different individuals).

251. See *supra* notes 199-203 and accompanying text (illustrating minimax criterion as one that requires a principle be accepted if other persons would accept it).

assume. If I negligently run you over with my car and cause you serious injuries, treating you equitably may require that I provide you with more expensive medical care than you would have been willing to provide for yourself, even assuming you and I have identical levels of pre-existing wealth.

But there are circumstances in which the single-person test *does* seem to provide a very good gauge of what equity requires. As noted above, when two people (or two groups of people) both derive meaningful and roughly equal benefit from an activity that imposes a significant burden on either or both of them, equity arguably requires that the burdens be, as far as possible, equally distributed between them.²⁵² As discussed, the typical workplace risk is a good example of such circumstances. In such circumstances, the single-person test can be useful in determining what an equitable distribution of the burden would look like.²⁵³

The single-person test raises an epistemic issue, however. What information should be assumed to be available to the hypothetical person when she makes the decision whether to assume the burden in question for her own benefit? One plausible answer is that the person should be assumed to know all and only the information that was knowable at the time the decision had to be made. This is, after all, the epistemic constraint on decisions made in real life.²⁵⁴ The single-person test could then be reformulated in the following way: how would the person have acted on their own behalf at the time of the decision, knowing all and only what was knowable at that time? In other words, knowing everything relevant and knowable at the time of the decision, would you have been willing to assume the burden in question for your own benefit?

Formulating the single-person test in this way allows us to understand why, for purposes of equity, *Trapped Miner* is importantly different from *Chromium (200,000)*. In *Trapped Miner*, at the time the payment decision is to be made, there is

252. See *supra* note 216 and preceding text.

253. See *supra* notes 167-68 and accompanying text (providing an example of typical workplace risk); see also *supra* notes 250-51 and accompanying text (showing single person test requires a balancing of the burdens to determine what is equitable).

254. See generally Hugh Courtney et al., *Deciding How to Decide*, HARV. BUS. REV., Nov. 2013, <https://hbr.org/2013/11/deciding-how-to-decide> [<https://perma.cc/J4KR-L6V6>] (examining process of decision making and how persons are limited by information known at the time).

information available to all persons—including the potential victim—about who it is that will die should the town residents choose *not* to pay the \$200-per-capita cost of the rescue: the person presently trapped at the bottom of the mine. Thus, were the town residents to apply the single-person test at the time their payment decision was to be made, the emphatic answer would be affirmative: knowing all and only what was knowable at that time, the would-be beneficiary of the payment decision (the trapped miner) would undoubtedly have been willing to assume the burden of paying \$200 in order to receive the benefit of avoiding certain premature death. Were the town residents to choose not to assume that burden, they would do so in the knowledge that, *at the very time they made that choice*, there was a person who, based on all and only the information available at that time, would have been willing to assume that burden on her own behalf. For that reason, equity would arguably require that each town resident make a \$200 payment to save the miner's life.²⁵⁵

By contrast, at the time of the regulatory decision at issue in *Chromium (200,000)*, no would-be beneficiary of a decision to regulate aggressively rather than moderately would, knowing all and only what was knowable at that time, have been willing to assume a \$200 burden in order to receive the corresponding benefit of eliminating a 1 in 100,000 mortality risk. This is precisely what *Arsenic (100,000)* established: if the most that any of 100,000 risk-bearers would be willing to pay to avoid a 1 in 100,000 mortality risk is \$90, then each risk-bearer would presumably choose to assume a risk of that magnitude rather than pay a cost of \$200.²⁵⁶ Assuming the same were true of the 100,000 workers in *Chromium (200,000)*, then requiring each of 100,000 cost-bearing consumers to pay \$200 to eliminate a 1 in 100,000 risk of death to those workers would mean requiring each consumer to make an investment in the workers' safety that, knowing all and only what was knowable at the time the investment decision had to be made, no worker would have been willing to make in her *own* safety. When I complain that you

255. This assumes that *Trapped Miner* is similar to the typical workplace risk scenario in that both the cost-bearers and the person in danger benefit meaningfully from the activity that gave rise to the risk. In *Trapped Miner*, the risky activity is the operation of the town mine, which could be presumed to provide jobs or other benefits to town residents.

256. See discussion *supra* Part II.D.2.

treated me inequitably by failing to assume a burden for the sake of providing me with a benefit, it seems morally relevant whether *you* would have been willing to assume the burden on your own behalf. If, knowing all and only what was knowable at the time the decision had to be made, you would *not* have been willing to assume the burden on your own behalf, this would seem to be at least *prima facie* grounds for believing that I would not be treating you inequitably by refusing to assume the burden on your behalf.

Thus, to the family of the unlucky worker killed when the chromium risk tolerated by moderate regulation materializes, one could justify the risk imposition by saying, “Based on what was known at the time, your loved one herself would not have been willing to invest in her own safety what each cost-bearer would have had to invest in order to prevent her death. Therefore, you cannot reasonably expect each cost-bearer to have made that investment on your loved one’s behalf.” And this is precisely what could *not* be said to the family of the unlucky miner in *Trapped Miner*. Thus, although equity requires saving the trapped miner, it arguably requires moderate, rather than aggressive, regulation in *Chromium (200,000)*.

Another way of putting this point is as follows. Equity weighs against aggressive regulation in *Chromium (200,000)* because the position each worker is put in by virtue of being exposed to a moderately regulated mortality risk is precisely the position the worker would have put *herself* in had all the burdens and benefits associated with the relevant decisions—whether to be exposed to the risk at all and how much to spend to reduce the risk—fallen on her.²⁵⁷ The worker would have chosen to be exposed to the mortality risk associated with chromium because this was the necessary cost of receiving the compensating benefit of gainful employment.²⁵⁸ Further, for the reasons discussed above in connection with *Chromium (2)*, were both the burdens and benefits of risk-regulation to accrue to the worker, she would

257. See discussion *supra* Part III.D.2 (showing aggressive regulation is not favorable because the cost of the regulation is not justified by the risk).

258. Obviously, the worker would prefer employment with no mortality risk, but the assumption here is that the worker’s decision to accept employment involving chromium exposure is a rational one, both because the benefit the worker receives from such employment exceeds the loss she experiences by virtue of the associated mortality risk and because the *net* benefit she realizes from such employment exceeds the net benefit she would realize from any other available employment.

have chosen moderate regulation, rather than no regulation or aggressive regulation.²⁵⁹ Thus, for a worker to expect consumers to bear the additional \$20 million cost of aggressive regulation is, simply by virtue of the fact that the costs of risk-reducing regulation happen to fall on someone else, to expect to be placed in a better position than that in which the worker would have placed herself. To do so would be to unfairly exploit a morally arbitrary feature of workplace risks: that those subject to the risk happen not to bear the costs of regulations that reduce it.

Thus, opting for aggressive rather than moderate regulation in *Chromium (200,000)* would be inequitable to each of the 100,000 consumers for precisely the same reasons that it would be inequitable to C in *Chromium (2)*. The fact that the expected cost of opting for moderate over aggressive regulation is greater in *Chromium (200,000)*—one death, rather than a 1 in 100,000 risk of one death—may strengthen the *welfare maximization* case for aggressive regulation, but it does not seem to weaken the *equity-based* case against such regulation.²⁶⁰ This suggests that application of the individual risk principle—which embodies both the single person question and contractualism's minimax requirement—achieves equity not only in the stylized two-person case of *Chromium (2)*, but also in the more realistic sort of workplace risk scenario presented in *Chromium (200,000)*.²⁶¹ Where a workplace risk falls on a large number of risk-bearers and thereby stands to result in one or more statistical deaths, equity requires that each cost-bearer be made to invest in risk-reduction up to the point at which further investment would reduce each cost-bearer's well-being by more than it would increase the expected well-being of each risk-bearer.²⁶² To allow regulatory investment to stop short of that point would be inequitable to the risk-bearer. To expect regulatory investment beyond that point would be inequitable to the cost-bearer.

The above analysis establishes that, when it comes to workplace risks, equity provides the normative basis for using

259. See *supra* note 257 and accompanying text.

260. See *supra* notes 193 & 257 and accompanying text.

261. Of course, the same three caveats to the individual risk principle discussed above in connection with *Chromium (2)* apply here as well. See *supra* text immediately following note 224.

262. See Waisman, *supra* note 19, at 1266.

willingness-to-pay as a guide for risk regulation. Of course, one might still think, in a case like *Chromium (200,000)*, the norm of welfare maximization trumps the norm of equity, requiring regulatory investment beyond the level dictated by the individual risk principle. Even if this were so, the fundamental point made above would still stand: to the extent willingness-to-pay is used as a guide for the regulation of workplace mortality risks, the basis for doing so is the norm of equity, interpreted as ex ante contractualism.

If one uses willingness-to-pay data as a basis for workplace risk regulation, the data must be employed in the way the individual risk principle contemplates: to inform a comparison between the amount each risk-bearer would be willing to pay for a given reduction in their mortality risk and the amount each cost-bearer would actually have to pay to accomplish that reduction.²⁶³ This comparison relies directly on the individual willingness-to-pay data underlying the VSL (e.g., an average person would be willing to pay \$90 to avoid a 1 in 100,000 mortality risk), rather than on the collective willingness-to-pay calculation the VSL itself reflects (e.g., \$9 million is the amount that 100,000 people would collectively be willing to pay to avoid a 1 in 100,000 mortality risk).²⁶⁴ This suggests that, when it comes to workplace risks, CBA must be modified in an important way. The next section explains how.

3. Another Realistic Scenario: Chromium (1.1 million)

Like the decision to regulate in *Arsenic (100,000)*, the decision to regulate aggressively rather than moderately in *Chromium (200,000)* would not be cost-justified based on a VSL of \$9 million. Both decisions would involve imposing a \$20 million regulatory cost for the sake of realizing a benefit worth only \$9 million (saving one statistical life). When a regulation is determined not to be cost-justified by the lights of CBA employing a VSL that is calculated based on willingness-to-pay, what is the normative significance of that determination? The answer is: it depends on the moral context.

263. See *supra* notes 221-23 and accompanying text.

264. See *supra* Part I.A (describing how the VSL is calculated).

Where those exposed to a mortality risk bear all or substantially all of the costs of regulation (as is true in the case of the typical consumption risk), a determination that a regulation is not cost-justified provides good grounds for believing that enacting the regulation would mean violating the personal autonomy of each risk-bearer.²⁶⁵ To say a regulation is not cost-justified is to say the total cost exceeds the total amount risk-bearers would collectively be willing to pay for the risk reduction the regulation accomplishes.²⁶⁶ If that is true, and if regulatory costs are equally distributed among the same people who are exposed to the risk, enacting a regulation that is not cost-justified (such as the regulation in *Arsenic (100,000)*) is tantamount to forcing each risk-bearer to invest more in their own safety than they would rationally wish to, thereby violating each risk-bearer's autonomy.²⁶⁷ This was the fundamental point established in Part II.

However, where those exposed to the risk do *not* bear the costs of regulation (as is true in the case of the typical workplace risk), whether or not the regulation is cost-justified seems to have no independent normative significance. That is, where workplace risks are concerned, there does not seem to be a normative basis for basing regulatory decisions on a comparison between the total cost of a proposed regulation and the total amount risk-bearers would *collectively* be willing to pay for the risk reduction the regulation accomplishes. In the case of a workplace risk, a comparison of this kind is a poor indicator of whether a particular regulation would promote autonomy, equity, overall well-being, or any other norm.

To illustrate, consider a workplace risk scenario that is identical to *Chromium (200,000)*, except that, instead of being distributed among 100,000 consumers, regulatory costs are distributed among 1 million consumers. (Call this *Chromium (1.1 million)*). This reflects a situation that often occurs in the real world, where the number of people on which a workplace risk falls is just a small fraction of the number of the people among

265. See *supra* notes 114-19 and accompanying text.

266. See *supra* text immediately following note 12 (describing cost-justified and providing an example).

267. See *supra* Part II.D.

whom the costs of risk-reducing regulation are spread.²⁶⁸ The chart below illustrates the individual and aggregate burdens associated with each regulatory option. Notice that each consumer's burden is 10% the size of the corresponding burden in *Chromium (200,000)*.

Chromium (1.1 million)

Regulatory Action	Each Worker's Ex Ante Burden	Each Consumer's Burden	Ex Post Burden on 100,000 Workers	Total Burden on 1 million Consumers
No regulation	1 in 50,000 risk (WTP \geq \$180)	\$0	2 expected deaths	\$0
Moderate regulation	1 in 100,000 risk (WTP = \$90)	\$4	1 expected death	\$4 million
Aggressive regulation	No risk (WTP = \$0)	\$24	0 expected deaths	\$24 million

Choosing aggressive over moderate regulation is not cost-justified here for the same reasons it was not cost-justified in *Chromium (200,000)*: it would mean imposing a \$20 million incremental cost in order to obtain a \$9 million incremental benefit (avoiding one additional premature death). However, from the standpoint of equity, there is a significant difference between *Chromium (1.1 million)* and *Chromium (200,000)*. Where regulatory costs are distributed among 1 million rather than 100,000 consumers, a decision to regulate aggressively rather than moderately would mean imposing an additional cost of \$20 on each consumer rather than an additional cost of \$200.

268. See Waisman, *supra* note 19, at 1266, 1311, 1313.

Because the resulting reduction in each worker's risk of premature death (from 1 in 100,000 to zero) is one for which a worker would be willing to pay \$90, the individual risk principle would require aggressive regulation in *Chromium (1.1 million)*. From the standpoint of equity, each consumer can reasonably be expected to bear the burden of an additional \$20 cost so that each worker can avoid a mortality risk imposition that is equivalent to bearing a cost over four times as great (\$90). Thus, although the total cost of regulation is unchanged from *Chromium (200,000)*, equity here dictates a different regulatory result by virtue of how the cost is distributed.

This shows that, when evaluating the regulation of workplace risks from the standpoint of equity, what matters is an individual-level comparison between the amount each cost-bearer would have to pay to subsidize regulation and the amount each risk-bearer would be willing to pay for the risk reduction the regulation accomplishes. Because equity is a norm concerned with the burdens and benefits experienced by each affected individual,²⁶⁹ rather than burdens and benefits considered in the aggregate, the type of comparison made in standard, VSL-based CBA tells us nothing about whether a given regulation will or will not promote equity.²⁷⁰ The reason aggressive regulation would be inequitable in *Chromium (200,000)* was *not* that its \$20 million incremental cost exceeded its \$9 million incremental benefit, but rather that the \$200 cost burden experienced by each consumer exceeded the \$90 each worker would be willing to pay for the resulting risk reduction. When, in *Chromium (1.1 million)*, the same \$20 million cost is distributed among a group of consumers ten times as large, aggressive regulation becomes the equitable result because the cost burden to each consumer falls to \$20. The comparison captured by standard VSL-based CBA—one between a regulation's total cost and the total amount risk-bearers would collectively be willing to pay for the risk reduction the regulation accomplishes—is irrelevant from the standpoint of equity.

But is this sort of comparison relevant from the standpoint of some other norm? One might think that a comparison between a regulation's total cost and total monetized benefit would be

269. See *supra* note 22 and accompanying text.

270. See *supra* Part I.A (explaining the comparisons made in the VSL calculations).

relevant to the norm of overall well-being. After all, if a regulation's total cost exceeds its total monetized benefit, doesn't this show that overall well-being would be reduced were the regulation enacted? For reasons explained above in Part II.B, it does not. Standard CBA captures the welfare effects of the regulation's monetary cost and the regulation's ex ante risk reduction benefit reasonably well.²⁷¹ It fails, however, to persuasively capture the welfare effect of the regulation's ex post death avoidance benefit.²⁷² For this reason, whether or not a given regulation is cost-justified under VSL-based CBA is a poor indicator of whether it would increase or decrease overall well-being on net.

Thus, where the regulation of workplace risks is concerned, the dictates of standard, VSL-based CBA have little normative significance. As the above analysis of *Chromium (1.1 million)* illustrates, a workplace risk regulation that is not cost-justified might still be morally justifiable on the basis of equity or overall well-being.²⁷³ Although standard CBA does not reliably track *any* norm, an individual-level comparison between each cost-bearer's share of regulatory costs and each risk-bearer's risk-reduction benefit at least provides a reliable indication of what the norm of equity would dictate in a given case. This suggests that, where workplace risks are concerned, CBA should be modified to focus on the type of comparison contemplated by the individual risk principle: a comparison between the amount each cost-bearer would have to pay to subsidize a given regulation and the amount each risk-bearer would be willing to pay for the risk reduction the regulation delivers. When regulators consider how stringently to regulate a workplace mortality risk, this sort of comparison will indicate when further regulation would be inequitable. Individualized CBA reliably tracks the norm of equity. By contrast, when it comes to workplace risks, standard CBA is a decision procedure in search of a normative basis.

271. See *supra* Part II.B.

272. See *supra* Part II.B.

273. See *supra* text accompanying note 260.

CONCLUSION

In this article, I have argued for the importance of taking moral context into account in regulating industrial mortality risks. Making mortality risk regulation sensitive to moral context means distinguishing sharply between cases in which the costs of regulation are borne by the same people who are exposed to the risk (consumption risks) and cases in which such costs are borne by others (workplace risks). Where consumption risks are concerned, the norm of personal autonomy provides a basis for using the VSL as a guide for risk regulation. In this moral context, following the dictates of standard, VSL-based CBA ensures that each risk-bearer is made to invest as much, but no more, in their own safety as they would rationally wish to. Where workplace risks are concerned, the norm of equity provides a basis for using individual willingness-to-pay, though not the VSL, as a guide for risk regulation. In that moral context, standard VSL-based CBA has no clear normative basis and should be subordinate to the individualized form of CBA represented by the individual risk principle.

What consumption risks and workplace risks have in common is that in both cases risk-bearers benefit meaningfully from the industrial activity that creates the risk. This shapes the moral context in two important ways. First, it makes the risk imposition itself morally defensible on the grounds that the imposition works to the *ex ante* advantage of those imperiled. Second, in cases where risk-bearers do not bear the costs of regulation, it provides a moral basis for limiting regulatory costs in an equitable manner.

Unlike consumption and workplace risks, environmental risks often fall on persons who neither benefit meaningfully from the industrial activity giving rise to the risk nor bear any appreciable share of the costs of reducing the risk. For example, the subsistence fishermen exposed to a small risk of death from eating fish filled with toxic chemicals deposited into fishing streams by nearby paper and pulp mills may not benefit in any meaningful way from the activity of paper production.²⁷⁴ Two

274. See generally *National Emission Standards for Hazardous Air Pollutants for Source Category: Pulp and Paper Production*, 63 Fed. Reg. 18,504 (Apr. 15, 1998) (to be codified at 40 C.F.R. pt. 63, 261, 430).

basic questions arise in this moral context. First, under what circumstances if any is it morally permissible to impose a risk of death on one group of people for the sake of providing a benefit to a different group of people? Second, assuming there are circumstances in which it is permissible to imperil the lives of some in order to benefit others, is it morally defensible to use the VSL or willingness-to-pay as a guide for risk regulation in such circumstances?

In future work, I hope to take up these questions. Tentatively, while it is possible to conceive of circumstances in which it would be morally permissible to impose a risk of death on some for the sake of benefitting others, it is far from clear that these circumstances actually obtain in the case of most environmental risks. And even where such circumstances do hold, using a VSL or the general concept of willingness-to-pay to guide risk regulation would not seem to be defensible on the basis of overall well-being, personal autonomy, equity, or any other norm.