Risk, Death and Harm: The Normative Foundations of Risk Regulation

Matthew D. Adler[†]

I.	Conceptual Preliminaries: The Nature of Well-Being and Risk	1303
A.	Well-Being	1303
B.	Risk	1310
C.	Consequentialism, Risk, and the Ex Ante/Ex Post Distinction	1316
II.	Is Death a Harm?	1321
III.	Is the Risk of Death a Harm? Risks as Relative Frequencies	1340
A.	What Is the Risk of Death? A Frequentist Account	1341
B.	Statistical Reference Classes	1348
C.	Physical Reference Classes	1357
D.	Conventional Reference Classes	1362
E.	Hybrid Accounts	1366
IV.	Is the Risk of Death a Harm? Bayesian Theories of Risk	1369
A.	First-Person Actual Beliefs	1370
B.	Fear	1375
C.	Other Bayesian Accounts: Third-Person Beliefs and Hypothetical First-Person Beliefs	1385
V.	Implications	1388
A.	Cost-Benefit Analysis	1389
1.	CBA: Monetized Version	1395
2.	CBA: Nonmonetized Version	1402

[†] Professor of Law, University of Pennsylvania Law School. Many thanks to Brian Bix, Jim Chen, Sherry Colb, Paul Edelman, Kristin Madison, Stephen Perry, Eric Posner, Arti Rai, Julia Rudolph, Reed Shuldiner, and audiences at Emory University School of Law, George Washington University Law School, University of Pennsylvania Law School, Vanderbilt University Law School, University of Virginia School of Law, and College of William and Mary School of Law for their help. All errors are my own.

B. Risk-Risk Analysis	1410
C. Health and Safety Thresholds	1414
D. Distributive Justice.....	1423
E. Comparative Risk Assessment.....	1431
F. Tort and Criminal Law	1436

INTRODUCTION

Is death a harm? Is the risk of death a harm? These two questions are fundamental to risk regulation.¹ Whether risky activities ought to be regulated and, if so, how stringently, largely depends on whether the activities are *harmful*—whether they produce welfare setbacks for humans or other welfare subjects—either because the activities cause some humans or other welfare subjects to die, or because they impose the risk of death on welfare subjects, or both.² Risk regulation is expensive, in many different ways. It limits the freedom of action of regulated parties; it imposes compliance costs on these parties, costs that are spread to consumers, workers, shareholders, and others engaged in commercial interactions with the regulatees; it consumes the money and effort of administrators, legislators, judges, and citizens.³ Surely the Environmental Protection Agency (EPA), the Occupational Safety and Health Administration (OSHA), the Food and Drug Administration (FDA), the National Highway Traffic Safety Administration (NHTSA), the Consumer Product Safety Commission (CPSC), and the host of other federal and state agencies that engage in risk regulation would not be justified in doing so absent some substantial nexus between risk, death, and well-being.⁴ Yet legal scholars and others who write about risk regulation have largely ignored these two questions.⁵

1. Throughout this Article, I use the term “risk” to mean a fatality risk.

2. See, e.g., Matthew D. Adler & Eric A. Posner, *Rethinking Cost-Benefit Analysis*, 109 YALE L.J. 165, 204-16 (1999) (arguing for the central role of overall well-being in determining normatively appropriate regulatory choice).

3. See Robert W. Hahn & John A. Hird, *The Costs and Benefits of Regulation: Review and Synthesis*, 8 YALE J. ON REG. 233, 271-78 (1991) (providing quantitative estimates for the costs and benefits of social regulation).

4. For an overview of federal risk regulation, see STEPHEN BREYER, *BREAKING THE VICIOUS CIRCLE: TOWARD EFFECTIVE RISK REGULATION* (1993).

5. Major scholarly works on risk regulation that are general in scope and normative in focus include BREYER, *supra* note 4; CARL F. CRANOR, *REGULATING TOXIC SUBSTANCES: A PHILOSOPHY OF SCIENCE AND THE LAW* (1993); RISK VERSUS RISK: TRADEOFFS IN PROTECTING HEALTH AND THE

Is death a harm? How could it not be? Uncontroversially, the death of some person can be a harm (welfare setback) for her friends, family, and others who survive her. The death may *pain* them; they may *disprefer* it; and the death of a friend, or family member, is plausibly counted as a welfare loss in a more *objective* sense, specifically as the ending of a close relationship. The standard constituents of well-being—pleasure and pain, preference, objective goods—are all present here.⁶ But is death a harm for the person who dies? That is the question I wish to focus on here and throughout this Article. Common sense urges an affirmative answer. If death is not a harm for the person who dies, why do most of us fear our own deaths quite intensely and take substantial steps to avoid fatal accidents, mortal illness, or dangerous substances? Yet common sense may well mislead here. Death itself (by contrast with the anticipation of death or the process of dying) is not painful for the person in that state. More generally, death itself (by contrast with these other states) is not *experienced* by the person in that state.⁷ Events or states can be harmful without being painful—if, for example, the events or states are dispreferred or are objectively bad. But can events or states be harmful to some person without being experienced by her? It is a widely shared intuition that posthumous events do not harm the deceased person: Slandering someone, ruining her business, and destroying her homestead after her death may wrong her but, intuitively, these are harmless wrongs.⁸ If this intuition is

ENVIRONMENT (John D. Graham & Jonathan Baert Wiener eds., 1995) [hereinafter RISK VERSUS RISK]; K.S. SHRADER-FRECHETTE, RISK AND RATIONALITY: PHILOSOPHICAL FOUNDATIONS FOR POPULIST REFORMS (1991); W. KIP VISCUSI, FATAL TRADEOFFS: PUBLIC AND PRIVATE RESPONSIBILITIES FOR RISK (1992); John S. Applegate, *The Perils of Unreasonable Risk: Information, Regulatory Policy, and Toxic Substances Control*, 91 COLUM. L. REV. 261 (1991); Clayton P. Gillette & James E. Krier, *Risk, Courts, and Agencies*, 138 U. PA. L. REV. 1027 (1990); Howard Latin, *Good Science, Bad Regulation, and Toxic Risk Assessment*, 5 YALE J. ON REG. 89 (1988); Christopher H. Schroeder, *Rights Against Risks*, 86 COLUM. L. REV. 495 (1986). None of these works give sustained attention to the questions "Is death a harm?" and "Is risk a harm?" Cass Sunstein's very important recent book on risk regulation (which was published after this Article was drafted) addresses numerous important questions, but not these. See CASS R. SUNSTEIN, RISK AND REASON: SAFETY, LAW, AND THE ENVIRONMENT (2002).

6. See *infra* text accompanying notes 36-60 (delineating elements of welfare).

7. See *infra* Part II (articulating and analyzing the arguments for the claim that death is not harmful).

8. See L.W. SUMNER, WELFARE, HAPPINESS, AND ETHICS 126-27 (1996)

accepted, it seemingly follows that death itself is also no harm. Posthumous events are powerless to affect the well-being of the deceased because they are experientially remote from her. They cannot alter what she experiences. Since death, too, is experientially remote in just this way, death is not a harm—or so a plausible philosophical argument goes.

This argument was first advanced by the Greek philosopher Epicurus.

Accustom thyself to believe that death is nothing to us, for good and evil imply sentience, and death is the privation of all sentience, . . . Death, therefore, the most awful of evils, is nothing to us, seeing that, when we are, death is not come, and when death is come, we are not. It is nothing, then, either to the living or to the dead, for with the living it is not and the dead exist no longer.⁹

Thomas Nagel, in his 1979 article crisply entitled *Death*, criticized Epicurus's argument but did not decisively refute it.¹⁰ Rather, *Death* triggered renewed philosophical interest in the Epicurean challenge to the common-sense view, and there is now a substantial philosophical literature devoted to the question, "Is death a harm?"¹¹ This literature includes spirited and sophisticated defenses of Epicurus's view by a number of academic philosophers.¹² Yet legal scholars, policy analysts, welfare economists, and others who have written normative scholarship about risk regulation have unwaveringly assumed that death itself is a serious welfare setback.¹³ The philosophical literature just described has been completely ignored by these other academic disciplines.

Is the risk of death a harm? Does imposing the risk of death on some person constitute a welfare setback for her, separate from whatever setback may inhere in death itself? This issue has received even less scholarly attention than the issue of death's harmfulness. Much philosophical work has been done explicating the nature of risk and probability.¹⁴ It

(discussing welfare impact of posthumous events).

9. Stephen E. Rosenbaum, *How to Be Dead and Not Care: A Defense of Epicurus*, in *THE METAPHYSICS OF DEATH* 119, 121 (John Martin Fischer ed., 1993) (alteration in original) (quoting Epicurus's "Letter to Menoeceus").

10. Thomas Nagel, *Death*, in *MORTAL QUESTIONS* 1 (1979), reprinted in *THE METAPHYSICS OF DEATH*, *supra* note 9, at 61.

11. Many of the important papers are collected in *THE METAPHYSICS OF DEATH*, *supra* note 9.

12. See sources cited *infra* note 110.

13. See, e.g., *supra* note 5 (citing major normative works on risk regulation by legal scholars and others).

14. Good overviews of this philosophical work include L. JONATHAN

turns out that risk is a many-headed monster. There are at least four standard accounts of what risk consists in (the frequentist account, the Bayesian account, the classical account, and the logical account) and numerous nonstandard accounts as well. The frequentist sees the risk of some event (for example, a death) as the frequency with which that type of event occurs in some large class of events (for example, all cases in which persons are exposed to a given toxic product).¹⁵ The Bayesian sees the risk of some event as someone's degree of belief that the event will occur; this account is seemingly more subjective than the frequentist account, because different persons can have different degrees of belief for the same event, while the relative frequency of a fatal event in a given reference class is a fact determined by the laws of physics, chemistry, and biology.¹⁶ Classicists and logicians, unlike Bayesians and frequentists, believe that risk is a conceptual rather than empirical matter. What the risk of death is, in a given case, depends not upon what people happen to believe (as per the Bayesian account), or what the empirically contingent relative frequencies happen to be (as per the frequentist account), but on the conceptual scheme that we use to understand the world and its future possibilities. Specifically, the risk of a given death is the ratio of the number of possible futures in which that death occurs to the number of all possible futures (this is the classical view), or some variation on that ratio (this is the logical view).¹⁷

Much philosophical work, too, has been devoted to the nature of well-being.¹⁸ Here, as with the nature of risk, there

COHEN, AN INTRODUCTION TO THE PHILOSOPHY OF INDUCTION AND PROBABILITY (1989); DONALD GILLIES, PHILOSOPHICAL THEORIES OF PROBABILITY (2000); HENRY E. KYBURG, JR., PROBABILITY AND INDUCTIVE LOGIC (1970); STORRS MCCALL, A MODEL OF THE UNIVERSE: SPACE-TIME, PROBABILITY, AND DECISION 141-47 (1994); ROY WEATHERFORD, PHILOSOPHICAL FOUNDATIONS OF PROBABILITY THEORY (1982); and Colin Howson, *Theories of Probability*, 46 BRIT. J. FOR PHIL. SCI. 1 (1995).

15. See *infra* text accompanying notes 77-82 (discussing the frequentist account of risk).

16. See *infra* text accompanying notes 73-76 (discussing the Bayesian account of risk).

17. See COHEN, *supra* note 14, at 43-47, 74-80 (discussing the classical and the logical accounts of risk); GILLIES, *supra* note 14, at 14-49 (same); KYBURG, *supra* note 14, at 29-39, 54-67 (same); WEATHERFORD, *supra* note 14, at 18-143 (same). The statement in the text is rough and, as discussed by the surveys just cited, better describes one well-known variant of the logical account (Rudolf Carnap's) than another (John Maynard Keynes's).

18. Good overviews include JAMES GRIFFIN, WELL-BEING 7-72 (1986);

are multiple standard accounts. I have already alluded to them: the hedonic account, which sees pleasure and pain as the source of welfare gains or losses; the preferentialist account, which looks instead to what people prefer or disprefer; and the objective-good account, which identifies certain "valuable" or "desirable" activities, occurrences, or states as welfare sources, independent of whether the activities, occurrences, or states are pleasurable or preferred.¹⁹ Each of the standard accounts has flaws, ably presented by the philosophical critics of that account. But this flurry of argument and counterargument has not led to some grand, scholarly synthesis. No single, consensus view of well-being has yet emerged—just as no single, consensus view of risk has. Instead, philosophical hedonists continue to defend hedonism, with sophisticated modifications designed to reduce the critics' damage, and the same is true of preferentialism and objectivism.²⁰ For example, the sophisticated hedonist construes "pleasure" as a desired or desirable mental state, rather than (more simplistically) a positive sensation.²¹ The sophisticated preferentialist looks to idealized preferences—to what someone would prefer under full information, say—in addition to, or in lieu of, actual preferences.²² And the sophisticated objectivist counts pleasure and preference satisfaction as objective goods in their own right—as sources, if not the exclusive sources, of well-being.²³

Are you confused? You should be. The question "Is risk a harm?" is intricate and tricky. Answering that question in a thoughtful and rigorous way means combining the insights of

DEREK PARFIT, REASONS AND PERSONS 493-502 (1984); T.M. SCANLON, WHAT WE OWE TO EACH OTHER 108-43 (1998); SUMNER, *supra* note 8, at 45-137 (1996); and Mozaffar Qizilbash, *The Concept of Well-Being*, 14 ECON. & PHIL. 51 (1998).

19. See *infra* text accompanying notes 37-39, 42-43, 48-56 (describing the hedonic, preferentialist, and objective-good accounts of well-being).

20. See *infra* text accompanying notes 37-60 (summarizing the debate about the standard accounts of well-being).

21. See *infra* text accompanying note 40.

22. See *infra* text accompanying note 45.

23. Martha Nussbaum's list of objective goods includes the use of the "senses, imagination, and thought," while James Griffin's includes enjoyment. See *infra* text accompanying notes 50-53. Both make clear that these goods include pleasure and/or the absence of pain. See MARTHA C. NUSSBAUM, WOMEN AND HUMAN DEVELOPMENT: THE CAPABILITIES APPROACH 78-79 (2000) ("Senses, [i]magination, and [t]hought" include "[b]eing able to have pleasurable experiences, and to avoid non-necessary pain."); JAMES GRIFFIN, VALUE JUDGEMENT: IMPROVING OUR ETHICAL BELIEFS 30 (1996) (stating, under the rubric of "[e]njoyment," that "[w]e value pleasures").

the two philosophical literatures just summarized. Even if the classical and logical accounts of risk are rejected, as I believe they should be,²⁴ we are still left with the multiple possible permutations between two quite different accounts of risk—Bayesian and frequentist—and a number of divergent views of well-being. Hedonism (as an account of well-being) plus Bayesianism (as an account of risk) may produce a different view of risk's harmfulness than hedonism plus frequentism, or preferentialism plus Bayesianism, or an objective-good view of welfare plus a frequentist view of risk. The task seems both demanding and yet also very inviting. Surely there are nonobvious and philosophically significant truths that will emerge by permuting accounts of risk with accounts of well-being. Surely, too, the question "Is risk a harm?" has much importance for risk regulation, as I will elaborate in a moment. Yet this question is largely overlooked—not only by lawyers, economists, and policy analysts, but also by philosophers. Neither the philosophical literatures on well-being or risk nor the legal, economic, and policy-analytic literatures on risk regulation have considered the harmfulness of risk in a sustained way.²⁵

I have already stated, emphatically, that the two questions mooted here—"Is death a harm?" and "Is the risk of death of harm?"—have much significance for risk regulation. Let me elaborate. Much modern legislation is targeted at fatal and probabilistically fatal activities, products, and substances.²⁶ If neither risk nor death is a harm, the justifiability of these risk-regulation statutes is open to serious question. Presumably they should just be repealed. Perhaps the supporter of regulation could respond that some of these statutes aim to diminish not only death and fatality risks, but also ecological harm or bodily injury. Ecological harm is indeed a distinct type of welfare setback²⁷ but one that generally has limited

24. See *infra* text accompanying notes 83-84.

25. There are a few substantial articles on this issue. See Stephen R. Perry, *Risk, Harm, and Responsibility*, in *PHILOSOPHICAL FOUNDATIONS OF TORT LAW* 321 (David G. Owen ed., 1995); Claire Finkelstein, *Is Risk a Harm?*, 151 U. PA. L. REV. (forthcoming 2003); Stephen Fogdall, *Risks as Harms* (unpublished manuscript, on file with author).

26. See, e.g., BREYER, *supra* note 4, at 8 ("[R]egulation designed to screen out risky substances . . . is embodied in many different regulatory programs—indeed, in at least twenty-six different statutes administered by at least eight different agencies.").

27. See, e.g., A. MYRICK FREEMAN III, *THE MEASUREMENT OF*

relevance for risk-regulation statutes administered by nonenvironmental agencies, such as OSHA, the NHTSA, the FDA, or the CPSC at the federal level. Further, "[t]he primary benefit of many important environmental statutes, as determined by the dollar value assigned by cost-benefit analysis, is the human lives that are saved."²⁸ Bodily injury absent death is clearly welfare reducing,²⁹ but what is far from clear is whether the large economic costs imposed by our current regime of workplace safety laws, motor vehicle safety laws, food and drug laws, consumer product safety laws, and the like would be warranted by a reduction in injury rates, taken alone.³⁰ A similar response could be made to the suggestion that these statutes reduce the third-party costs of death: the pain and misery death causes for the deceased's friends and family. At a minimum it seems quite plausible that the scope and stringency of the current risk-regulation regime—now thought to be warranted primarily by the reduction in first-party costs that result from fatalities and fatality risks—would change dramatically.³¹

More generally, both the scope and structure of statutes regulating fatal and probabilistically fatal activities, and the proper exercise of administrative discretion pursuant to these statutes, will depend crucially on which box in the matrix below turns out to be correct.

ENVIRONMENTAL AND RESOURCE VALUES: THEORY AND METHODS (1993) (discussing valuation of both ecological harm and death).

28. Richard L. Revesz, *Environmental Regulation, Cost-Benefit Analysis, and the Discounting of Human Lives*, 99 COLUM. L. REV. 941, 943 (1999).

29. Cf. VALUING HEALTH FOR POLICY (George Tolley et al. eds., 1994) (discussing valuation of setbacks to bodily integrity resulting both from disease and from death).

30. Cf. BREYER, *supra* note 4, at 11-19 (assuming that death is a harm, but suggesting that substantial changes in regulatory programs would occur if we properly valued death).

31. Cf. Robert W. Hahn, *Regulatory Reform: What Do the Government's Numbers Tell Us?*, in RISKS, COSTS, AND LIVES SAVED: GETTING BETTER RESULTS FROM REGULATION 208, 219 (Robert W. Hahn ed., 1996) [hereinafter RISKS, COSTS, AND LIVES SAVED] (reviewing environmental, health, and safety rules issued by OSHA, the EPA, the CPSC, the NHTSA, and the Mine Safety and Health Administration, and concluding that "[a]bout 60 percent of the total benefits results from reductions in the risk of death, disease, and injury").

Is the Risk of Death a Harm?

		<i>No</i>	<i>Yes</i>
Is Death a Harm?	<i>No</i>	Null View	Risk View (Frequentist and/or Bayesian)
	<i>Yes</i>	Death View	Hybrid View (Frequentist and/or Bayesian)

Consider the justifiability of regulatory statutes, and of the directives issued by regulatory agencies pursuant to these statutes, with respect to the moral criterion of *overall well-being*.³² The Null View denies that either death itself, or the risk of death, is a first-party cost for purposes of determining overall well-being. Various harms and benefits separate from fatalities (including ecological harms and bodily injury harms), plus third-party harms resulting from fatalities, are relevant, but the first-party effects of death or the risk of death have zero relevance for the welfare calculus. The Death View changes this picture by giving weight to the first-party effects of death. Since the first-party cost of a single death is now appraised by regulators as somewhere in the vicinity of \$6 million,³³ the practical significance of moving from the Null View to the Death View is huge. The Risk View omits the \$6 million per death first-party cost from the calculus of overall well-being, but adds in a (presumably much smaller) cost for each person at risk, in the appropriate sense, of dying. Finally, the Hybrid View stipulates *both* that death is a harm *and* that risk is an independent harm. Any regulatory option that involves fatalities and fatality risks is attributed a \$6 million cost for each fatality, plus a separate cost for each person at risk (in the appropriate sense) of dying.³⁴

32. See Adler & Posner, *supra* note 2, at 209-16, 225-38 (defending the moral relevance of overall well-being, and conceptualizing cost-benefit analysis as a good if not perfect proxy for overall well-being).

33. See Frank Ackerman & Lisa Heinzerling, *If It Exists, It's Getting Bigger: Revising the Value of a Statistical Life* 1-2 (Global Dev. & Env't Inst., Working Paper No. 01-06, Oct. 2001), <http://www.ecologicaleconomics.org/documents/valueoflife.pdf>.

34. Strictly, this is not correct, since the "value of statistical life" (VOSL) method used to calculate the current \$6 million valuation of death is inconsistent with the "Bayesian" type of Hybrid View I argue for in this

In this Article, I defend a Hybrid View of risk regulation. More precisely, I defend a *Hybrid (Bayesian) View* of risk regulation. First, I argue against Epicurus and the modern philosophical Epicureans and contend that death is indeed a welfare setback for the person who dies. Second, I consider the two main viable accounts of risk, frequentist and Bayesian, and argue that risk in the Bayesian sense is indeed a harm for the person at risk, but that risk in the frequentist sense is not. Death and Bayesian risk should be counted as real costs and as separate costs in the regulatory calculus of overall well-being. By contrast, a high level of frequentist risk is not a welfare cost and is irrelevant to policy formulation insofar as policies are shaped by the aggregate-welfare criterion. More generally, death and Bayesian risk (but not frequentist risk) should be counted as separate types of welfare setbacks with respect to every moral criterion that incorporates welfare considerations. For example, if regulatory agencies are properly attentive to the "distributive impacts" of their decisions³⁵—to the effects of regulation on the distribution of well-being across the population—both death and the Bayesian risk of death should figure separately in the distributive analysis. Someone who dies, and was previously at Bayesian risk of dying, fares worse (*ceteris paribus*) than someone who was put at Bayesian risk but survives, and this survivor in turn fares worse (*ceteris paribus*) than someone who neither dies nor was put at Bayesian risk of dying. However, the fact that some person or group is exposed to a high relative frequency of premature death has no bearing on distributive justice.

This Article has five parts. Part I clarifies the concepts of risk and well-being. Part II considers whether death is a harm. Parts III and IV consider whether risk is a harm, in first the frequentist and then the Bayesian sense. The upshot of the analysis set forth in Parts II, III, and IV is the Hybrid (Bayesian) View of risk just described: Death is a harm, risk in the Bayesian sense is a harm, but risk in the frequentist sense is not. Finally, Part V surveys the large legal implications of this account. As I demonstrate in Part V, the issues mooted

Article. See *infra* text accompanying notes 302-05 (criticizing the VOSL method). Still, the \$6 million figure may be approximately correct, see *id.*, and in any event nothing in my Hybrid (Bayesian) View suggests that the correct monetary valuation of death is insubstantial.

35. See *infra* Part V.D. (discussing the role of distributive considerations within risk regulation).

here have very substantial implications for a range of regulatory institutions and practices: cost-benefit analysis, both monetized and nonmonetized; risk-risk analysis; the proper interpretation of the safety "thresholds" that statutes governing risk regulation commonly create; distributive assessment pursuant to "environmental justice" mandates; comparative risk assessment; and tort and criminal law.

I. CONCEPTUAL PRELIMINARIES: THE NATURE OF WELL-BEING AND RISK

A. WELL-BEING

What is well-being or, synonymously, welfare? Traditionally, philosophers have offered three kinds of answers to this question: a *hedonic* answer, a *preferentialist* answer, and an *objectivist* answer.³⁶

The *hedonic* account of welfare, in the simple version of this account exemplified by the work of Jeremy Bentham,³⁷ states that welfare consists in pleasure and the avoidance of pain, with "pleasure" understood as a positive feeling or sensation experienced by a conscious subject and about whose existence and sensational qualities the subject is infallible. "Pain" is understood similarly, except that the feeling or sensation is negative.³⁸ Simple hedonism is vulnerable to the following type of counterexample: "At the very end of his life, Freud, ill and in pain, refused drugs except aspirin. 'I prefer,' he said, 'to think in torment than not to be able to think clearly.'"³⁹ Freud seems to have benefited, all things considered, from a mental life that was, all things considered, more painful in the Benthamite sense—more painful in how it felt, in the sensations that it contained.

The Freud example suggests a more sophisticated brand of hedonism: define welfare-enhancing mental states as more valuable (or as desired) mental states rather than as more pleasurable mental states. Indeed, John Stuart Mill and

36. For good overviews of the philosophical literature on well-being, see sources cited *supra* note 18.

37. See SUMNER, *supra* note 8, at 87-89 (describing Bentham's version of hedonism).

38. See *id.* at 83-92.

39. GRIFFIN, *supra* note 18, at 8 (citing ERNEST JONES, *THE LIFE AND WORK OF SIGMUND FREUD*, 655-56); see also Jeffrey Goldsworthy, *Well-Being and Value*, 4 UTILITAS 1, 6-11 (1992) (further discussing Griffin's example).

another famous hedonist, Henry Sidgwick, modified Benthamite hedonism in just this way.⁴⁰ But now a different challenge to hedonism arises: Welfare cannot be reduced to mental states, however defined. Hedonists, simple or sophisticated, are committed to the thesis that welfare *supervenes* on mental states: Two outcomes in which a given person has the same mental states must be the same for his welfare. Here is one plausible counterexample to the supervenience thesis: If Joan prefers that her husband remain faithful to her, then she is worse off if he betrays her, even if the treachery never comes to her attention and her beliefs, feelings, and other mental states are no different from what they would have been had her husband remained faithful.⁴¹

The *preferentialist* account of welfare allows nonexperiential features of outcomes to come into play.⁴² In its simplest version, the preferentialist account says that a person is benefitted by one outcome, relative to a second, if she prefers the first to the second.⁴³ Note that a person can prefer one outcome to another even though her mental states are just the same in both outcomes. For example, Joan can prefer that her husband remain faithful and that she (correctly) believe him to be so, as against her husband betraying her and her (incorrectly) believing him to be faithful. Simple preferentialism, unlike hedonism, cannot be criticized for tying well-being too tightly to our experiences. But it can be criticized on other grounds.⁴⁴ First, a person can prefer an

40. See SUMNER, *supra* note 8, at 90-91 (describing Mill's and Sidgwick's versions of hedonism).

41. For a famous critique of the supervenience thesis, see ROBERT NOZICK, *ANARCHY, STATE, AND UTOPIA* 42-45 (1974).

42. For general discussions of preference-based or, equivalently, desire-based theories of well-being, see, for example, GRIFFIN, *supra* note 18, at 10-39; PARFIT, *supra* note 18, at 494-99; SCANLON, *supra* note 18, at 113-23; SUMNER, *supra* note 8, at 113-37; and Qizilbash, *supra* note 18, at 58-63. On the nature of a "preference," see, for example, S.L. HURLEY, *NATURAL REASONS* 55-83 (1989); Richard J. Arneson, *Liberalism, Distributive Subjectivism, and Equal Opportunity for Welfare*, 19 *PHIL. & PUB. AFF.* 158, 161-64 (1990); Arthur Ripstein, *Preference*, in *VALUE, WELFARE, AND MORALITY* 93, 93-111 (R.G. Frey & Christopher W. Morris eds., 1993); and the papers in *PREFERENCES* (Christoph Fehige & Ulla Wessels eds., 1998).

43. See, e.g., Mark C. Murphy, *The Simple Desire-Fulfillment Theory*, 33 *NOUS* 247, 269 (1999) (providing a qualified defense of the "simple desire-fulfillment theory" of well-being).

44. See, e.g., DANIEL M. HAUSMAN & MICHAEL S. MCPHERSON, *ECONOMIC ANALYSIS AND MORAL PHILOSOPHY* 71-83 (1996); SCANLON, *supra* note 18, at 113-23; SUMNER, *supra* note 8, at 113-37; Tyler Cowen, *The Scope and Limits*

outcome that is, intuitively, bad—not just morally bad, but bad for her, or at least no better. Assume that I am a sadist and prefer to torture Phil. The world in which I am allowed to satisfy this preference is, of course, worse for Phil, but it also seems worse or at least no better for me. Second, a person can prefer an outcome that is, intuitively, unrelated to her own life. If I have a conversation with a stranger on a train, learn of various projects in which he is engaged, develop a mild preference that the projects succeed, and never see the stranger again, then the projects' success does not benefit me even though this is the outcome I prefer. Third, simple preferentialism seems to "overcorrect" the error in hedonism. While hedonism binds welfare too tightly to mental states, preferentialism arguably makes the link too loose. Preferentialism, in its simple version, does not limit the features of outcomes that can ground preferences and thereby welfare differentials. Thus it allows that an outcome can benefit some person, relative to some alternative, even though the only difference between the outcomes is something quite remote from the person's experience—at the limit, something which it would be impossible for that person to experience, as in the case of dispreferred posthumous events, which according to simple preferentialism are welfare setbacks for the deceased.

The first of these criticisms might be met by idealizing preferences, in some way. Indeed, many preferentialists now stipulate that the preferences which ground welfare must be "fully informed."⁴⁵ Satisfying my desire to torture Phil does not benefit me because that desire is nonideal. If I were to reflect, with full information and due deliberation, on what torturing Phil involved for him and for me, I would not desire it. However, the move to "fully informed" preferences does nothing to address the second and third just-listed criticisms,⁴⁶ and it arguably fails to answer the first criticism as well. The *reason* that torturing Phil is bad or at least no better for me is that

of *Preference Sovereignty*, 9 *ECON. & PHIL.* 253 (1993).

45. See Don Loeb, *Full-Information Theories of Individual Good*, 21 *SOC. THEORY & PRAC.* 1 (1995) (surveying full-information accounts of well-being and presenting criticisms); Connie S. Rosati, *Persons, Perspectives, and Full Information Accounts of the Good*, 105 *ETHICS* 296, 296-99 (1995) (same); David Sobel, *Full Information Accounts of Well-Being*, 104 *ETHICS* 784, 790-96 (1994) (same).

46. See, e.g., SCANLON, *supra* note 18, at 115 ("The objects of a person's informed desires are likely to include many things that are not related to the quality of the desirer's own life, intuitively understood.").

this activity lacks value; it is undesirable. We have a sense of what a good human life consists in, reflected in (or constituted by) the evaluative concepts with which we perceive and describe human activity: concepts like "accomplishment," "friendship," "leadership," "participation," "intimacy," "love," and so on. With full information, I would disprefer torturing Phil because I would bring these evaluative concepts into play and come to recognize that sadistic pleasure makes no positive contribution to my life. To say that some outcome improves my welfare if I would prefer it under full information gets things backwards. Rather, it improves my welfare if it is desirable or preferable—*worthy* of desire or preference.⁴⁷

We have now moved to the third type of welfare account typically advanced within the philosophical literature: an objective-good account. Objectivism about welfare goes back to Aristotle⁴⁸ and, more recently, has been defended by John Finnis, Martha Nussbaum, George Sher, and other distinguished philosophers.⁴⁹ Objectivists are typically pluralists; they typically offer a list of "values" or "goods" that represent different dimensions of human welfare, different ways in which a human life can go well or badly. For example, Finnis claims that these goods are life itself, knowledge, play, aesthetic experience, sociability, practical reasonableness, and religion.⁵⁰ Nussbaum's list includes: life, bodily health, bodily integrity, the uses of the "senses, imagination and thought," the emotions, practical reason, affiliation, interaction with other species, play, and control over one's environment.⁵¹ Derek Parfit, describing (without endorsing) objectivism, writes that "[t]he good things might include moral goodness, rational

47. See, e.g., *id.* at 119. Scanlon states,

It may be true that something contributes to one's well-being only if one has reason to desire it. But even when this is so, what makes this thing good will not be the fact that it would satisfy that hypothetical desire but rather those considerations, whatever they may be, that provide reasons for desiring it.

Id.

48. See, e.g., SUMNER, *supra* note 8, at 69-72 (summarizing Aristotle's view of well-being).

49. See JOHN FINNIS, *NATURAL LAW AND NATURAL RIGHTS* 59-99 (1980); THOMAS HURKA, *PERFECTIONISM* 9-143 (1993); NUSSBAUM, *supra* note 23, at 34-110; GEORGE SHER, *BEYOND NEUTRALITY: PERFECTIONISM AND POLITICS* 199-244 (1997). James Griffin's most recent work on well-being has a decidedly objectivist flavor. See GRIFFIN, *supra* note 23, at 19-36.

50. See FINNIS, *supra* note 49, at 85-90.

51. See NUSSBAUM, *supra* note 23, at 78-80.

activity, the development of one's abilities, having children and being a good parent, knowledge, and the awareness of true beauty."⁵² James Griffin lists accomplishment, autonomy, understanding, enjoyment, and deep personal relations.⁵³ In short, there is substantial disagreement within the objectivist camp as to the *content* of the canonical list of goods or values, and there is also substantial disagreement as to the *basis* for that list. Some objectivists argue that human values or goods are grounded in the human essence—in the properties, such as a capacity for rational belief and action, that a being necessarily possesses if she is human.⁵⁴ A different proposal is that "what unifies the diverse elements of a good life is their connection(s) to near-universal, near-unavoidable goals."⁵⁵ This, in turn, differs from an objectivist view that places more emphasis on culture and language: humans, living together, develop shared understandings of what valuable and worthless lives consist in.⁵⁶ Despite these differences, objectivists all concur in the claim that welfare depends upon some set of *objective* goods or values: objective in that what is truly good or valuable for a given person can differ from what she believes to be good or valuable, or what she prefers.

Objectivism, in theory, can remedy the various deficits in preferentialism mentioned above. It will circumscribe the features of outcomes that count as good or bad, welfare-enhancing or -reducing, for a given person. Features that are undesirable in any life (sadistic pleasures), or unconnected to a given person's life (the stranger on the train), or too experientially remote (posthumous events), will be ruled out. And objectivism, unlike hedonism, is not committed to the view that experience is all that matters for our welfare.

Notwithstanding these attractions, objectivism has hardly defeated preferentialism or hedonism within the philosophical literature on well-being. Why not? To begin, it seems implausible that a person can be benefited by some outcome which she never prefers or comes to prefer.⁵⁷ Opera might be better than sitcoms, but if I like and continue to like sitcoms, my watching opera instead of sitcoms does not benefit me. A

52. PARFIT, *supra* note 18, at 499.

53. GRIFFIN, *supra* note 23, at 29-30.

54. See HURKA, *supra* note 49, at 9-51.

55. SHER, *supra* note 49, at 229.

56. See Qizilbash, *supra* note 18, at 63 (discussing Griffin's objectivism).

57. See Adler & Posner, *supra* note 2, at 200-02.

possible response is to develop a mixed theory of well-being which states: an outcome benefits some person if and only if it is both (1) good for him and (2) preferred by him.⁵⁸ But this mixed theory will be unpersuasive to some preferentialists, who will insist that "goodness" or "value" is irrelevant to welfare and should be abandoned in favor of the concept of "fully informed" (or otherwise idealized) preferences.

As for the continued appeal of hedonism, some philosophers continue to find compelling the thought that welfare *does* supervene on our experience.⁵⁹ Consider the concepts of happiness, welfare, and value. Happiness, most would agree, supervenes on experience; value, most would agree, does not.⁶⁰ Is welfare nearly the same as happiness, or is it nearly the same as value? Because answers to this question differ, it is not surprising that reactions to the supervenience thesis also do.

In short, the philosophical literature has failed to coalesce around a particular account of well-being, and indeed there is persisting and deep-rooted scholarly disagreement about which general type of account (hedonic, preferentialist, or objectivist) is correct. Given this disagreement, how should we analyze the two issues at the heart of risk regulation, namely, "Is death a harm?" and "Is risk a harm?" Rather than grounding the analysis on a particular welfare view, which would render the analysis unpersuasive to those who reject the view, or considering the two issues in light of all possible welfare views—an impossibly large task—my approach will be as follows. I will consider how the main plausible elements of a welfare account, those that have received substantial and continuing support within the philosophical literature, bear upon the harmfulness of risk and death. I take the "main plausible elements of a welfare account" to be the following:

58. See Matthew D. Adler, *Beyond Efficiency and Procedure: A Welfarist Theory of Regulation*, 28 FLA. ST. U. L. REV. 241, 265 (2000).

59. For a recent defense of experientialism about welfare, see Mark Bernstein, *Well-Being*, 35 AM. PHIL. Q. 39 (1998). See also Goldsworthy, *supra* note 39, at 3, 6-20 (presenting a qualified defense, which does not choose conclusively between hedonism and the "non-cognitivist" position that well-being is not a factual matter); Matthew Silverstein, *In Defense of Happiness: A Response to the Experience Machine*, 26 SOC. THEORY & PRAC. 279 (2000) (contending that Nozick's experience-machine example fails to refute hedonism about well-being); *id.* at 281 n.7 (citing scholarly defenses of hedonism).

60. On the concept of happiness, see SUMNER, *supra* note 8, at 140-47.

Experience, Preference, Value, and Integration.

Experience: As we have seen, hedonists insist that welfare supervenes on experience: A given person cannot fare better in one outcome than another unless his experience or mental states are different in the two outcomes. Experience might be incorporated into welfare accounts in weaker or, for that matter, more robust ways than through a supervenience requirement. The basic thought behind any experience requirement is that welfare concerns how well *a person's life* goes, and the "boundary" of her life—the distinction between those states, events, etc., sufficiently linked to her to affect her welfare, and those that are too remote—is best delineated in terms of her experience.

Preference: Here, as with experience, there are many ways to incorporate the general concept of "preference" into a welfare account. The preferences highlighted by a given welfare account might be prospective or retrospective, more or less tightly connected to cognitions, more or less tightly connected to choice, and so forth. The basic function of preferences, within a welfare account, is to provide a ranking of outcomes from the subject's point of view. The thought here is that welfare concerns how well a person's life goes *for her*—and thus that an outcome cannot benefit a person unless she herself favors that outcome (at some level).

Value: What is valuable, i.e., what goes on the list of objective goods or values? What makes something valuable: its connection to human perfection, to unavoidable human goals, to our shared sense of a valuable life, or something else? As already noted, these questions remain the subject of ongoing debate among objectivists. Still, there is a common idea that motivates the inclusion of objective goods or values within a welfare account: Since welfare concerns *how well* a person's life goes (or goes for her), an outcome that is less valuable for some person cannot be better for that person's welfare.

Integration: Welfare changes are, intuitively, changes in the subject's life. They are not remote changes in the world that she (or someone else) prefers or that are good in a general sense. Call this "integration": a welfare account will typically try to delineate the boundary between the subject's life and the outside, and to ensure that welfare changes occur within the boundary. This boundary might be delineated with the help of an experience requirement, or a list of objective goods (chosen

to be sufficiently "personal" to the subject), but it need not be.⁶¹ The Integration component of welfare does not necessarily reduce to Experience or Value.

Experience, Preference, Value, and Integration, or some subset of these, will be the central constituents of any colorable welfare view. They will therefore be central to my analysis of the harmfulness of risk and death.

B. RISK

At the threshold, let me distinguish between *probabilistic* and *nonprobabilistic* conceptions of risk.⁶² Probabilistic conceptions are dominant within economics,⁶³ philosophy,⁶⁴ and within the more technical literature on risk assessment,⁶⁵ and if not dominant are certainly very important within scholarship on risk regulation.⁶⁶ On the probabilistic view, the magnitude of the risk of some given adverse outcome equals the probability of the outcome. Probabilists will say that Jim is subject to a higher *risk* of death in one scenario, as opposed to a

61. Many within the preferentialist tradition have suggested that the category of welfare-relevant preferences needs to be "restricted" in some way—so that, for example, the satisfaction of someone's moral preferences doesn't count as enhancing her welfare. See Adler & Posner, *supra* note 2, at 199-200 & n.93 (citing sources). Mark Carl Overvold has done the most sustained philosophical work in developing a theory that is preferentialist, not hedonic or objectivist, but distinguishes between welfare-relevant and welfare-irrelevant preferences. See Mark Carl Overvold, *Morality, Self-Interest, and Reasons for Being Moral*, 44 PHIL. & PHENOMENOLOGICAL RES. 493, 499-501 (1984); Mark Carl Overvold, *Self-Interest and Getting What You Want*, in THE LIMITS OF UTILITARIANISM 186 (Harlan B. Miller & William H. Williams eds., 1982); Mark Carl Overvold, *Self-Interest and the Concept of Self-Sacrifice*, 10 CAN. J. PHIL. 105, 114-17 (1980) [hereinafter Overvold, *Self-Interest*]; see also David Sobel, *Well-Being as the Object of Moral Consideration*, 14 ECON. & PHIL. 249, 266-69 (1998) (criticizing Overvold's account).

62. See Clayton P. Gillette & James E. Krier, *Risk, Courts, and Agencies*, 138 U. PA. L. REV. 1027, 1071-85 (1990) (distinguishing between "expert" and "lay," i.e., probabilistic and nonprobabilistic, conceptions of risk); Ortwin Renn, *Concepts of Risk: A Classification*, in SOCIAL THEORIES OF RISK 53, 58-72 (Sheldon Krinsky & Dominic Golding eds., 1992) (surveying the conceptions of risk adopted by scientists, economists, psychologists, and sociologists); Paul B. Thompson & Wesley Dean, *Competing Conceptions of Risk*, 7 RISK 361, 367-71 (1996) (distinguishing between "probabilist" and "contextualist" conceptions of risk).

63. See, e.g., W. KIP VISCUSI, *FATAL TRADEOFFS: PUBLIC AND PRIVATE RESPONSIBILITIES FOR RISK* 19-20 (1992).

64. See, e.g., David McCarthy, *Rights, Explanation, and Risks*, 107 ETHICS 205, 210-11 (1997); Perry, *supra* note 25, at 322.

65. See Renn, *supra* note 62, at 58-61.

66. See, e.g., BREYER, *supra* note 4, at 3-6.

second scenario, just in case his *probability* of death in the first scenario is greater. A nonprobabilistic conception of risk makes the risk level depend, in part, on factors *other than* the probability of the adverse outcome at issue—such as whether the relevant persons have voluntarily chosen to take a chance of incurring that outcome; whether the outcome is particularly “dreaded” or feared by them; and whether the outcome, and the causal mechanisms leading to it, are well or poorly understood by these persons.⁶⁷ Nonprobabilists need not say that Jim’s risk of death is higher in the scenario where his probability of death is higher. For example, the risk of Jim’s death created by his 1 in 10,000 probability of dying from radiation leakage from the nearby nuclear plant might be higher than the risk of his death created by a 1 in 5,000 probability of dying in a car accident that results from his own negligence. Nonprobabilistic conceptions of risk have been influential within the experimental psychology⁶⁸ and sociological literatures on risk and risk perception,⁶⁹ and have also played some role in scholarship about risk regulation.⁷⁰

This Article focuses on the probabilistic conception of the “risk” of death. Although the analysis presented will have important implications for nonprobabilistic accounts of risk, it is beyond the scope of the Article to delineate those implications in detail.

The probabilist equates “risk of death” with “probability of death.” This equation immediately leads us to ask: What is probability? There are four standard theories of probability. This Article focuses on two (the Bayesian and frequentist

67. See, e.g., PAUL SLOVIC, *THE PERCEPTION OF RISK*, at xxi-xxxvii (2000) (surveying Slovic’s scholarship, which advances a nonprobabilistic conception of risk).

68. See *id.*; SCOTT PLOUS, *THE PSYCHOLOGY OF JUDGMENT AND DECISION MAKING* 138-40 (1993).

69. See Roger E. Kasperson, *The Social Amplification of Risk: Progress in Developing an Integrative Framework*, in *SOCIAL THEORIES OF RISK*, *supra* note 62, at 153-78 (reviewing part of the sociological literature on risk); Ingar Palmlund, *Social Drama and Risk Evaluation*, in *SOCIAL THEORIES OF RISK*, *supra* note 62, at 197-212 (same); Steve Rayner, *Cultural Theory and Risk Analysis*, in *SOCIAL THEORIES OF RISK*, *supra* note 62, at 83-115 (same); Ortwin Renn, *The Social Arena Concept of Risk Debates*, in *SOCIAL THEORIES OF RISK*, *supra* note 62, at 179-96 (same).

70. See, e.g., Cass R. Sunstein, *Which Risks First?*, 1997 U. CHI. LEGAL F. 101, 122-29 (suggesting that some but not all of the factors that underlie nonprobabilistic, lay judgments of risk should be given weight in risk regulation independent of the criterion of “decently-livable life-years”).

theories) and ignores the other two (the classical and logical theories).⁷¹ All four theories agree on the following points. Probabilities are numbers that can be attached to propositions, or to certain types of propositions. Probabilities are real numbers between zero and one. Probabilities obey a set of mathematical rules known as the "probability calculus."⁷² Take $p(S)$ to be the probability of proposition S . Then, for example, the probability calculus states that

- $p(S) = 1 - p(\text{not } S)$
- $p(S \text{ or } T) = p(S) + p(T) - p(S \text{ and } T)$
- $p(S/R)$, the so-called "conditional probability" of S given R , equals $p(S \text{ and } R)/p(R)$
- If $p(S/R) = p(S)$, i.e., if S and R are "independent,"
 $p(S \text{ and } R) = p(S) \cdot p(R)$

The Bayesian theory sees the probability of a proposition as someone's "degree of belief" in that proposition. The probability that "The President of the United States had an intimate relationship with Monica Lewinsky," for Fran, is Fran's degree of belief that the President had such a relationship. If Fran's degree of belief in that proposition is $1/3$, then according to the Bayesian theory her degree of belief in its negation is $2/3$, and her degree of belief in the proposition "the President of the United States either had or did not have an intimate relationship with Monica Lewinsky" is 1. In effect, probabilities for the Bayesian are a numerical measure of a given person's certainty or confidence in various propositions.⁷³

How, precisely, are degrees of belief assigned to propositions? Different Bayesian theories give different answers to this question. A person might have primitive judgments of the likelihood of various propositions, which

71. For good overviews of the scholarship on the nature of probability, see sources cited *supra* note 14.

72. For introductions to the probability calculus, see IAN HACKING, *AN INTRODUCTION TO PROBABILITY AND INDUCTIVE LOGIC* 23-78 (2001); MICHAEL D. RESNIK, *CHOICES: AN INTRODUCTION TO DECISION THEORY* 45-80 (1987); BRIAN SKYRMS, *CHOICE AND CHANCE: AN INTRODUCTION TO INDUCTIVE LOGIC* 130-59 (2d ed. 1975).

73. For overviews of Bayesian theories of probability, see COHEN, *supra* note 14, at 58-70; GILLIES, *supra* note 14, at 50-87; COLIN HOWSON & PETER URBACH, *SCIENTIFIC REASONING: THE BAYESIAN APPROACH* 75-97 (2d ed. 1993); KYBURG, *supra* note 14, at 68-76; WEATHERFORD, *supra* note 14, at 219-42; and Howson, *supra* note 14, at 2-13. An important, recent collection of essays exploring the Bayesian view is *SUBJECTIVE PROBABILITY* (George Wright & Peter Ayton eds., 1994).

under certain conditions can be represented numerically.⁷⁴ Alternatively, her degree of belief in a proposition might track the odds at which she is willing to make a small monetary bet for or against that proposition;⁷⁵ or, it might be a number which we can derive once we know the "utility" she attaches to different outcomes and choices, since the utility for Fran of a choice depends on the utility of the various outcomes that might result from the choice, plus her beliefs about the relative likelihood of these various outcomes.⁷⁶

The frequentist account of probability looks not to the beliefs of actual or hypothetical persons, but rather to the *frequency* with which objects possess a certain property. The basic idea behind the frequentist account is as follows: Start with a collection or "reference class" of objects. This might be a class of persons ("All men over the age of thirty"), or it might be a class of events, which are a kind of object ("All firings of a revolver," "All releases of radiation from a water-cooled nuclear plant"), or it might be a class containing some other kind of object. Then we can define the probability that an object in the reference class has the property as the *proportion of objects in the reference class with that property*. For example, the probability that a man over the age of thirty is bald equals the proportion of men over the age of thirty (the reference class of persons) with the property of baldness. The probability that a release of radiation from a water-cooled nuclear plant exceeds 100 curies is the proportion of radiation releases from water-cooled nuclear plants (the reference class of events) with the property of exceeding 100 curies.

The frequentist account of probability was first rigorously developed some seventy years ago by the statistician Richard von Mises.⁷⁷ Von Mises's version of frequentism focuses

74. See, e.g., SIMON FRENCH, *DECISION THEORY: AN INTRODUCTION TO THE MATHEMATICS OF RATIONALITY* 222-32 (1988).

75. The pioneer of this approach is de Finetti. See Bruno de Finetti, *Foresight: Its Logical Laws, Its Subjective Sources*, in *STUDIES IN SUBJECTIVE PROBABILITY* 53, 93 (Henry E. Kyburg, Jr. & Howard E. Smokler eds., 1964).

76. The classic Bayesian works that develop probability and utility measures in tandem are RICHARD C. JEFFREY, *THE LOGIC OF DECISION* 59-163 (2d ed. 1983); LEONARD J. SAVAGE, *THE FOUNDATIONS OF STATISTICS* 6-104 (Dover 1972) (1954); and Frank Plumpton Ramsey, *Truth and Probability*, in *STUDIES IN SUBJECTIVE PROBABILITY*, *supra* note 75, at 63. An excellent, philosophically sophisticated survey of these works and expected utility theory more generally is ELLERY EELLS, *RATIONAL DECISION AND CAUSALITY* 65-86 (1982).

77. See RICHARD VON MISES, *PROBABILITY, STATISTICS AND TRUTH* 1-65

specifically on events, rather than other kinds of objects, and requires that the reference class of these events be infinite and have certain other characteristics—it must be sequentially ordered and it must be “random.” Von Mises then, famously, defines probability as *the limit* of the relative frequency of an attribute with respect to a random, sequentially ordered, infinite reference class of events.⁷⁸ But no reference class of actual events—events that have actually occurred in this world—is infinite. Von Mises’s infinite reference class is, thus, best understood as an abstract, mathematical entity that (in some sense) constitutes a good theoretical representation of a class of actual events. Other frequentists have dispensed with the requirement of an infinite reference class, and instead define probability as the relative frequency of an attribute in a sufficiently large and temporally sequenced class of actual events.⁷⁹ Still others look to large or infinite classes of hypothetical events: events that *would occur* were a certain kind of experiment to be repeated over and over again.⁸⁰ In short, there are important differences of detail between frequentist views of probability. All such views, however, share the following characteristics. Probability is defined *relative to* a reference class of (actual, hypothetical, or mathematical) objects, such as events. This reference class must be *general*; it must be infinite or at least very large. The attribute in question must also be general; it must be the kind of attribute that many things, or at least more than one thing, could possess.

Because it is structured around general reference classes, general attributes, and relative frequencies, the frequentist account is unable to attach a probability number to so-called “singular” propositions absent some restructuring of such propositions in general terms.⁸¹ A singular proposition is one that predicates some attribute of a *particular* person, event or other object: for example, “The card face down is a spade” or

(2d rev. ed. 1957).

78. See *id.* at 28-29.

79. For good discussions of von Mises’s view and the variations advanced by other frequentists, in particular Hans Reichenbach, see GILLIES, *supra* note 14, at 88-112; KYBURG, *supra* note 14, at 40-53; and WEATHERFORD, *supra* note 14, at 144-218. Reichenbach’s account is presented in *THE THEORY OF PROBABILITY* (2d ed. 1949), particularly at 337-83.

80. See Karl R. Popper, *The Propensity Interpretation of Probability*, 10 *BRIT. J. FOR PHIL. SCI.* 25, 37-38 (1960).

81. See, e.g., GILLIES, *supra* note 14, at 119-25.

“John will die of cancer” or “Bush will win the next Presidential election.” The frequentist account directly attaches a probability number *only* to a proposition that (1) describes a suitably general reference class of objects (i.e., an infinite or very large class, depending on the theory), and (2) predicates a suitably general attribute of objects within the class—for example, “a man over the age of seventy is infertile.” The frequentist probability of that proposition is q , where q is the proportion of infertile men within the large or infinite class of men over seventy. But a singular proposition cannot (without reinterpretation) be assigned a relative frequency, since the reference class here has only one member, the particular object, which either has the attribute in question or lacks it. This feature of the frequentist account will be very important when we consider whether risk (in the frequentist sense) is indeed harmful.⁸²

Neither the frequentist account nor the Bayesian account is a perfectly satisfying account of probability. Intuitively, probability is both *physical* (not mind dependent) and *predicable of individuals* (not merely of classes). The probability of “this rat dying from a 10-centigray dose of radiation” seems to be a hard, physical number, like the pH of the rat’s blood, the weight of its body, or the amount of radiation released. It also seems to be a number which attaches to this particular exposure and this particular rat, like the pH, weight, and amount of radiation. However, no rigorous account of probability that satisfies both these intuitive desiderata has yet been developed. The frequentist theory is mind independent—the frequency with which some attribute occurs in a reference class is a purely physical fact—but essentially links probability to classes, not individuals. The Bayesian theory has no trouble with individualized probability ascriptions—it is perfectly intelligible to speak of some person’s degree of belief that this particular rat has a given attribute—but reduces all probabilities to what actual or hypothetical persons would believe. In short, neither theory *dominates* the other, and I will therefore consider both the Bayesian and the frequentist theories of probability in analyzing whether risk is harmful.

The classical and logical accounts of probability are not

82. See *infra* text accompanying notes 151-59.

considered here.⁸³ Both accounts insist that the probability of a proposition is established *a priori*—that a given proposition has one and the same probability in all possible worlds. But this seems wildly implausible, at least for the conception of probability relevant to risk regulation. Surely the probability that Jim's exposure to a particular dose of radiation will cause his premature death depends at least in part upon the empirical correlation between such exposures and death (on the frequentist view), or on someone's beliefs about that correlation (on the Bayesian view), or on some other, empirically contingent features of the world in which the exposure occurs. For this reason, and for others explained in the margin,⁸⁴ we can safely ignore the classical and logical construals of the probability of death and focus our attention on the Bayesian and frequentist construals.

C. CONSEQUENTIALISM, RISK, AND THE EX ANTE/EX POST DISTINCTION

As will become clear, the analysis of the questions "Is death a harm?" and "Is risk a harm?" offered in this Article is consequentialist in its presuppositions. In Part II, I ask whether two outcomes or "possible worlds" differentiated by the fact that some person (P) dies earlier in one outcome than the other are different for P's well-being.⁸⁵ In Parts III and IV, I ask a similar question for pairs of outcomes differentiated by the existence of a risk of death for P, in either the frequentist or the Bayesian sense of risk.⁸⁶

What justifies this consequentialist approach? I believe, and have argued elsewhere, that the moral criteria bearing on regulatory choices include at least some consequentialist criteria.⁸⁷ A criterion (C) is consequentialist if the right thing for the agent to do, in a choice situation, in light of C is to promote the occurrence of good outcomes—where the goodness

83. See sources cited *supra* note 14 (discussing classical and logical accounts).

84. The so-called Principle of Indifference, which assigns equal probabilities to equally possible cases, and which is central to the classical account, generates paradoxes that afflict that account and perhaps the logical account as well. See COHEN, *supra* note 14, at 43 (presenting the principle of indifference); GILLIES, *supra* note 14, at 37-49 (discussing the paradoxes).

85. See *infra* text accompanying notes 101-09.

86. See *infra* text accompanying notes 147-50.

87. See Adler, *supra* note 58, at 302-13 (arguing for moral relevance of overall well-being); Adler & Posner, *supra* note 2, at 209-16 (same).

of outcomes is specified in a suitably impartial (technically, "agent neutral") way.⁸⁸ Consequentialist criteria are not necessarily welfarist in focus. For example, a criterion that ranks outcomes in light of the ecological complexity of the outcomes, or their aesthetic value, may give a higher ranking to outcomes that are worse with respect to overall well-being, the equal distribution of well-being, or some other welfare-focused measure.⁸⁹ Still, there is at least one, morally weighty consequentialist criterion that (1) does in fact bear upon regulatory choice, or so I have argued; and (2) is welfarist in structure, indeed welfarist in the strong sense that outcomes differing in light of the criterion must differ for someone's welfare. This is the criterion of overall welfare.⁹⁰ The consequentialist criterion of equal welfare has also seemed plausible to many philosophers and political theorists, as have the related criteria of equal responsibility-adjusted welfare; equal resources for welfare; and a criterion that looks to overall weighted welfare, with the welfare of poorer persons given greater weight.⁹¹

Might the moral relevance of welfare and welfare changes, for regulators, transcend the consequentialist criteria bearing on their choices? Perhaps. Imagine that the set of criteria binding them includes at least one criterion that is both

88. For an accessible discussion of the concept of consequentialism and an overview of the debate between consequentialists and nonconsequentialists ("deontologists"), see SHELLY KAGAN, *NORMATIVE ETHICS* 59-77 (1998). For sources containing more technical discussions, see Adler, *supra* note 58, at 314 nn.188-89.

89. See KAGAN, *supra* note 88, at 59; Adler, *supra* note 58, at 315-17.

90. See Adler, *supra* note 58; Adler & Posner, *supra* note 2.

91. See Adler, *supra* note 58, at 302-13, 315-19 (discussing distributive views). Distributive views are usually consequentialist in their structure—the criteria of appropriate distribution articulated by these views typically provide an agent-neutral ranking of outcomes—although the foundations of such views are often contractarian, and for that reason they are sometimes characterized as "nonconsequentialist." See KAGAN, *supra* note 88, at 17-22, 189-303 (distinguishing between the "factors" set forth by a moral theory and its foundations, and discussing both "teleological" and "deontological" types of foundations).

Distributive views that look to the distribution of welfare resources, or to the distribution of welfare adjusted for responsibility, are welfarist not in the strong sense that the goodness of outcomes supervenes on welfare, but in the weaker sense that welfare considerations are central to these views. See Matthew D. Adler, *Legal Transitions: Some Welfarist Remarks*, J. CONTEMP. LEGAL ISSUES (forthcoming 2003) (manuscript at 4-8, on file with author) (discussing variants of welfarism).

deontological and *welfarist*: deontological because it does not enjoin the agent to promote good outcomes characterized in an agent-neutral way, welfarist because it defines morally required or permitted actions partly in terms of their welfare impact.⁹² This is a conceptual possibility. (Here is a silly example: a deontological criterion that prohibits the agent from acting in a way that reduces the welfare of anyone situated exactly ten miles from the agent at the time of choice.) Nor would I want to argue, here, that deontological criteria do not exist. But I tend to doubt that such deontological criteria as do exist make reference to welfare. Consider the classic example of a deontological constraint: the putative constraint on the actor *intentionally harming* another person (as opposed to merely acting in a way that foreseeably harms her).⁹³ I seriously doubt that “harming” for purposes of this constraint, assuming it exists, means “reducing welfare.” Imagine that I intentionally pick a flower from the common garden, just to spite you. Have I committed a deontological wrong? Intuitively not. Plausibly, “harming” here means producing some kind of physical impact on the victim which is neither necessary nor even sufficient for a welfare reduction. Or consider the plausible, deontological constraints on lying and promise breaching:⁹⁴ welfare-irrelevant lies and breaches would arguably violate these constraints, and (obviously) many welfare-reducing actions would satisfy them.

Perhaps I am wrong about the deontological irrelevance of welfare. Assume that a deontological constraint prohibits “intentional harming,” where harmings are (1) physical impacts of a certain sort that (2) reduce welfare. If so, we would need to understand whether deaths or death risks are welfare reductions for purposes of this constraint—whether they are welfare reducing in a deontological, not a consequentialist, sense. The consequentialist analysis presented in this Article does not address that question. Still, as I have explained, it is far from clear whether the question needs to be addressed.

Return, then, to consequentialism. Any consequentialist criterion C provides a ranking of outcomes—more precisely, a ranking of possible worlds, that is, of maximally specified

92. See KAGAN, *supra* note 88, at 70-137 (describing plausible deontological constraints).

93. See *id.* at 100-05.

94. See *id.* at 106-25.

outcomes.⁹⁵ These outcomes or worlds are ranked as better, worse, equally good, or perhaps incommensurably good.⁹⁶ This C-ranking of outcomes, in turn, helps determine the C-ranking of actions for a given agent, in a given choice situation. One moves from a view about the goodness of outcomes to a view about the rightness of the alternative actions that the agent might perform.⁹⁷ The agent ought to choose the action that best “promotes” good outcomes. But what does *that* mean? Different variants of consequentialism offer different answers to this question. Here are six different possibilities: in a given choice situation (1) the agent should choose the action the actual outcome of which is better or at least not worse than the actual outcomes of the other actions available for choice; (2) the agent should choose the action with the highest Bayesian probability of having the best actual outcome; (3) the agent should choose the action with the highest frequentist probability of having the best actual outcome; (4) the agent should choose the action with the largest *Bayesian expected goodness*, where the Bayesian expected goodness of an action equals $p_1 * g(O_1) + p_2 * g(O_2) + \dots + p_n * g(O_n)$, such that $g(O_i)$ is a numerical measure of the goodness of outcome O_i and p_i is a Bayesian probability that the action results in O_i ; (5) the agent should choose the action with the largest *frequentist expected goodness*, where the frequentist expected goodness of an action equals $p_1 * g(O_1) + p_2 * g(O_2) + \dots + p_n * g(O_n)$, such that $g(O_i)$ is a numerical measure of the goodness of O_i and p_i is a frequentist probability that the action results in O_i ; and (6) the agent should follow a “maximin” approach, i.e., choose the action whose worst possible outcome is best.⁹⁸

95. See MICHAEL J. LOUX, *METAPHYSICS: A CONTEMPORARY INTRODUCTION* 176-214 (2d ed. 2002) (discussing possible worlds).

96. On incommensurability, see Matthew Adler, *Incommensurability and Cost-Benefit Analysis*, 146 U. PA. L. REV. 1371, 1383-89 (1998) [hereinafter Adler, *Incommensurability*]; Matthew Adler, *Law and Incommensurability: Introduction*, 146 U. PA. L. REV. 1169 (1998) [hereinafter Adler, *Law*].

97. On this issue see RICHARD A. FUMERTON, *REASON AND MORALITY: A DEFENSE OF THE EGOCENTRIC PERSPECTIVE* 92-113 (1990); KAGAN, *supra* note 88, at 64-69. For a general discussion of moral decision making under uncertainty, see TED LOCKHART, *MORAL UNCERTAINTY AND ITS CONSEQUENCES* (2000).

98. See FUMERTON, *supra* note 97, at 92-113 (distinguishing between “actual consequence,” “probable consequence,” and “value adjusted possible consequence” conceptions of consequentialist choice). Since the probabilities employed in the latter two conceptions might be either Bayesian or frequentist, Fumerton’s tripartite distinction generates the first five of the six

As this summary suggests, risk or probability might be morally relevant, in two quite different ways, for any consequentialist criterion *C*. First, the fact that someone is at risk of some adverse (or for that matter favorable) occurrence might bear upon the goodness of outcomes. Imagine two possible worlds, O_1 and O_2 , identical except that in O_1 some event occurs and imposes a high risk of death (in either the Bayesian or frequentist sense) on *P*, while in O_2 the event does not occur. Then it might well be the case that O_1 is a worse outcome, in light of *C*, than O_2 . Second, risk or probability might be relevant in moving from the *C*-ranking of outcomes to the *C*-ranking of *actions* or choices available to a given agent in a given choice situation. Consider, for example, the Bayesian expected goodness account of consequentialist rightness. On this account, what the agent ought to do in a given choice situation, faced with possible actions $\{A_1 \dots A_m\}$, is the following: (1) for each action A_j and each outcome O_i from all possible worlds $\{O_1 \dots O_n\}$, there is a Bayesian probability, in some sense, that A_j will result in O_i ; (2) using these probabilities plus the "goodness" numbers measuring the goodness of each possible world, the agent determines the Bayesian expected goodness of each possible action; (3) the agent chooses the action with the highest Bayesian expected goodness. Within the Bayesian expected goodness account of consequentialist choice, probability numbers linking actions and worlds play an important role in determining what choices actors ought to perform, quite independent of any role that probability might have in shaping the goodness of worlds. The same is true of a frequentist expected goodness account and of the accounts that specify the consequentially right action as the action with the highest probability (frequentist or Bayesian) of having the best outcome.

Economists frequently distinguish between welfare maximization "ex post" and "ex ante."⁹⁹ The "ex post"/"ex ante" distinction, I suggest, is plausibly mapped onto the distinction between the goodness of outcomes and the goodness of actions that I have just described. To talk about "ex post" welfare

variants of consequentialism listed here. The last variant, maximin, derives from the literature on rational choice. See FRENCH, *supra* note 74, at 36.

99. See, e.g., LOUIS KAPLOW & STEVEN SHAVELL, FAIRNESS VERSUS WELFARE 437-44 (2002); Philippe Mongin & Claude d'Aspremont, *Utility Theory and Ethics*, in 1 HANDBOOK OF UTILITY THEORY 371, 437-44 (Salvador Barberà et al. eds., 1998).

maximization is really just to specify how all the different possible outcomes of one or another choice are ranked in light of the overall welfare criterion. That criterion provides a ranking or, technically, an "ordering" (partial if not complete)¹⁰⁰ of the set of possible worlds $\{O_1 \dots O_n\}$. To talk about "ex ante" welfare maximization is to specify how one moves from the overall-welfare ordering of outcomes to the overall-welfare ranking of actions $\{A_1 \dots A_m\}$ available to a given agent at a given time.

The focus of this Article is on the "ex post" inquiry. I ask whether death or risk is a feature of an outcome that diminishes the welfare of the dying or at-risk person, in that outcome, and thereby affects the goodness of that outcome in light of one or another consequentialist criterion that incorporates welfare considerations. I largely ignore the "ex ante" inquiry—for there is, after all, only so much one can do in a single article. So the picture of the normative foundations of risk regulation painted here is, in an important respect, incomplete. A complete picture would specify how one moves from the outcome ranking to the action ranking, and what role risk (in a Bayesian or frequentist sense) plays in that move separate from any role it plays in shaping the goodness of outcomes. Still, the conclusions I reach in this Article have much importance for risk regulation, as elaborated in Part V. *Whatever* one's view about the move from outcomes to actions—about the "ex ante" perspective, as it were—my claim that death and Bayesian risk but not frequentist risk affect the welfare goodness of outcomes will be highly relevant in determining what regulators ought to do.

II. IS DEATH A HARM?

Is death a harm? It is, I will claim; and this claim will prove to have much importance for risk regulation. At the outset, however, let me explain as clearly as I can what it is I will be claiming.

First, "harm" is an ambiguous term which is sometimes used to mean a welfare setback, sometimes not. For purposes

100. See Adler, *Incommensurability*, *supra* note 96, at 1401-05 (suggesting that outcomes might be incomparable with respect to overall well-being rather than better, worse, or precisely equal); Adler, *supra* note 58, at 328-30 (same); Amartya Sen, *Social Choice Theory*, in 3 HANDBOOK OF MATHEMATICAL ECONOMICS 1073, 1127-28 (Kenneth J. Arrow & Michael D. Intriligator eds., 1986) (summary of technical, social-choice literature on partial orderings).

of this Article, "harm" is an exact synonym for "welfare setback."

Second, as already explained, the spirit of my analysis is consequentialist. I analyze the impact of P's premature death on his welfare by asking whether two outcomes or "possible worlds" differentiated by the time of his death are different for P's welfare. Third, and relatedly, my claim about the harmfulness of death concerns comparative well-being. I will argue that an outcome in which a given person, P, dies at a particular point in time can be worse for P than an outcome in which he does not die at the time, but lives longer.¹⁰¹

Fourth, my claim about the harmfulness of death concerns harm to P, not harm to third parties. No one disputes that P's premature death can adversely affect the welfare of his friends, family, and other survivors. What is controversial is whether premature death can affect P's own welfare.¹⁰² Fifth, the claim concerns death itself, not the process of dying.¹⁰³ No one disputes that an outcome in which P dies painlessly or instantly, as opposed to one in which he dies a slow, painful death, can be better for P. The harder question is whether two outcomes differentiated by the fact that P continues to live in one outcome, and not in the other, can differ for P's welfare even if the process by which P dies in the latter outcome is not painful, anxiety provoking, or otherwise experientially bad.

Sixth, my claim is that death *can* be a harm, not that it *need* be. If death is necessarily harmful, then (for example) an outcome in which P dies at a certain point will be worse for him than an outcome in which he survives, but in a permanent

101. See, e.g., Fred Feldman, *Some Puzzles about the Evil of Death*, in THE METAPHYSICS OF DEATH, *supra* note 9, at 308 (arguing that the harmfulness of death consists in a comparison between the individual's actual life history and the alternative history that would have occurred had the individual not died); Nagel, *supra* note 10, at 67 (same). I am not sure what it would mean (at least within a consequentialist framework) to describe death as noncomparatively bad for a person, or to describe life as noncomparatively good. Such statements, I suggest, are not to be taken literally but actually involve covert comparisons of the premature-death outcome or continued-life outcome to some baseline outcome. See JOHN BROOME, *ETHICS OUT OF ECONOMICS* 162-73 (1999) (arguing that goodness is reducible to betterness and that Epicurus's experientialist argument against the harmfulness of death is therefore "trivially mistaken"); see also Ruth Chang, *Comparison and the Justification of Choice*, 146 U. PA. L. REV. 1569, 1569-72 (1998) (presenting a comparativist account of justified choice).

102. See Nagel, *supra* note 10, at 62.

103. See FRED FELDMAN, *CONFRONTATIONS WITH THE REAPER* 130 (1992).

vegetative state. That strikes me as untrue. Bare life—life without experiences, activities, and the other features that subserve human welfare—certainly need not be better than death.¹⁰⁴ Indeed, a welfare theorist who places emphasis on Experience will say that continuing P's life in a vegetative state is never better for him than early death; while a theorist who places emphasis on Preference would at most say that continuing P's life in a vegetative state is better if P (prior to the loss of consciousness) preferred that.¹⁰⁵

Seventh, my claim is that death can be a harm to P *in light of events that would have occurred had P continued to live*. Let us call the outcome in which P dies earlier O_E , and the outcome in which he dies later O_L . P dies at time T in O_E ; he does not die at that time in O_L , but continues to live until some later time. Obviously, if O_E and O_L differ in events that occur *before* T, then O_E can be worse for P than O_L . Assume that, in O_E , P lives in poverty and ignorance, has no friends or meaningful work, and dies at the age of forty, while in O_L P has sufficient resources, is well educated, has lots of friends and an engaging career, and dies at the age of seventy-five. It would be absurd to suggest that P here is no worse off in O_E than in O_L . The proponents of the view that "death is no harm" do not mean to assert *that*. The much more plausible and interesting assertion is this: if O_E and O_L differ only in events that occur after T, the time of P's early death in O_E , then O_E is not worse for P than O_L . And this is the assertion that I mean to dispute. Assume that

104. As Nagel puts it, "[A]lmost everyone would be indifferent . . . between immediate death and immediate coma followed by death twenty years later without reawakening." Nagel, *supra* note 10, at 62.

105. The kind of harm involved in death is therefore subtly different from the harm at issue in my discussion, below, of frequentist and Bayesian risks. In that discussion, I focus on whether frequentist and Bayesian risks are *intrinsic* harms—features of outcomes that, *ceteris paribus*, make the outcomes worse with respect to someone's welfare. But in claiming that "death is a harm," I am not claiming that death is an intrinsic harm. If O and O* differ merely in P's being alive (as would be the case where P dies at T in O, and lives longer in a vegetative state in O*), then the two outcomes may well be identical for P's welfare. Rather, my claim is that death can reduce the deceased's well-being by depriving him of intrinsic benefits he would realize if he lived longer.

Note however that many and perhaps all such benefits necessitate being alive—for example, the benefits of having certain experiences or engaging in various activities. Insofar as some welfare benefit does necessitate being alive, life is an intrinsic benefit as a component of a hybrid, and death necessarily deprives the person of that benefit. So, in this sense, the harmfulness of death is not very different from intrinsic harmfulness.

P has a sociable and professionally engaging life in O_E and dies a sudden death at the age of forty, in 1980; in O_L , he lives exactly the same life up until 1980, but does not die and continues to prosper until the age of seventy-five. Is O_E here worse for P than O_L ? I believe it is, while modern philosophical Epicureans deny that it is. That is the crux of the dispute.

More generally, where the earlier-death world O_E and the later-death world O_L differ in events that occur both before and after T, the time of P's death in O_E , then the later-occurring events, like the earlier occurring events, can make a difference to P's welfare. Assume P leads a life of poverty in O_E , and dies at the age of forty, while in O_L he leads a sociable and professionally engaged life until the age of seventy-five. Then O_L is better for P than O_E , both because of the better life P leads in O_L up to the age of forty, *and* because of the additional, fruitful years of life he enjoys in O_L after the age of forty. Longer life can be one feature of an outcome (among others) that makes that a better outcome for some person. This is the form of my "death is harmful" claim that will have the most practical relevance for regulators. Regulatory choices typically produce or prevent deaths that occur or would occur at some point in time subsequent to the choices; thus, regulators are typically comparing outcomes in which some persons live longer, as against outcomes in which those persons' lives are both shorter and otherwise different. The pure case in which P's life in the earlier-death world O_E and the later-death world O_L is just the same, up until the time T that P dies in O_E , is an analytic device designed to sharpen understanding of death's harmfulness, not a realistic scenario. However, because this pure case is particularly clear and simple, that case will be the focus of my discussion in this Part.

Finally, my claim concerns *lifetime well-being*, not momentary well-being. Lifetime well-being is, I believe, a meaningful construct.¹⁰⁶ A person's (actual or possible) life history can be compared with another (possible) life history, and one can often if not always determine that the first life history is better or worse, with respect to the welfare of that person, than the second. Momentary well-being may also be a meaningful construct. Intuitively, one can ask and often determine how well a person is doing at one moment, as

106. See John Bigelow et al., *Death and Well-Being*, 71 PAC. PHIL. Q. 119, 120-23 (1990); J. David Velleman, *Well-Being and Time*, in *THE METAPHYSICS OF DEATH*, *supra* note 9, at 329.

compared to another moment later or earlier in her life, or in a different possible life. However, momentary well-being is not my concern here. It is a serious mistake to think of the harmfulness of death in terms of momentary well-being. Imagine O_E and O_L that are identical up until T , the time of P 's death in O_E . At times prior to T , P 's momentary well-being is the same in both worlds. At T and at later times, P does not exist in O_E , and so lacks a momentary well-being in that world. Therefore P 's momentary well-being at or after T in O_L is *not* greater than her momentary well-being at or after T in O_E . The comparison is meaningless. Think of the parallel case of happiness: Since the dead cannot be happy or sad during post-mortem moments, it is meaningless to ask whether the dead are less happy at those moments, than they would be had they continued to live.¹⁰⁷

The upshot is that P 's momentary well-being in O_L is never greater than her momentary well-being in O_E ; neither during the moments before T (in the pure case we are considering), nor at or after T . Thus my claim that O_L can be better for P than O_E is, and must be, a claim about lifetime well-being. But does not lifetime well-being reduce to momentary well-being? The answer is no. As David Velleman explains,

[T]he welfare value of a life is not in general determined by, and cannot be inferred from, the amount of momentary well-being that the life contains.

....

Consider two different lives that you might live. One life begins in the depths but takes an upward trend Another life begins at the heights but slides downhill Surely, we can imagine two such lives as containing equal sums of momentary well-being. Your retirement is as blessed in one life as your childhood is in the other; your nonage is as blighted in one life as your dotage in the other.

Yet even if we were to map each moment in one life onto a moment of equal well-being in the other, we would not have shown these lives to be equally good. For after the tally of good times and bad times have been rung up, the fact would remain that one life gets progressively better while the other gets progressively worse To most people, I think, the former story would seem like a better life-story . . . in the sense that it is the story of a better life.¹⁰⁸

P 's life in O_L , her overall package of experiences, activities, accomplishments, relationships and so on, can be better than

107. But see Bigelow et al., *supra* note 106, at 134-39 (arguing that death typically lowers "temporal well-being" as well as "global well-being").

108. Velleman, *supra* note 106, at 330-31.

her life in O_E , even though there is no moment at which P's happiness or momentary well-being is greater.¹⁰⁹

To sum up, the claim I will defend in this Part (for short, "death can be harmful") is as follows: *It is possible that the lifetime well-being of a person who dies prematurely is lower than it would have been, had she continued to live, in virtue of events that would have occurred during that additional period of life.* The loss of life years can diminish a person's lifetime or overall well-being.

Isn't this obvious? In a word, no—not if one is sensitive to the thorny conceptual and normative issues that are posed by death, welfare, and their intersection. In the remainder of this Part, I will seek to accomplish two goals in tandem: first, to clarify why my claim is philosophically controversial; and, second, to defend it. Epicurus and the modern philosophical Epicureans¹¹⁰ raise serious challenges to the common-sense view that death can be harmful, but these challenges are ultimately unsuccessful. Indeed, the claim that death can be harmful is true across reasonable theories of well-being. This claim does not presuppose a commitment to a hedonic theory that gives great weight to Experience; a preferentialist theory that gives great weight to Preference; or an objectivist theory that gives great weight to Value.¹¹¹ On this score, the philosophical literature is quite misleading. Philosophical

109. I should note that the view about lifetime well-being presupposed by this Article is not identical to Velleman's view. This Article presupposes, and Velleman would agree, that lifetime well-being is a coherent and morally important concept distinct from momentary well-being. Further, my discussion of legal implications in Part V assumes that it is lifetime well-being, not momentary well-being, that is relevant to legal standards (such as cost-benefit analysis) that incorporate welfare considerations and do so in a consequentialist way. This is surely a contestable assumption, but unfortunately not one that I have space to defend within the confines of this Article. Velleman, by contrast, does not argue that momentary well-being has less moral and legal importance than lifetime well-being. See Velleman, *supra* note 106, at 347-54 (arguing for the "independent validity of momentary perspectives"). Finally, Velleman contends that two lives "containing equal sums of momentary well-being" may differ for lifetime well-being—that lifetime well-being also depends upon, as it were, the narrative structure of a life. *Id.* at 331. This last contention may well be true, but is not important for my purposes, since in standard cases of premature death the two outcomes will not contain the same aggregate momentary well-being.

110. See, e.g., Walter Glannon, *Epicureanism and Death*, 76 *MONIST* 222 (1993); O.H. Green, *Fear of Death*, 43 *PHIL. & PHENOM. RES.* 99 (1982); Stephen Hetherington, *Deathly Harm*, 38 *AM. PHIL. Q.* 349 (2001); Rosenbaum, *supra* note 9, at 117.

111. See *infra* text accompanying notes 115-35.

critics of the Epicureans have typically argued for the harmfulness of death by advancing a specific welfare theory: hedonism, preferentialism, objectivism, or some hybrid. As we shall see, the claim that death can be harmful is much more robust than these critics have suggested.

Why do Epicureans deny that death can be harmful? Let us start with Epicurus's own words on the subject and those of an early disciple, Lucretius. Epicurus argues,

Accustom thyself to believe that death is nothing to us, for good and evil imply sentience, and death is the privation of all sentience, . . . Death, therefore, the most awful of evils, is nothing to us, seeing that, when we are, death is not come, and when death is come, we are not. It is nothing, then, either to the living or to the dead, for with the living it is not and the dead exist no longer.¹¹²

Lucretius explains,

Death therefore to us is nothing, concerns us not a jot, . . . For he whom evil is to befall, must in his own person exist at the very time it comes, if the misery and suffering are haply to have any place at all; but since death precludes this, and forbids him to be, upon whom the ills can be brought, you may be sure that we have nothing to fear after death, and that he who exists not, cannot become miserable . . .¹¹³

One clear theme in these passages concerns the experiential element of well-being. Epicurus contends that "good and evil imply sentience, and death is the privation of all sentience"; Lucretius reasons that death precludes harm ("evil") because it precludes "misery and suffering," that is, unpleasant, dispreferred, or undesirable mental states. Modern philosophical Epicureans, too, suggest that death cannot be harmful for the person who dies because, at death, that person loses the capacity to have experiences—the capacity to feel pains or pleasures, to have beliefs, or to possess other types of mental states. Steven Rosenbaum, a prominent contemporary defender of the Epicurean view, writes,

I offer [and defend as cogent] the following reconstruction of Epicurus's argument. In formulating the arguments as I do, I attempt to do justice to Epicurus's philosophical insight, caring less for historical accuracy than for verisimilitude. The reconstruction

112. See Rosenbaum, *supra* note 9, at 121 (alteration in original) (quoting Epicurus's "Letter to Menoeceus").

113. See FELDMAN, *supra* note 103, at 129 (first alteration in original) (quoting Lucretius's *De Rerum Natura*).

runs as follows:

A state of affairs is bad for person *P* only if *P* can experience it at some time. . . .

. . . *P*'s being dead is not a state of affairs that *P* can experience at some time.

....

THEREFORE, *P*'s being dead is not bad for *P*.¹¹⁴

Like Epicurus and Lucretius before him, Rosenbaum is in effect voicing some version of the following general argument: Because the right theory of well-being includes an Experience element, death cannot be harmful for the person who dies.

Is the argument cogent? As I have already noted, the Experience element of welfare might be specified in different ways.¹¹⁵ Standardly, that element is specified as a *supervenience* requirement, which states that two outcomes cannot differ, with respect to the welfare of some person *P*, if *P* has the same mental states in both outcomes.¹¹⁶ Note, however, that my "death can be harmful" claim is *consistent* with a supervenience requirement. To say that "death can be harmful" is to say that *O_E* (the earlier-death outcome) can be worse for *P* than *O_L* (the later-death outcome) even if *O_E* and *O_L* are identical in features that occur before *T*, the time of *P*'s early death in *O_E*. Imagine, therefore, that *O_E* and *O_L* are identical in their pre-*T* features but differ in their post-*T* features as follows: In *O_E*, *P* is dead and experiences nothing, while in *O_L*, *P* continues to live and experiences various pains, pleasures, and other mental states. In more colloquial terms, imagine a case in which *P* would continue to possess a mental life (in *O_L*), were he not to die (in *O_E*). This is hardly an esoteric case. Although premature death does not necessarily cut short a mental life that would otherwise continue—the premature death of a comatose person, who would die later without regaining consciousness were the premature death not to occur, shortens life but not mental life—it is certainly normal, unsurprising, and in any event quite possible for premature death to deprive a person of experiences she would otherwise have. So it is normal, unsurprising, and in any event quite possible for *P* *not* to have the same mental states in *O_E* and *O_L*. But if *P* does *not* have the same mental states in *O_E* and *O_L*, the

114. Rosenbaum, *supra* note 9, at 121-22.

115. See *supra* text accompanying notes 60-61.

116. See, e.g., KAGAN, *supra* note 88, at 180-81; Bernstein, *supra* note 59, at 42-47.

supervenience requirement is satisfied. As Jeff McMahan puts it: "[A]lthough one does not experience death, it does affect one's experience—by limiting or ending it."¹¹⁷

McMahan's comment here about "experienc[ing] death" suggests a different specification of the Experience element—one that makes a particular type of mental state, namely awareness or cognizance, a precondition for welfare changes.¹¹⁸ Similarly, when Rosenbaum writes that "a state of affairs is bad for person *P* only if *P* can experience it at some time" and "*P*'s being dead is not a state of affairs that *P* can experience at some time,"¹¹⁹ he is seemingly using the term "experience" as synonymous with awareness or cognizance. Imagine, then, that our welfare theory includes an "awareness requirement" which stipulates that some feature (*F*) of an outcome can make a difference to *P*'s well-being only if *P* is aware of *F* when or after it occurs. If this latter requirement obtains, then death cannot be harmful. Why? Consider our test case for the harmfulness of death: two outcomes *O_E* and *O_L*, where *P* dies at *T* in *O_E* and the outcomes are identical in their pre-*T* features. The crucial point to see is that *P* in *O_E* cannot be cognizant of its post-*T* features (specifically, the fact that she is not alive and does not engage in the various activities, relationships, etc. that she would have pursued had she not died). But the post-*T* features of *O_E* are the only features that differentiate it from *O_L*. Therefore, given the awareness requirement, *O_E* cannot be worse for *P* than *O_L*.¹²⁰

117. Jeff McMahan, *Death and the Value of Life*, in THE METAPHYSICS OF DEATH, *supra* note 9, at 234; see also BROOME, *supra* note 101, at 170-73 (arguing that experientialism about welfare permits death to be harmful, since death can deprive a person of experiences he would have if alive).

118. See McMahan, *supra* note 117, at 234-45 (distinguishing between "Narrow Experience Requirement" and "Wide Experience Requirement").

119. Rosenbaum, *supra* note 9, at 121-22. Rosenbaum goes on to state that experience has a causal aspect, and that the concept in one sense entails the subject's awareness but in another sense simply entails a causal connection between the "experienced" event or state and the subject. *Id.* at 124. I do not believe that redefining Experience to include a causal element undermines the argument presented in this section for the harmfulness of death, since what is at issue is comparative lifetime well-being across different possible worlds, not some putative effect that death causes and the dead person "experiences" (in the causal sense) after her death.

120. The proposition that *O_L* is better for *P* than *O_E* is, at first blush, consistent with the awareness requirement, since *P* in *O_L* is alive after *T* and can be cognizant of the features that differentiate that outcome from *O_E*. One might then (1) permit the asymmetry of "better" and "worse" and conclude that *O_L* really is better for *P* even though *O_E* is not worse; or (2) demand

How should the defender of the claim that death can be harmful respond to this powerful, Epicurean line of argument? Unless cognizance and awareness somehow continue after death—and my assumption here is that they do not, that death necessarily means the end of the person's existence and therewith his awareness and his mental states more generally—the argument from an awareness requirement to an Epicurean conclusion seems ironclad.¹²¹ I will therefore need to show that the awareness requirement itself is mistaken. Let me postpone the issue, for the moment,¹²² and examine whether the Preference component of welfare also argues for Epicureanism. As we shall see, Preference might be understood to generate an "endorsement requirement" that is closely related to the awareness requirement. I will delineate this new requirement, and then try to demonstrate why the twin requirements, endorsement and awareness, are too robust to be incorporated in the correct theory of welfare.

The Preference component of welfare stipulates that the welfare ranking of outcomes, for a given person, depends at least in part on the person's *own* ranking of the outcomes—on his preferences as between the outcomes. The stipulation, thus stated, is quite general. Are the preferences actual or hypothetical? Are they retrospective or prospective? Are they closely or only tenuously linked to choice? Need they have a certain kind of cognitive or affective component? Welfare theorists who share a basic commitment to Preference can disagree about the answers to these more specific questions.¹²³ The crucial question, for our purposes here, concerns the timing of the preferences that are incorporated in the welfare theory.

Let us first consider the case of a welfare theory that incorporates a Preference component, and specifies in some way which preferences are determinative of welfare—for example, by stipulating that such preferences are actual preferences which are sufficiently informed and sufficiently grounded in deliberation—but does not define the category of

symmetry and conclude that O_L is not really better for P than O_E since (given the arguments stated in the text) O_E is not worse; or (3) demand symmetry and conclude that O_E is really worse for P since O_L is better. On these issues, see McMahan, *supra* note 117, at 235-40.

121. *But see supra* note 120 (describing a line of argument that, together with the putative symmetry of "better" and "worse," seeks to square the awareness requirement with the harmfulness of death).

122. *See infra* text accompanying notes 127-33.

123. *See supra* text accompanying notes 60-61.

welfare-determinative preferences in a way that excludes prospective preferences. Clearly, this type of preferentialist welfare theory permits death to be harmful. Assume that P prospectively prefers continued life, to premature death, except under unusual circumstances where continued life would be miserable or painful. P carefully and with good information considers how he would rank a world like O_E , as compared to a world like O_L , and determines that he would give a lower ranking to O_E . Further, assume that this preference or ranking of P's is his ranking at all times, both in O_E and in O_L . Finally, assume that the actual world is O_E . P ends up dying prematurely at time T, as opposed to enjoying continued life. Under these circumstances, premature death can be worse for P than continued life even though the premature-death world (O_E) and the continued-life world (O_L) are identical in their pre-T features. P had a preference, prior to his premature death, not to die prematurely and forego the experiences and activities he would otherwise enjoy. Although this preference was prospective, it was actual, informed, and deliberative, and thus counts as a welfare-determinative preference within the theory just described.¹²⁴

By contrast, if the Preference element of welfare is specified so that the only welfare-determinative preferences are retrospective, not prospective, death cannot be harmful. Specifically, imagine that our welfare theory includes an "endorsement requirement," which states that a feature F can make an outcome better (or worse) for some person P only if P

124. See generally Bernard Williams, *The Makropulos Case: Reflections on the Tedium of Immortality*, in *THE METAPHYSICS OF DEATH*, *supra* note 9, at 73 (presenting a preferentialist account of death's harmfulness). As Williams notes, "[A] man could . . . have a reason from his own point of view to prefer a possible world in which he went on longer to one in which he went on for less long . . ." *Id.* at 78.

Isn't it incoherent to imagine a human person, with the epistemic and conceptual limitations characteristic of humans, actually having preferences with respect to entire possible worlds? This question raises deep issues for the actual Preferentialist that I cannot, and need not, consider here. Whatever the answer, my claim that a Preference requirement permits death to be harmful, if the relevant preferences can be prospective, holds true. If the answer to the question is "yes," then the Preferentialist will either shift entirely to idealized rather than actual preferences, or instead will weaken her demand with respect to actual preferences (for example, by requiring that the person actually prefer some of the features differentiating the two worlds in question). In either case, O_E can be worse than O_L , assuming prospective preferences.

(dis)endorses F during or after the time F occurs.¹²⁵ Consider now our two outcomes O_E and O_L . The feature of O_E that putatively makes it worse for P is the fact that P does not enjoy various welfare goods, which he would have enjoyed (in O_L) had he not died prematurely. But this feature of O_E does not occur until P dies. Therefore, P cannot have a retrospective preference with respect to this feature. Since he is dead after T, he cannot have *any* preferences, actual or hypothetical. P might prefer, before his premature death, not to die prematurely; but he cannot prefer, after his premature death, not to have died prematurely.¹²⁶

The endorsement requirement is one possible specification of the Preference element of welfare. The awareness requirement, discussed earlier, is one possible specification of the Experience element of welfare.¹²⁷ But these two requirements are closely linked. Both are based on the intuitive idea that P's own perspective is crucial to his welfare.¹²⁸ Both requirements are also grounded in the thought that the authoritative perspective is backward looking, not forward looking. It is only *after* experiences, activities, and other putative constituents of welfare have occurred (or failed to occur) that we can really assess their welfare impact—or so the argument goes. P might prospectively prefer or disprefer some feature F of an outcome, but this preference (since prospective) merely concerns a possible occurrence. Possible features, when actualized, can surprise even the best informed. It is a truism that people often end up disliking what they imagined they would like, or endorsing what they imagined they would disprefer.¹²⁹ Thus our prospective likes and dislikes with respect to possible occurrences should be given much less weight, in a welfare theory, than our retrospective likes and

125. Cf. SUMNER, *supra* note 8, at 122-83 (emphasizing the importance of retrospective evaluations rather than prospective desires to well-being).

126. The symmetry of benefit and harm might be adduced to show that death can be harmful, consistent with an endorsement requirement, since life can be beneficial. See *supra* note 120 (noting a parallel argument in the context of the awareness requirement). Because, I believe, the endorsement requirement, like the awareness requirement, can successfully be attacked head-on, see *infra* text accompanying notes 127-33, I do not pursue the symmetry issue here.

127. See *supra* text accompanying notes 118-20.

128. SUMNER, *supra* note 8, at 26-80.

129. See *id.* at 129-33 (arguing against the equation of welfare with the satisfaction of antecedent desires).

dislikes with respect to what actually happens. Until P has both become aware of F and developed a retrospective preference for or against F, F cannot make an outcome better or worse for P.

The line of reasoning just articulated has real force, but ultimately fails to persuade. Retrospective awareness and preference are *not* preconditions for welfare. Consider the plausible "idealizing" conditions, other than temporal position, that make a perspective or point of view relevant for welfare assessment. Plausibly, the point of view must be one where the person, P, is sufficiently informed¹³⁰ about the features being assessed; those features must be sufficiently vivid to her;¹³¹ she must be sufficiently thoughtful about the features; she must be sufficiently engaged by the ranking exercise;¹³² and so on. Virtually all the plausible "idealizing" conditions I can think of—even conditions like vividness and engagement—can be true of prospective as well as retrospective points of view. One exception is an *actuality* condition: a requirement that the features under assessment be actual, not merely possible, features. A moment's thought, however, will show that the actuality condition is too demanding. When P prospectively evaluates a possible life history O, as against another possible life history O*, he is comparing one group of possible features to another. When P retrospectively evaluates his actual life history O, as against a possible life history O* that did not occur, he is comparing one group of actual features to a group of possible features. If features must be actual to be properly evaluated, then retrospective evaluation, like prospective evaluation, falls short. Intrapersonal judgments of comparative well-being *always* involve comparisons of multiple outcomes, at least some of which must be merely possible, not actual.

Note further that the awareness requirement and endorsement requirement make prospective points of view flatly irrelevant to welfare evaluation, rather than merely—and more plausibly—discounting their relevance vis-à-vis retrospective points of view. Imagine that P prospectively prefers one outcome to another, and then later, retrospectively,

130. Cf. sources cited *supra* note 45 (surveying, but criticizing, full-information preferentialist accounts of welfare).

131. RICHARD B. BRANDT, *A THEORY OF THE GOOD AND THE RIGHT* 111-12 (1979).

132. On these last two elements, see GRIFFIN, *supra* note 18, at 26-31.

switches his ranking. Perhaps we might give priority to the retrospective ranking, in determining which outcome is really better for P. But the premature-death case is different. Here, we have (or may have) a prospective ranking by P of continued life over premature death, and no retrospective ranking whatsoever by P if he does die prematurely. Surely this prospective ranking could be good enough—sufficiently informed, engaged, vivid, etc.—to be authoritative for P, in the absence of a conflicting retrospective ranking.

Finally, note that the awareness requirement and endorsement requirement would apply to all changes in a person's circumstances that preclude awareness and endorsement—not just death. The awareness and endorsement requirements, taken separately or together, imply that no such changes can be harmful. Consider the family of experiential changes that preclude awareness and endorsement. In O, P lives a normal life. In O*, P continues to live, but suffers catastrophic brain damage which lowers his IQ to fifty and prevents him being aware of this damage and from having all but the most infantile preferences.¹³³ Or, O* is an outcome in which P continuously experiences terrible pain—pain so intense that more advanced cognitive states like belief (awareness) and preference are simply crowded out. Or, O* is an outcome in which P lives but is unconscious. Or, in O* P develops a severe and rapidly progressing case of Alzheimer's disease. Intuitively, these kinds of radical, negative changes in the quality of someone's experiential life are very bad, indeed, for her welfare. And I see no real reason why a welfare theory should insist otherwise. What changes is experience (so the Experience requirement is satisfied); the persons involved might strongly and with good reason prefer, beforehand, that the changes not occur (so Preference is satisfied); and the changes either introduce bad things into the person's life, or deprive her of good things (so Value is satisfied, as is Integration). Yet an awareness requirement and/or an endorsement requirement would *inter alia* lead to the conclusion that catastrophic brain damage, radical and permanent pain, continuing unconsciousness, severe Alzheimer's, and the like are not welfare setbacks.

Let me summarize the analysis to this point. We are

133. See Nagel, *supra* note 10, at 65-66 (inferring death's harmfulness from the harmfulness of severe brain damage).

considering whether death can be a harm: whether a life history O_E ending in premature death for P at T can be worse for P than an alternative life history O_L differing only in events that occur after T . I suggested that the Experience component of welfare, if specified in certain ways, is consistent with the claim that “death can be harmful” and the same holds true for Preference. Experience and Preference can be specified in ways that contradict the “death can be harmful” claim, but these specifications—the awareness and endorsement requirements, respectively—were shown to be mistaken. I see no other plausible specifications of Experience and Preference that contradict the claim, and so will now consider whether Value or Integration do so.

The Value element of welfare poses no difficulty for the “death can be harmful” claim. Different theorists generate different lists of objective “goods” and “bads” and adduce different rationales for these lists.¹³⁴ Still, any viable list and rationale will count *certain* activities and experiences as valuable; and death, then, can be harmful within the framework of this list and rationale, because death precludes activity or experience. For example, Parfit offers “moral goodness, rational activity, the development of one’s abilities, having children and being a good parent, knowledge, and the awareness of true beauty”¹³⁵ as a paradigmatic list of goods that an objectivist might adduce. Premature death prevents the dying person from performing any additional morally good acts, from continuing to engage in rational activity, from developing his abilities any further, from having more children or engaging in good parenting, from acquiring more knowledge, or from being aware of true beauty.

While the Value element of welfare doesn’t provide Epicureans with ammunition for their view that death is harmless, the Integration element clearly does. There are a cluster of difficult metaphysical problems raised by the “death can be harmful” thesis—problems concerning the *subject* of the putative harm and the *time* of the putative harm. Epicurus alludes to these problems when he writes: “Death, therefore, the most awful of evils, is nothing to us, seeing that, when we are, death is not come, and when death is come, we are not.”¹³⁶

134. See *supra* text accompanying notes 48–56.

135. PARFIT, *supra* note 18, at 499.

136. Rosenbaum, *supra* note 9, at 121.

As Thomas Nagel explains,

[T]here are special difficulties, in the case of death, about how the supposed misfortune is to be assigned to a subject at all. There is doubt both as to *who* its subject is and as to *when* he undergoes it. So long as a person exists, he has not yet died, and once he has died, he no longer exists; so there seems to be no time when death, if it is a misfortune, can be ascribed to its unfortunate subject.¹³⁷

Consider, once more, my claim that an earlier-death outcome O_E can be worse for a person P than a later-death outcome O_L , even if the two outcomes are identical in events that occur before the time T of P 's death in O_E . Isn't this claim metaphysically incoherent? If P dies prematurely, at what time is she worse off than if she had continued to live? In short, at what time does premature death harm P ? Presumably not before time T , because P has not died yet. Yet it also appears that premature death does not harm P at any time *after* T , because at those later times P does not exist and cannot be harmed. Finally, it seems odd to think that the harm of premature death is an instantaneous harm, which occurs neither before T , nor after T , but precisely at the instant of P 's death.

The problems I have just sketched are, in effect, problems of Integration. My claim is that P 's well-being is lower in one outcome (O_E), as compared to another outcome (O_L). Integration demands that the difference between O_E and O_L be a difference that "concerns" P . This difference must not be merely a difference in the general or impersonal goodness, beauty, or value of the outcomes. It is not enough to show that O_E is worse than O_L from some detached viewpoint. O_E must be worse than O_L *for* P .¹³⁸ More precisely, the Integration element of welfare plausibly includes the following requirement: A person's lifetime well-being is lower or higher in the actual world O , as compared to a counterfactual alternative O^* , only if P in O (at some time) has the property of being worse or better off than her counterfactual self in O^* . Yet there is, apparently, no time in the premature-death world O_E where P would have the property of being worse off than her counterfactual longer-lived self in O_L .

Much philosophical ink has been spilled proposing, and criticizing, various solutions to the related problems of identifying (1) the time at which premature death is harmful

137. Nagel, *supra* note 10, at 64.

138. See *supra* text accompanying note 61.

and (2) the subject of that harm. Indeed, more attention has been given to these problems than to any other issue discussed in the "Is death harmful?" literature. One general approach to solving the timing and subject problems is to argue that dead people do, in some sense, exist. Harry Silverstein and Palle Yourgrau have taken this approach.¹³⁹ The other general approach is to concede that dead people do not exist, but then either (I) identify a discrete moment or stretch of time at which premature death is harmful, or (II) argue that there is no discrete moment or stretch of time at which death is harmful. Joel Feinberg, George Pitcher, and Julian Lamont fall in category I. Feinberg and Pitcher argue that premature death is harmful for the dying person *before* the premature death occurs.¹⁴⁰ Lamont argues that premature death is harmful for the dying person at the point of death and thereafter.¹⁴¹ Fred Feldman, Thomas Nagel, and William Grey fall in category II. Feldman argues that premature death is harmful "eternally," i.e., at all times:

[W]hen we say that [the dying person's] death is bad for her, we are really expressing a complex fact about the relative values of two possible worlds. If these worlds stand in a certain value relation, then (given that they stand in this relation at any time) they stand in that relation not only when [the person] exists, but at times when she does not.¹⁴²

Nagel argues that premature death is harmful at no time, while Grey argues that the harm in premature death *does* have a temporal location, but this location is vague—it cannot be precisely defined.¹⁴³

One important point has been missed, or at least under-emphasized, in this philosophical back and forth. The specific problem of explaining how premature death can be harmful for the dying person is simply one instance of a more general problem, namely explaining how a person's lifetime well-being can *ever* be higher or lower than a counterfactual alternative.

139. See Harry S. Silverstein, *The Evil of Death*, in *THE METAPHYSICS OF DEATH*, *supra* note 9, at 95, 110-15; Palle Yourgrau, *The Dead*, in *THE METAPHYSICS OF DEATH*, *supra* note 9, at 137, 142-45.

140. See Joel Feinberg, *Harm to Others*, in *THE METAPHYSICS OF DEATH*, *supra* note 9, at 171, 186-88; George Pitcher, *The Misfortunes of the Dead*, in *THE METAPHYSICS OF DEATH*, *supra* note 9, at 159, 168.

141. Julian Lamont, *A Solution to the Puzzle of When Death Harms Its Victims*, 76 *AUSTRALASIAN J. PHIL.* 198, 209-12 (1998).

142. Feldman, *supra* note 101, at 321.

143. See William Grey, *Epicurus and the Harm of Death*, 77 *AUSTRALASIAN J. PHIL.* 358, 363-64 (1999); Nagel, *supra* note 10, at 67.

To see this, imagine that we are comparing the actual world *O*, to a counterfactual alternative *O** in which *P* dies at exactly the same time as in *O*. The issue of premature death is thus eliminated. But here, too, the Integration requirement seemingly cannot be satisfied. The Integration requirement stipulates that *P*'s lifetime well-being is higher or lower in *O*, as compared to *O**, only if *O* is better or worse *for P* than *O**—only if the difference between the outcomes is expressible as a relational property of *P*, such that *P* in *O* has the property of being better or worse off than his counterfactual counterpart in *O**. Yet at what time is *O* better or worse for *P* than *O**? Seemingly not before *P*'s death. At that point, *P*'s life is not over, and his actual lifetime well-being has not yet been determined. How can we say that *P*'s life has in fact been better or worse for him, as compared to a counterfactual alternative, until *P* has died—until we know what path his life history has in fact taken? On the other hand, the time at which *O* is better or worse for *P* than *O** cannot, seemingly, be the time *P* dies or thereafter, because *P* does not exist at or after that time.

In short, if it is metaphysically incoherent to say that premature death is harmful for the dying person, then more generally it is metaphysically incoherent to make statements to the effect that one life history is better or worse for a person than another life history. This leaves three options. The first is to give up the concept of lifetime well-being. But lifetime well-being is a coherent concept, and indeed is crucial to moral assessment. As David Velleman and others have persuasively argued, judgments of well-being cannot be reduced to judgments of momentary well-being.¹⁴⁴ The first option should be rejected.

The second option is to give up the Integration requirement. If the Integration requirement is abandoned, then to say that *P*'s lifetime well-being in *O* is higher or lower than her lifetime well-being in *O** does not imply that *O* is better or worse *for P* than *O**. The third option is to preserve the Integration requirement, and show how it is possible for *P* in one world to be better or worse off—with respect to lifetime well-being—than her counterfactual counterpart in another outcome. This is where the various proposed solutions to the “timing” and “subject” puzzles advanced by Silverstein,

144. See *supra* notes 106-09 and accompanying text.

Yourgrau, Feinberg, Pitcher, Lamont, Feldman, Nagel and Grey come into play.¹⁴⁵ These “solutions” are different ways to pursue the third option—different ways to show that one life history can be worse for the subject than the alternative. Note that this third option *must* be correct if the first two are properly rejected. If lifetime well-being is a coherent concept, and the Integration requirement obtains, then one of the proposed solutions to the timing and subject puzzles must be cogent.

Which one? I cannot hope to address that extremely difficult metaphysical question here. My contribution, if any, is to have shown that the timing and subject puzzles are generally implicated in comparisons of lifetime well-being, and reciprocally to have argued that the coherence of such comparisons plus an Integration requirement implies the solubility of these puzzles. I do not know when, if ever, it becomes true that P in O_E , the premature-death world, is worse off than P in O_L , the continued-life world.¹⁴⁶ What I have argued is this: If P in O_E is never worse off than P in O_L , and if this implies that premature death is not a harm for the dying person, then it follows that nothing which alters someone’s life-plan can be harmful or beneficial for him. Because this is an absurd conclusion, we should infer that an Integration requirement poses no obstacle to the claim that death can be harmful. Although the metaphysicians continue to debate how to reconcile the harmfulness of death with the Integration requirement, what is clear is that some such reconciliation

145. Since Nagel and Feldman do claim that death can be harmful for the dead person, I am inclined to interpret them as pursuing the third option (preserving Integration) rather than the second (abandoning it) even though they do not identify a discrete stretch of time during which the harm occurs.

146. I am inclined to think that those who die prematurely are worse off, as compared to hypothetical alternative life histories, at the point of death and thereafter (and not before). This is Lamont’s view, and it rests upon the quite plausible premise that nonexistent things can have relational properties. See David Hillel Ruben, *A Puzzle About Posthumous Predication*, 97 PHIL. REV. 211, 222-36 (1988). But see Jack Li, *Commentary on Lamont’s When Death Harms Its Victims*, 77 AUSTRALASIAN J. PHIL. 349, 352-53 (1999) (criticizing Lamont). If indeterminism is true and our futures are indeed open, it would not be true that the person who ends up dying prematurely was actually worse off, earlier on, than in the counterfactual outcome where he continues to live. The “no time” solutions are, I think, inconsistent with Integration: If P has the relational property of being worse off than some counterfactual counterpart, then she possesses this property at some time. I cannot, however, defend Lamont’s position here, and—as explained in the text—do not need to do so in order to argue that death is harmful.

must be possible if any judgments of comparative lifetime well-being are.

III. IS THE RISK OF DEATH A HARM? RISKS AS RELATIVE FREQUENCIES

Is the risk of death a harm? Here, as with the question "Is death a harm?", I use the term "harm" to mean a setback to well-being. Here, as there, I adopt a consequentialist approach and focus on what economists would call "ex post" welfare impacts—namely, whether the risk of death is a feature of an outcome or possible world that changes a person's welfare in that world.¹⁴⁷ Finally, my focus remains comparative rather than noncomparative well-being, lifetime rather than momentary well-being, and harm to the person herself (here, the person at risk) rather than harm to third parties.¹⁴⁸

In short, the question I intend to ask is this: Does (or can) a risk of death lower the lifetime well-being of a person who is subject to that risk, as compared to an outcome in which she is not at risk? In answering this question, I will focus on probabilistic conceptions of risk and, within that category, on the two most attractive accounts of probability: the relative-frequency account and the Bayesian account. Frequentist accounts of risk are discussed in this Part and Bayesian accounts in Part IV.

Throughout this Part and the next, I will assume that risks of death occur as a result of *events*.¹⁴⁹ A risky event might be simple in structure (a nuclear reactor near to P's home releases some radiation at a particular place and at a particular time), or it might be more complex (P's supervisor misinforms him

147. See *supra* Part I.C. (explaining the consequentialist and "ex post" approach of this Article).

148. See *supra* text accompanying notes 101-09 (explaining that the analysis of death's harmfulness focuses on comparative well-being, lifetime welfare setbacks, and harm to the person who dies).

149. Probabilities, as I have explained, are numbers attached to propositions. See *supra* text accompanying notes 71-72. But those propositions—at least on a frequentist account—are, paradigmatically, propositions that predicate some property of some event. See, e.g., COHEN, *supra* note 14, at 40-51 (summarizing the frequentist view). And presentations of the probability calculus that are agnostic as between frequentist, Bayesian, classical, and logical accounts of probability are often formulated in terms of events. See HACKING, *supra* note 72, at 38 ("Most statisticians and most textbooks of probability talk about the probability of events.").

about workplace safety procedures regarding certain chemicals; P, over the course of his adult life, consumes a diet that makes him particularly susceptible to those chemicals; and P's coworkers, at different points, engage in negligent behavior in an area of the factory close to the chemicals). Risky events include natural events as well as human actions and—within the latter category—P's own actions as well as the actions of others.

Events that produce premature death risks also frequently lead to other kinds of welfare setbacks. Imagine that a toxic waste spill not only produces a death risk for P, but damages his property and injures his pets. Or, imagine that the consumption of a dangerous food additive increases P's risk of death, and also causes P physical discomfort or injury of some kind. For that matter, imagine that the risk of premature death imposed on P by some event is *realized*, so that P suffers not only the harm of risk (such as it may be) but the additional and separate harm of death itself.

The general issue under consideration here is whether the imposition of a risk of death on P by an event E is a feature of an outcome that can make the outcome worse for P's lifetime well-being, as compared to an outcome in which P suffers no such risk. I will sharpen and simplify analysis of this issue by focusing on the simple case where the outcomes under comparison are identical *except for* the risk of death.¹⁵⁰ In this simple case, O_R , the riskier world, contains some event (E) that imposes a substantial risk of premature death on P; O_S , the safer world, does not contain that event; and O_S and O_R are otherwise identical with respect to P's well-being. The question I will address, in detail, is whether—given the multiple considerations plausibly relevant to well-being, specifically Preference, Experience, Value and Integration—P's lifetime well-being in O_R is lower than his lifetime well-being in O_S . In short, is the risk of death itself a harm? As we shall see, the answer is “no” on a frequentist account of risk, but “yes” on a Bayesian account.

A. WHAT IS THE RISK OF DEATH? A FREQUENTIST ACCOUNT

At the threshold, a conceptual problem arises. Is the risk of death an intelligible concept for the frequentist? O_R is a

150. See *supra* text accompanying notes 105-06 (adopting a similar procedure for analyzing whether death is harmful).

world in which P is at substantial risk of premature death from some event E. More precisely, O_R is a world in which E occurs and the proposition "P will die earlier than he would have had E not occurred" has a high probability. But does the frequentist account permit us to attach a probability number to this kind of proposition—a so-called singular proposition?

Some frequentists (call them "hard-core" frequentists) would say no. Consider what Richard von Mises, the statistician and philosopher whose book *Probability, Statistics, and Truth* (1928) is a landmark in the development of the frequentist view of probability, had to say about the probability of death.

The rational concept of probability, which is the only basis of [the] probability calculus, applies only to problems in which either the same event repeats itself again and again, or a great number of uniform elements are involved at the same time. Using the language of physics, we may say that in order to apply the theory of probability we must have a practically unlimited sequence of uniform observations [i.e., a suitably large reference class].

....

... When we speak of 'the probability of death,' the exact meaning of this expression can be defined in the following way only. We must not think of an individual, but of a certain class as a whole, e.g., 'all insured men forty-one years old living in a given country and not engaged in certain dangerous occupations.' A probability of death is attached to this class of men or to another class that can be defined in a similar way. We can say nothing about the probability of death of an individual even if we know his condition of life and health in detail. *The phrase 'probability of death,' when it refers to a single person, has no meaning for us.* This is one of the most important consequences of our definition of probability . . .¹⁵¹

The hard-core frequentist has a straightforward answer to the question "Is the risk of death a harm?" The answer is no, she will claim, because the putative harm (the harm to a particular person arising from his risk of death) presupposes a "risk" or "probability" (the risk or probability of a particular person dying) that in fact does not exist.¹⁵² According to von Mises, "[I]t is utter nonsense to say . . . that Mr. X . . . has [a given] probability . . . of dying in the course of the next year;"¹⁵³ a fortiori, for von Mises and other hard-core frequentists, it should be "utter nonsense" to say that Mr. X's risk of death harms him.

151. VON MISES, *supra* note 77, at 11 (emphasis added).

152. *See id.* at 16-18 (discussing the probability of death).

153. *Id.* at 17-18.

Yet not all frequentists adopt such a position. For example, Hans Reichenbach, who along with von Mises is standardly seen as a leading contributor to the frequentist school, suggests that a proposition which ascribes an attribute to a particular object rather than a class of objects might be assigned a probability as follows: first, subsume the object under an appropriate reference class; second, calculate the frequency of the attribute in that reference class. According to Reichenbach, "If we are asked to find the probability holding for an individual future event, we must first incorporate the case in a suitable reference class."¹⁵⁴ He elaborates,

I regard the statement about the probability of the single case, not as having a meaning of its own, but as representing an elliptic mode of speech. In order to acquire meaning, the statement must be translated into a statement about a frequency in a sequence of repeated occurrences. The statement concerning the probability of the single case thus is given a *fictitious meaning*, constructed by a *transfer of meaning from the general to the particular case*. The adoption of the fictitious meaning is justifiable . . . because it serves the purpose of action to deal with such statements as meaningful.¹⁵⁵

Let us call this the "soft-core" frequentist view. The soft-core frequentist thinks that probability numbers *are* meaningfully attachable to singular propositions such as "P will die" or "P will die as a result of E" or "P will die earlier than he would have had E not occurred." Consider the proposition "P will die earlier than he would have had E not occurred." The probability of that proposition would be seen, by the soft-core frequentist, as *the frequency with which people like P die prematurely when events like E occur*.

To be a bit more precise, we can imagine a large or infinite reference class of events $\{E_1, E_2 \dots E_n\}$ that includes the particular event of interest E.¹⁵⁶ We can also imagine a large or infinite class of persons $\{P_1, P_2 \dots P_n\}$, where each P_i is alive

154. REICHENBACH, *supra* note 79, at 374.

155. *Id.* at 376-77.

156. As I have mentioned, the events or, more generally, the objects that the frequentist uses to determine relative frequencies and therewith probabilities might be conceptualized in different ways: as actual objects, hypothetical objects, or mathematical constructs that represent actual or hypothetical objects. See *supra* text accompanying notes 77-80. Thus the class of events $\{E_1 \dots E_n\}$ and matching class of persons $\{P_1 \dots P_n\}$ employed to determine the frequentist probability that E results in P's premature death might be classes of actual or hypothetical events and persons or of mathematical constructs that represent them. As far as I can tell, nothing in my argument against the harmfulness of frequentist probability depends on how the classes are specified.

when the matching event E_i occurs. In each case we could, in principle, determine whether the given P_i dies prematurely, i.e., dies earlier than he would have had E_i not occurred. The relative frequency of premature death is the number of premature deaths, divided by the total number of events in the reference class $\{E_1 \dots E_n\}$ —or the limit of the number of premature deaths divided by the total number of events, in the case of an infinite reference class $\{E_1 \dots E_n \dots\}$.¹⁵⁷

A simple example might help clarify how this soft-core frequentist strategy for ascribing death probabilities would work. A particular person P eats a food item containing the food additive sodium nitrite. This event E (P 's ingestion of the food item) is alleged to impose a high probability of death on P . A soft-core frequentist would subsume E within some broader class. For example, E falls within the large class $\{E_i\}$ including every ingestion within actual human history of a sodium nitrite-containing food item. For each E_i within this class, the matching P_i might be defined as the person who ingested the food item. For each such E_i , we can ask whether the person P_i who ingested the item died "prematurely," i.e., earlier than he would have had the ingestion not occurred. Imagine that $\{E_i\}$ includes 10 million events, and that in 20,000 cases the ingester P_i died earlier than he would have had the ingestion not occurred. Thus, the relative frequency of premature death within this reference class $\{E_i\}$ is $20,000/10 \text{ million} = 1/500$. This relative frequency number, $1/500$, might be taken as the probability of the proposition " P will die earlier than he would have had E not occurred."

In sum, the soft-core frequentist would characterize the difference between the safer world O_s and the riskier world O_r as follows:

Riskier and Safer Outcomes: A Frequentist Account

- O_s is a world in which E does not occur.
- O_r is a world in which (1) E occurs and (2) the frequency with which people like P , exposed to events like E , die prematurely, is high.
- O_r and O_s are otherwise identical with respect to P 's welfare.

157. To put this another way, given our particular event E and person P , we are subsuming the ordered pair (E, P) in a class of ordered pairs $\{(E_1, P_1), (E_2, P_2), \dots (E_n, P_n)\}$. I am indebted to Paul Edelman for a helpful discussion on this issue.

Yet there remains a crucial incompleteness in this characterization of O_s and O_R . Given a particular event E , there are many different reference classes in which E might be subsumed. E has multiple characteristics; each such characteristic, and each conjunction of E 's characteristics, defines a different reference class that E is "similar to," in a certain way. For example, a particular release of uranium by a nuclear plant upriver from a major city falls in the reference class of "releases of toxic material," "releases of radioactive material," "releases of uranium," "releases of uranium by nuclear plants proximate to large cities," "releases of toxic material by facilities proximate to large cities," and so on. An analogous point can be made about the person P exposed to the "risk" of death. That risk depends upon the frequency of premature death within a class $\{P_i\}$ including P , but of course there are many (indeed, an infinite number) of such classes.

[Consider] the probability of a particular man aged 40 living to be 41. Intuitively the probability will vary depending on whether we regard the individual merely as a man or more particularly as an Englishman; for the life expectancy of Englishmen is higher than that of mankind as a whole. Similarly, the probability will alter depending on whether we regard the individual as an Englishman aged 40 or as an Englishman aged 40 who smokes two packets of cigarettes a day, and so on.¹⁵⁸

Because P and E can be characterized in different ways, there are multiple different probabilities that E will result in P 's premature death—corresponding to the multiple reference classes of events $\{E_i\}^A, \{E_i\}^B \dots$ and matching classes of persons $\{P_i\}^A, \{P_i\}^B \dots$ that contain E and P , with each such pair of classes generating a different relative frequency of premature death. This point might be seen to have quite immediate and dire implications for the claim that risk, in the relative-frequency sense, is harmful. After all, if the frequentist probability that E causes P 's premature death is high, relative to one description of E and P , but low, relative to another

158. GILLIES, *supra* note 14, at 119. This crucial problem for frequentist theory—the problem of the reference class and the closely related problem of singular probabilities—has been much discussed in the literature on probability theory. See, e.g., *id.* at 119-25; HOWSON & URBACH, *supra* note 73, at 338-42; KYBURG, *supra* note 14, at 47-51; WESLEY C. SALMON, *THE FOUNDATIONS OF SCIENTIFIC INFERENCE* 90-94 (1966); WEATHERFORD, *supra* note 14, at 165-67; Howson, *supra* note 14, at 21-23. The multiplicity of reference classes subsuming a particular event is crucial to Stephen Perry's argument that risk in the frequentist sense is not harmful. See Perry, *supra* note 25, at 330-39.

description, how can it be said that E's occurrence has really harmed P? I do not believe, however, that the dependence of frequentist probabilities on how we describe particular events, etc., *necessarily* makes these probabilities irrelevant to our welfare. The soft-core frequentist can argue as follows: There is a *welfare-relevant* way to characterize the event of interest E and the person of interest P. When we characterize E and P in the welfare-relevant way, we come up with classes $\{E_i\}^{\text{WELFARE}}$ and $\{P_i\}^{\text{WELFARE}}$; the frequency with persons in $\{P_i\}^{\text{WELFARE}}$ die prematurely as a result of events in $\{E_i\}^{\text{WELFARE}}$ might be high or low; if it is high, we can say that the welfare-relevant characterization of P and E generates a high probability of premature death; and in this case P has indeed been harmed, regardless of whether a different (welfare-irrelevant) characterization of P and E generates a different probability of premature death.

Does this line of argument work? Is there indeed a welfare-relevant way to characterize events and persons such that a high frequency of premature death, relative to this characterization scheme, is indeed harmful? I suggest that the answer is "no," and will try to show why by examining a number of different ways in which events and persons can be assigned to reference classes.¹⁵⁹

At this juncture, I should register both my great debt to, but also partial disagreement with, Stephen Perry's scholarship on the harmfulness of risk. In a seminal article, *Risk, Harm and Responsibility*, Perry rigorously analyzes the harmfulness of risk for purposes of tort law.¹⁶⁰ Perry distinguishes between Bayesian (what he calls "subjectivist") and frequentist conceptions of risk;¹⁶¹ focuses on frequentist risk (on the grounds that Bayesian risk seems unlikely to have moral significance);¹⁶² explains how ascriptions of frequentist risk are relative to reference classes;¹⁶³ and concludes that frequentist risk is not harmful, at least in a deterministic world.¹⁶⁴ The core of Perry's argument on this last point, I take it, is that (1) in a deterministic world a person who is not physically injured by some action is at zero frequentist risk of injury, relative to

159. See *infra* Part III.B–D.

160. Perry, *supra* note 25, at 322–39.

161. *Id.* at 322–27.

162. *Id.* at 327–28.

163. *Id.* at 333–35.

164. *Id.* at 334–36.

what I will term a “physical” reference class, namely one specifying all the causally relevant features of the action and person; and therefore (2) the fact that the person has nonzero frequentist of injury relative to a less finely specified reference class cannot itself be a welfare setback for her.

[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 . . . The basis for saying that a person is at risk of suffering a certain type of physical harm is that he or she belongs to a particular reference class with which is associated a known objective [i.e., frequentist] probability of harm of that type. But there is nothing magical about the particular reference class selected: generally it will simply be the narrowest class, given the current state of our knowledge, for which we are able to determine with some degree of accuracy the relative frequency of the type of harm in question. If we were in a position to describe the objective probabilities associated with a still narrower reference class to which we knew the person in question also belonged, presumably we would do so. Moreover we know, given the assumption of determinism, that it is in principle possible to partition any such reference class into two subclasses, one for which the probability of physical harm is one and one for which the probability is zero. Thus, there is simply no ground for the claim that the objective [frequentist] probability of harm, associated with the particular reference class that we are currently in a position to describe, itself constitutes a distinct form of damage.¹⁶⁵

I agree with Perry’s conclusion, here, but think the argument is too quick. To determine whether someone is harmed by a high frequentist risk of injury or premature death, we need to consider different plausible accounts of well-being and different possible specifications of the reference class generating the high risk. Perry does not do this. His assumption, as suggested by the above quotation, seems to be that our interest in frequentist risk is only epistemic—as an indicator of whether the person at risk will be injured or not. However, frequentist risk relative to some reference class (for example, a conventional class) might be welfare relevant on some welfare theory (for example, a theory that emphasizes Value and the role of conventional ascriptions in fixing Value) independent of its epistemic function. Further, the fact that frequentist risk is epistemically irrelevant to perfectly informed agents in a deterministic world does not show that it lacks epistemic and thereby welfare relevance for imperfectly informed agents. Why look, in determining welfare, to what

165. *Id.* at 336.

the person whose welfare is at issue would believe or prefer under hypothetical conditions of full information, rather than to what she actually believes and therefore prefers?

In point of fact, I *do* think that full-information preferences are an important part of the welfare analysis, and in part for that reason ultimately concur in Perry's conclusion. Yet, I believe, we cannot establish the harmlessness of risk without explicitly considering the whole range of plausible welfare theories (more precisely, the whole range of plausible elements of welfare theories, namely Experience, Preference, Value and Integration.) So I will, in my own fashion, reanalyze the question that he addresses in *Risk, Harm, and Responsibility*. My analysis does, at important junctures, overlap with his, and at every such juncture I have (I hope) noted the debt.

B. STATISTICAL REFERENCE CLASSES

Reichenbach suggests the following principle for subsuming a particular event within a reference class: choose "the narrowest class for which reliable statistics can be compiled."¹⁶⁶

Reichenbach's principle, and variations thereof, is a standard proposal within the frequentist literature.¹⁶⁷ This proposal might be used to generate a probability of premature death, as follows. Given a particular event E and a particular person P, we consider progressively more detailed (joint) characterizations of E and P. Imagine that E is the ingestion, at 10 p.m., on a hot summer day, of an all-beef hot dog containing sodium nitrite, by a thirty-five-year old Caucasian man P weighing 200 pounds, with a family history of liver cancer, and who tends to eat red meat and other high-fat, high-

166. REICHENBACH, *supra* note 79, at 374. For a similar suggestion, see A.J. Ayer, *Two Notes on Probability*, in *THE CONCEPT OF A PERSON AND OTHER ESSAYS* 202 (1963). Reichenbach's approach need not generate a determinate probability, since there may be multiple classes, none narrower than any of the others, for which reliable statistics are available and which include a given event. See GILLIES, *supra* note 14, at 121. In my analysis below, I ignore this difficulty, assume that a determinate probability is generated, and show why—nonetheless—frequentist risk relative to a Reichenbachian reference class is not harmful.

167. See, e.g., SALMON, *supra* note 158, at 91 (modifying Reichenbach's proposal by stipulating that the reference class should be "homogeneous"—not further specifiable in a way that changes the frequency of the relevant attribute—and that the broadest such class should be used to determine a single-case probability).

cholesterol foods. We can characterize E as "the ingestion of a sodium nitrite-containing food item by a man," and P as "the man who ingests the item." We can, with greater specificity, characterize E as "the ingestion of a sodium nitrite-containing food item by a man with a family history of cancer" and P as "the man, with a family history of cancer, who ingests the item." With even greater specificity, we can describe E as "the ingestion of a sodium nitrite-containing food item by a man P with a family history of cancer and a high-fat diet" and P as "the man, with a family history of cancer and a high-fat diet, who ingests the item." For some characterizations of E and P, reliable statistics will presumably be available at or around the time of E, while for other characterizations such information will be unavailable.¹⁶⁸ For example, there might be reliable statistics about the frequency with which the ingestion of a sodium nitrite-containing food item results in premature death, but not about the frequency with which the ingestion of sodium nitrite-containing hot dogs by large Caucasians with a history of liver cancer results in premature death. Choose the most specific characterization of E and P for which reliable statistics are available. This characterization generates a reference class for E and P, which I will call the "statistical" class and abbreviate $\{E_i\}^{\text{STAT}}$ and $\{P_i\}^{\text{STAT}}$. If the most specific such characterization of E and P is, say, "the ingestion of a sodium nitrite-containing food item by a man with a history of liver cancer," then $\{E_i\}^{\text{STAT}}$ is the set of all ingestions¹⁶⁹ of sodium nitrite-containing food items by men with a history of liver cancer, and $\{P_i\}^{\text{STAT}}$ is the set of men who ingest those items.

168. See REICHENBACH, *supra* note 79, at 377-78. Reichenbach states, "We can ask only for the best reference class available, the reference class that, on the basis of our present knowledge, will lead to the greatest number of successful predictions, whether they concern hits of bombs, cases of disease, or political events. If no statistics are available for the common class A.C., we shall base our probability calculations on the reference class A, and must renounce improvement in the success ratio that might result from the use of the reference class A.C."

Id.

169. More precisely, $\{E_i\}^{\text{STAT}}$ could be: (1) the finite set of all such ingestions that have actually occurred; (2) the infinite set of hypothetical ingestions that would occur, were the right sort of event-generating scenario (here, having a man with a history of liver cancer eat a sodium nitrite-containing food item) to occur repeatedly ad infinitum; or (3) a set of mathematical constructs representing the first or the second. See *supra* notes 77-80 (discussing whether the relevant objects, for the frequentist, are actual, hypothetical, or mathematical).

These classes can then be used to calculate what I will call the "statistical" probability that our particular event E will result in the premature death of our particular person P , namely, the frequency with which persons within $\{P_i\}^{\text{STAT}}$ die earlier than they would have had the matching events within $\{E_i\}^{\text{STAT}}$ not occurred.

Consider, now, our two outcomes, safer O_s and riskier O_r , on the assumption that the high probability of death in O_r is a *statistical* probability.

Riskier and Safer Outcomes: A Frequentist Account, with Statistical Probabilities

- O_s is a world in which E does not occur.
- O_r is a world in which (1) E occurs and (2) the frequency with which people statistically similar to P , exposed to events statistically similar to E , die prematurely, is high. In other words, the relative frequency of premature death by persons within $\{P_i\}^{\text{STAT}}$, as a result of events within $\{E_i\}^{\text{STAT}}$ (where $\{P_i\}^{\text{STAT}}$ and $\{E_i\}^{\text{STAT}}$ are the narrowest classes for which reliable statistics can be compiled) is high.
- O_r and O_s are otherwise identical with respect to P 's welfare.

Is O_r worse for P than O_s ? I think not.

First, O_r need not be worse than O_s with respect to the Experience component of welfare. An event E creates a high statistical probability of P 's premature death just in case the most specific characterization of E and P (given currently available information) generates a high frequency of premature death. The occurrence of such an event does not entail any change in P 's mental states.¹⁷⁰ E itself could be an event external to P 's mind (either an event that occurs in the outside world, or an event that involves bodily but not mental changes in P). And the statistical classes $\{E_i\}^{\text{STAT}}$ and $\{P_i\}^{\text{STAT}}$ are the narrowest classes for which reliable statistics can be compiled; this may not be information about which P herself is aware. For example, a dangerous event such as a toxic release, unsafe behavior by co-workers, or a change in supervisory practices, could occur in P 's workplace without P 's knowledge or awareness, and without her knowledge or awareness of the frequency with which statistically similar events result in

170. Cf. Perry, *supra* note 25, at 338 (arguing that frequentist risk is not harmful in itself, but could be once discovered by the person at risk).

premature death to statistically similar persons. In such cases, the claim that O_R is worse for P than O_S runs afoul of the standard Experientialist principle that welfare supervenes on mental states.¹⁷¹

To be sure, one can imagine unusual cases where O_R is statistically riskier for P than O_S , where the two outcomes are otherwise identical for P's welfare, but where O_R and O_S still differ with respect to P's experience. This will occur when the event of interest, E, itself involves a change in P's mental life. For example, E might be a shockingly scary nightmare, and it might emerge that the statistical frequency with which shockingly scary nightmares result in the premature death (by heart attack, say) of the persons experiencing such nightmares is high. Even here, however, an Experientialist would deny that the statistical probability of premature death is itself harmful. What makes O_R worse for P, if anything, is the scary nightmare—not the statistical probability of death thus produced, i.e., the fact that statistically similar nightmares often result in premature death to persons statistically similar to P. That fact, itself, need not enter P's experience. Further, if we imagine a case in which P does become aware of the statistics about the risky event, then we have added something to O_R *other than* the high statistical probability of P's premature death—namely, P's awareness of that probability—and it is this additional feature, not the probability itself, that makes O_R worse for P.

Second, the claim that statistical probabilities are harmful runs afoul of the Preference component of welfare. Why? After

171. Conceivably, the Experience component of welfare might be specified in a form weaker than a supervenience requirement. See *supra* text accompanying note 61 (noting this possibility). For example, the Experientialist might stipulate that two outcomes can differ for P's welfare only if they differ in some way accessible to P's experience. The Experience requirement, thus weakened, is consistent with the harmfulness of statistical and more generally frequentist risk. But is the requirement, in this form, a plausible part of a welfare account? Why hew to the intuition that welfare has something to do with our mental states and yet flesh it out in such a minimal way? I do not pursue these issues here because even if the view that statistical and, more generally, frequentist risk is harmful is consistent with the best construal of the Experience requirement, it is problematic on other grounds. See *infra* text accompanying notes 172-83, 186-90, 195-200 (arguing that frequentist risk relative to statistical, physical, and conventional reference classes does not satisfy the Value, Preference, and/or Integration elements of welfare, and raising further objections to the claim that frequentist risk in the physical or conventional sense is harmful).

all, people can prefer or disprefer anything. P might prefer the nonoccurrence of events which impose a high statistical probability of premature death on him. If P has this preference, then P prefers O_s to O_R , and won't the Preferentialist therefore conclude that O_s is better for P's welfare than O_R ? The problem is that Preferentialists, quite typically (and quite plausibly), stipulate that welfare-constitutive preferences must be *fully informed*.¹⁷² Yet it is very hard to see why P, under conditions of full information, would prefer O_s to O_R .

Consider once more the case where E is the ingestion, at 10 p.m., on a hot summer day, of an all-beef hot dog containing sodium nitrite by a thirty-five-year old Caucasian man P weighing 200 pounds, with a family history of liver cancer, and who tends to eat red meat and other high-fat, high-cholesterol foods. The narrowest reference class $\{E_i\}^{STAT}$ for which reliable statistics are available might be all "ingestions of a sodium nitrite-containing food item by a man with a history of liver cancer." Yet P, under full information, would also know the relative frequency of premature death given a much more specific reference class including all the causally relevant features of E—namely, "ingestions of sodium nitrite-containing food items by Caucasian men who weigh 200 pounds, who subsist on high-fat diets, who have a family history of liver cancer, and who" I will use the term "physical probability" to mean frequencies relative to classes of events and persons, $\{E_i\}^{PHYS}$ and $\{P_i\}^{PHYS}$, that share *all* the causally relevant features of E and P—not merely the features about which good statistical data is actually available.¹⁷³ I submit that P's preferences, under conditions of full information, would be driven by the physical probability of death and not the statistical probability.

Statistical and physical probabilities are gauges or indicators of P's security; they indicate whether E is a substantial or insubstantial threat to P's continued existence. It is plausible that P (both in his actual state and his fully informed state) would prefer his continued existence not to be *insecure*. Thus P would plausibly prefer that the best available gauge not indicate a high level of insecurity. But the physical

172. See sources cited *supra* note 45.

173. For rigorous discussions of what I am calling physical probability, see GILLIES, *supra* note 14, at 125-29; HOWSON & URBACH, *supra* note 73, at 338-42.

probabilities are a better gauge of P's security than the statistical probabilities. The statistical reference class $\{E_i\}^{\text{STAT}}$ includes events that are roughly similar, in their causal features, to E; the physical reference class $\{E_i\}^{\text{PHYS}}$ includes events that are perfectly similar, in their causal features, to E.¹⁷⁴ Any deviation between the statistical and the physical probabilities reflects a lacuna in the stock of information used to generate the statistical probabilities. Society, at or around the time of E, simply lacks reliable relative-frequency data about the correlation between some of E's casually relevant features, and premature death; and thus those relevant-but-poorly-understood features are ignored in generating the statistical reference class $\{E_i\}^{\text{STAT}}$.¹⁷⁵

The third plausible component of well-being, along with Experience and Preference, is the Value component. Here, too, serious difficulties arise for the claim that a high statistical probability of premature death is harmful. How can I say anything definitive about the comparative *value* of O_r and O_s ? After all, various philosophers of well-being including Martha Nussbaum, John Finnis, George Sher, and others have proposed diverse lists of putative welfare values and have grounded these lists in different foundational considerations (the human essence, near-universal and -unavoidable goals, shared understandings of the good life).¹⁷⁶

Notwithstanding this continuing scholarly disagreement about the content of welfare-values, and about their basis, it is

174. Reichenbach himself was quite straightforward about the fact that statistical probabilities reflect our imperfect information. See REICHENBACH, *supra* note 79, at 375 ("[T]he probability of a single case [generated using the narrowest class for which reliable statistics can be compiled] is . . . dependent on our state of knowledge . . ."). Perry generally argues that frequentist probabilities, other than physical probabilities, are not welfare relevant because our interest in such probabilities does not persist as our information increases. See Perry, *supra* note 25, at 334-37.

175. I am conceding here that a high physical probability of premature death does satisfy the Preference component of well-being—more precisely, that two outcomes differentiated by a high physical death probability are (consistent with Preference) different for the subject's welfare. That concession may be unwarranted. Arguably the preferences over outcomes of a fully informed agent, concerned about her own physical security, would be driven neither by the statistical probability of premature death, nor by the physical probability, but just by the occurrence of premature death itself. See *infra* note 185. But this point, if true, hardly rehabilitates the case for the welfare relevance of statistical probability.

176. See *supra* text accompanying notes 48-56 (discussing these proponents of a Value-based theory of welfare).

(I think) possible to make a claim that all value theorists would assent to, namely, that welfare values are *objective*.¹⁷⁷ Values do not merely reflect what P himself prefers or wants; they have a more robust normative status than that.¹⁷⁸ In turn, it is quite plausible that there is a deep connection between objective values and the preferences of idealized agents. Plausibly, one outcome is objectively more valuable than another only if all persons (or perhaps all persons within a certain type of group, for example a society), under ideal conditions, would prefer the first outcome. The philosopher Michael Smith, among others, has argued persuasively for this account of "objective value."

[W]hat it is desirable for us to do is what we would desire that we do if we were fully rational. In other words . . . what it is desirable for us to do in certain circumstances—let's call these circumstances the 'evaluated possible world'—is what we, not as we actually are, but as we would be in a possible world in which we are fully rational—let's call this the 'evaluating possible world'—would want ourselves to do in those circumstances. . . .

Typically, of course, the evaluated world will be the actual world. Thus, what it is desirable for us to do in our actual circumstances is what our more rational selves, looking down on ourselves as we actually are from their more privileged position, would want us to do in our actual circumstances.¹⁷⁹

Assume Smith's account of objective value is correct. Then O_s is better for P than O_r only if all persons, fully informed (or at least the appropriate group of persons, fully informed) would

177. See, e.g., HURKA, *supra* note 49, at 5 (describing his account as "objective"); PARFIT, *supra* note 18, at 499 (presenting a list of plausible welfare goods under the rubric of an "Objective List Theory" of well-being); SHER, *supra* note 49, at 229 (describing his account as "objective"); SUMNER, *supra* note 8, at 45-80 (discussing and criticizing objective theories of well-being, including Finnis's and Nussbaum's). Griffin declines to describe his list of welfare values as "objective," meaning that these obtain independent of desires, but he would, I think, agree that the values do not reduce to the subject's own desires and, in that sense, are objective. See GRIFFIN, *supra* note 23, at 29, 35-36.

178. On the nature of value, see generally GRIFFIN, *supra* note 23, at 52-67; JOEL J. KUPPERMAN, VALUE . . . AND WHAT FOLLOWS 1-83 (1999); NOAH M. LEMOS, INTRINSIC VALUE: CONCEPT AND WARRANT 3-100 (1994); RAMON M. LEMOS, THE NATURE OF VALUE: AXIOLOGICAL INVESTIGATIONS 1-71 (1995).

179. MICHAEL SMITH, THE MORAL PROBLEM 151-52 (1994); cf. David Lewis, *Dispositional Theories of Value*, in PROCEEDINGS OF THE ARISTOTELIAN SOCIETY 63 Supplemental Vol. 113 (1989) (defining value in terms of second-order desires); Connie S. Rosati, *Internalism and the Good for a Person*, 106 ETHICS 297, 307 (1996) (defining the good for a person P as roughly what P hypothetically would care about under idealized conditions that P actually regards as welfare relevant).

prefer to be P in O_s rather than P in O_r . But I have already argued that P, if fully informed, would focus on the physical probability of his premature death rather than the statistical probability. The physical probability of P's premature death is a better indicator of P's security than the statistical probability of his premature death. For just the same reason, other fully informed agents (if contemplating the choice between being P in O_r and being P in O_s) would care about the physical and not the statistical probability of premature death. Therefore, the fact that O_r is statistically riskier for P than O_s does *not* make O_r less *valuable* for P.¹⁸⁰

We come now to the fourth *and* final component of welfare: Integration. Integration demands that welfare changes occur "within" the subject's life.¹⁸¹ They must be changes in the world that make it worse "for the subject," not worse (merely) for someone else or in some detached, impersonal way. The Integration requirement poses real difficulties for the claim that probabilities, in the frequentist sense, can be harmful or beneficial.

Consider O_s and O_r . What differentiates the two outcomes is the following fact: in O_r , an event E occurs such that a high proportion of events "similar to" E result in premature death to people "similar to" P. The notion of "similarity," here, may be cashed in statistical terms, in physical terms, or in some other terms. *However* we generate our classes $\{E_i\}$ and $\{P_i\}$, the essential point remains: A high "probability" of premature death for P, in the frequentist sense, means a high frequency of premature death within some *group of persons*, $\{P_i\}$, that includes P. But this "group fact" does not change P's own life—or so it seems.¹⁸² Although a group including P may have a high frequency of some harmful attribute, what matters for P's well-being is whether she herself has the attribute. If P has the attribute, then she is harmed, and the fact that only a few other members of the group share the attribute does not lessen

180. *But see* Finkelstein, *supra* note 25 (manuscript at 13) (arguing that the risk of harm in the frequentist sense is itself a harm, in the sense of being a setback to an "interest," where "an agent's interests [are] an expanded set of goods derivable in some way from the preferences the agent actually has").

181. *See supra* text accompanying note 61.

182. *See* Perry, *supra* note 25, at 335 ("[Frequentist probabilities] have been defined with respect to a certain *class* of persons But it is individuals who are said to suffer risk damage, not classes of persons, and therein lies the fallacy of the claim that risk damage constitutes injury in its own right.").

the harm. Likewise, if P lacks the attribute, then she is unharmed, and the fact that many other members of the group have the attribute does not amount to a harm for P herself. The high relatively frequency of premature death, or some other negative attribute, within a group of persons $\{P_i\}$ that happens to include P, is a feature of an outcome that (seemingly) is too remote from P's own life to satisfy Integration.

P is, let us imagine, a redhead; the first letter of his last name is "S"; and he owns a truck. Imagine a world in which many redheaded truck owners whose last names begin with "T" suffer from terrible cluster headaches. Imagine another world in which few or no redheaded truck owners whose last names begin with "T" incur such headaches. In short, in the first world there is a high relative frequency of cluster headaches among a group of persons excluding P, and in the second world there is a low relative frequency of cluster headaches among that group. Does this difference between the two outcomes change P's well-being? No—at least not if the Integration requirement has any bite.

Now change the example slightly so that P's last name begins with "T." Thus, in the first world there is a high relative frequency of cluster headaches among a group of persons *including* P, and in the second world there is a low relative frequency of cluster headaches among that group. Does the sheer fact of P's inclusion in the group satisfy the Integration requirement, so that we now can say the second world is better *for P* than the first world? All that P shares with his "fellow" group members is truck ownership, red hair, and a last name beginning with T. In the first example, an increase in the rate of cluster headaches did not suffice to change P's own life. Why, in the second example, should an increase in the rate of cluster headaches among a group of persons tenuously linked to P (redhaired, T-named truck owners) suffice to do so?

The lesson of these examples, I submit, is that a high relative frequency of premature death or some other negative attribute within some group $\{P_i\}$ does not satisfy the Integration requirement, with respect to P, merely because P is included within the $\{P_i\}$. On the other hand, we can imagine much stronger links between P and the $\{P_i\}$ than mere inclusion.¹⁸³ For example, imagine that P *identifies with* the

183. See Owen M. Fiss, *Groups and the Equal Protection Clause*, 5 PHIL. &

$\{P_i\}$. P is a male Sephardic Jew; a high proportion of male Sephardic Jews suffer cluster headaches; and P self-identifies as a male, Sephardic Jew. In this case, the Integration requirement arguably permits the high rate of cluster headaches among the $\{P_i\}$ to be counted as a harm to P himself. The “group fact” about male Sephardic Jews is also, arguably, a fact about P’s own life. Even if P avoids the headache, he himself is arguably worse off in a world in which many of his fellow male Sephardic Jews—a group that not only includes P, but is defined by characteristics that are highly salient to P—suffer the headache.

Now back to statistical probabilities. O_R is a world in which the relative frequency of premature death, among the class of persons $\{P_i\}^{STAT}$ who are “statistically” similar to P, is high. What P shares with the other members of $\{P_i\}^{STAT}$ is having a particular set of characteristics whose correlation with premature death is relatively well documented. P need not identify with $\{P_i\}^{STAT}$. Indeed, where an event generates a high statistical probability of premature death for P, P need not (and, in the normal case, will not) have any idea what the defining characteristics of the statistical similarity class are. For that matter, P need not (and, in the normal case, will not) have any idea what a statistical similarity class is. The linkage between P and the $\{P_i\}^{STAT}$ is not personal (P’s own identification), nor is it social (how ordinary members of P’s society would characterize P), nor is it physical (the group sharing all of P’s causally relevant characteristics). The linkage is, rather, *forensic*— $\{P_i\}^{STAT}$ is the grouping that would be salient to qualified statisticians, operating at a particular point in time, with a particular information base—and this linkage is probably not strong enough to satisfy any non-minimal Integration requirement.

C. PHYSICAL REFERENCE CLASSES

Is a high physical probability of premature death a welfare setback?¹⁸⁴ Consider our two outcomes, O_S and O_R , with the

PUB. AFF. 107, 148-49 (1976) (defining “social group” and arguing that harms to social groups are harms to their members).

184. Perry suggests that a high physical probability of premature death, in an indeterministic universe, could well be a welfare setback. See Perry, *supra* note 25, at 336-37. I disagree, for reasons elaborated in the text immediately below. See *infra* text accompanying notes 185-88. In brief, the problems of Experience and Integration do not disappear with the shift from statistical to

difference between the two now specified in terms of physical (rather than statistical) probability.

Riskier and Safer Outcomes: A Frequentist Account, with Physical Probabilities

- O_s is a world in which E does not occur.
- O_R is a world in which (1) E occurs, and (2) the physical probability that E results in P's premature death is high. In other words, the relative frequency of premature death within $\{P_i\}^{PHYS}$, as a result of events within $\{E_i\}^{PHYS}$ (where $\{P_i\}^{PHYS}$ is the class of persons sharing all of P's causally relevant characteristics, and $\{E_i\}^{PHYS}$ the class of events sharing all of E's causally relevant characteristics) is high.
- O_R and O_s are otherwise identical with respect to P's welfare.

In my discussion above, I suggested that physical probabilities of premature death *are* plausibly dispreferred (even under full information) and that these physical probabilities are plausibly value reducing. In short, O_R is plausibly worse for P than O_s , with respect to the Preference and Value components of well-being.¹⁸⁵ Should we conclude that risk, in the physical-probability sense, is indeed a welfare setback?

This conclusion would, I think, be too hasty. First, the physical probability of premature death—like the statistical probability of premature death—is not a setback with respect to the Experience component of well-being. Unless E is itself a

physical probabilities.

185. It could be argued that O_R is not worse than O_s with respect to Preference, to the extent that the preferences must be fully informed, and therefore not worse with respect to Value either. It seems quite plausible that a fully informed agent would ex ante disprefer the occurrence of an event involving a high physical probability of her premature death, as compared to nonoccurrence—but this is not the same as saying that she would disprefer an outcome characterized by a high physical probability of premature death, as compared to an otherwise-identical outcome characterized by a lower physical probability. Insofar as she is motivated by physical security, her preference against the occurrence of such an event is purely instrumental, not intrinsic, and thus outcomes differentiated solely by the occurrence of that event are just the same for her welfare. A possible response here, I suppose, is that P's preferences over outcomes are determined by looking to her fully-informed ex ante preferences. This response raises deep issues about the nature and determination of preferences over outcomes—over whole possible worlds—which I will not pursue, since even if O_R and O_s are different with respect to Preference and Value, the putative welfare relevance of physical probability encounters further difficulties discussed in this section.

mental event (say, a shockingly scary nightmare), P's mental states will not be different in O_S and O_R . And in the case where E is itself a mental event, it will be that experiential feature of O_R , not the physical probability of premature death associated with E, that (according to the Experientialist) makes O_R worse for P than O_S .

Second, it is doubtful, or at least very unclear, whether O_R is worse than O_S with respect to the Integration component of well-being. In O_R , an event E occurs such that physically similar events (events within $\{E_i\}^{PHYS}$) frequently cause premature death to persons physically similar to P (persons within $\{P_i\}^{PHYS}$). Is this "group fact" about O_R sufficient to make O_R worse for P? In discussing statistical probabilities, I suggested that a high frequency of harm within some group $\{P_i\}$ could not be taken as a harm to P—given the Integration requirement—merely because $\{P_i\}$ includes P. I conceded, however, that a high frequency of harm within some group $\{P_i\}$ could be taken as a harm to P—consistent with that requirement—if P identifies with the $\{P_i\}$.¹⁸⁶ Clearly, P need not identify with the class of persons physically similar to him. A high physical probability of premature death does not entail P's awareness of that probability; likewise, it does not entail P's awareness of the causally relevant characteristics that he shares with the other members of $\{P_i\}^{PHYS}$, or his identification with that group.

Perhaps my demand that P identify with the $\{P_i\}^{PHYS}$ is too robust. Could there be some linkage between a person and a group suffering a high frequency of harm, weaker than identification, which suffices for that group fact to be counted as a welfare loss for the person himself? One such linkage might be social. If P is conventionally or socially linked to the group, then arguably a high frequency of harm within the group is sufficiently integrated within P's life to be welfare reducing for him, even if P himself does not identify with the group.¹⁸⁷ Does the linkage of physical similarity between P and a group also function to integrate the group's harm into P's own

186. See *supra* text accompanying note 183.

187. Cf. Fiss, *supra* note 183, at 148. Fiss points to individual identification in determining an individual's membership in a "social group" but also suggests that social identification plays a role: "Blacks are viewed as a group; they view themselves as a group; their identity is in large part determined by membership in the group; their social status is linked to the status of the group; and much of our action, institutional and personal, is based on these perspectives." *Id.*

life?

I doubt it. Imagine that P has two genetically identical siblings, Q and R, who were separated from P at birth and live in a different country. P considers Q and R to be virtual strangers to him; they mean little more to him than a stranger he might meet on the train.¹⁸⁸ In short, P does not identify with Q and R; nor do the members of P's society identify him with Q and R, given their residence elsewhere. We might imagine that P prefers that Q and R not suffer harm, just as we might imagine that P prefers the stranger he meets on the train not to suffer harm. Even if the stranger on the train does suffer harm, and P learns of that, P himself will be no worse off (if welfare includes an Integration requirement). Imagine, instead, that Q and R are harmed, and P learns of that. What differentiates Q and R from the stranger on the train is their genetic similarity to P; but neither P nor others in P's society see this similarity as creating a stronger bond between P and Q/R, than between P and the stranger. If so, I suggest, harms to Q and R are just as remote from P's own life as harms to the stranger. To be sure, physical similarity and genetic similarity are conceptually different, but it is very hard to see why physical similarity absent individual identification or social linkage should satisfy Integration, while genetic similarity absent individual identification or social linkage does not.

Third and finally, even if the objections from Experience and Integration are answered, and a high physical probability of premature death is properly viewed as a welfare setback, it is unclear how much relevance this type of harm would have for risk regulators. Why? Consider the following crucial feature of physical probabilities. *Physical probabilities can lie between 0 and 1 only if causal laws are indeterministic.*¹⁸⁹ In a world with deterministic causal laws, where every occurrence is in principle predictable, with certainty, from preexisting conditions, the physical probability of an event having some attribute is either 1 (if the event ends up having the attribute) or 0 (if it does not).

To see this point, imagine that the pitch of a canary's song is completely determined by the canary's weight, length, and

188. See *supra* text accompanying note 44 (discussing the "stranger on the train" example, which illustrates that people can have preferences for outcomes detached from their own welfare, viz., the stranger's success).

189. See sources cited *supra* note 173 (discussing the link between physical probabilities and indeterminism).

eye color. For a given combination of weight, length, and eye color, the casual laws necessitate that the canary's pitch will either be high or low. What is the *physical probability* that the song a particular canary, Candy, sings tomorrow will have a high pitch? Note that the statistical probability of Candy's song having a high pitch might lie between 0 and 1. It might be, say, 0.3. If we only have reliable statistics linking weight and pitch, and if the frequency with which canaries statistically similar to Candy—canaries weighing the same amount—have a high-pitched song is 0.3, then the statistical probability of Candy's song having a high pitch is 0.3. However, the physical probability of Candy's song having a high pitch is necessarily either 1 (just in case the song will, in fact, turn out to be high-pitched) or 0 (just in case the song will, in fact, turn out to be low-pitched). Assume the song will have a high pitch. Then Candy has a combination of weight, length, and eye-color characteristics necessitating a high-pitched song. The class of canaries physically identical to Candy is the class of canaries that are the same weight, length, and eye color as Candy. And the relative frequency of high-pitched songs within this class is 1 (since the causal laws necessitate, for each and every member of this group, that its song be high pitched). Similarly, if Candy's song will turn out to be low pitched, the relative frequency of low-pitched songs within the class of physically similar canaries is 1, and the relative frequency of high-pitched songs within that class is 0.

Modern physics tells us that determinism does not obtain at the level of quantum mechanics. But it could well be the case that causal laws at the level of human physiology—causal laws linking the large-scale events of interest to risk regulators, with premature human death—are deterministic. As Stephen Perry explains,

It seems clear . . . that true indeterminism exists at the level of sub-atomic particles, but there is some reason to think that random deviations in different directions compensate for one another, in effect washing indeterminism out of the system at the macroscopic level. If that were the case then interactions among what J.L. Austin called medium-sized dry goods would, for all intents and purposes, be deterministic in character.¹⁹⁰

190. Perry, *supra* note 25, at 337; see HOWSON & URBACH, *supra* note 73, at 341 (similarly distinguishing between macro-level determinism and micro-level indeterminism); Kenneth F. Schaffner, *Causing Harm: Epidemiological and Physiological Concepts of Causation*, in ACCEPTABLE EVIDENCE: SCIENCE AND VALUES IN RISK MANAGEMENT 204, 208 (1991) (same); Alvin Goldman,

In short, notwithstanding the indeterminacy of quantum mechanics, it is quite unclear whether the physical probability of some macroscopic event *E* resulting in a given person's premature death—where *E* is, say, a toxic release, the ingestion of some food, an industrial practice, the emission of some pollutant or some such event targeted by regulatory agencies—ever takes a numerical value other than 1, in the case where death results, or 0, in the case where it doesn't. If the causal laws relevant to risk regulators are indeed deterministic, it is incoherent to conceptualize a high physical probability of premature death as a harm that is separate from, and additional to, the harm of death itself.

D. CONVENTIONAL REFERENCE CLASSES

The probability that *E* results in *P*'s premature death could be determined by placing *E* and *P* within classes of events and persons that are *conventionally* similar to *E* and *P*, rather than statistically similar or physically similar.¹⁹¹ We might ask how an ordinary person, within *P*'s society, would characterize *E* and *P*—more precisely, how an ordinary person, concerned to gauge the risk that *E* imposes on *P*, would characterize *E* and *P*. For example, if *E* is the ingestion, at 10 p.m., on a hot summer day, of an all-beef hot dog containing sodium nitrite by a thirty-five-year old Caucasian man *P* weighing 200 pounds, with a family history of liver cancer, and who tends to eat red meat and other high-fat, high-cholesterol foods, then the ordinary person might see this as a case of a large white man eating a meat product. $\{E_i\}^{CONV}$, the class of events conventionally similar to *E*, would be the class of all eatings of meat products by large white men; $\{P_i\}^{CONV}$ is the class of men who eat these products. Clearly, $\{E_i\}^{CONV}$ and $\{P_i\}^{CONV}$ are different from $\{E_i\}^{STAT}$ and $\{P_i\}^{STAT}$, and $\{E_i\}^{PHYS}$ and $\{P_i\}^{PHYS}$, respectively.¹⁹² The "conventional" probability that *E* results in

Quasi-Objective Bayesianism and Legal Evidence 14-15 (unpublished manuscript, on file with author) (same).

191. Sociologists, in particular, have emphasized the conventionality of risk ascriptions. For reviews of the sociological literature on risk, see SHRADER-FRECHETTE, *supra* note 5, at 27-52; Kasperson, *supra* note 69; Palmlund, *supra* note 69; Rayner, *supra* note 69; Renn, *supra* note 69.

192. The huge literature on the discrepancy between "lay" and "expert" risk perceptions bears out my point that the conventional (i.e., lay) and statistical (i.e., expert) reference classes used to generate probabilities in the frequentist sense can differ. See, e.g., Gillette & Krier, *supra* note 5, at 1071-85 (discussing the lay/expert discrepancy and citing sources); Richard H. Pildes &

P's premature death is, let us say, the frequency with which persons within $\{P_i\}^{\text{CONV}}$ die earlier than they would have had the matching events within $\{E_i\}^{\text{CONV}}$ not occurred.

Is O_R worse than O_S for P, when the difference between the two outcomes is specified in terms of conventional probability?

Riskier and Safer Outcomes: A Frequentist Account, with Conventional Probabilities

- O_S is a world in which E does not occur.
- O_R is a world in which (1) E occurs, and (2) the conventional probability that E results in P's premature death is high. In other words, the relative frequency of premature death within $\{P_i\}^{\text{CONV}}$, as a result of events within $\{E_i\}^{\text{CONV}}$, is high.
- O_R and O_S are otherwise identical with respect to P's welfare.

Note that the conventional probability of premature death, like the statistical probability of premature death—and unlike the physical probability of premature death—can take a value other than 0 or 1, and can be distinguished from the actual occurrence of premature death, even in a world where causal laws are deterministic. Note too that a high conventional probability of premature death—unlike a high statistical or physical probability—arguably satisfies the Integration requirement. As I have already suggested, it is arguable that a high frequency of harm to some group is sufficiently integrated into P's own life, to count as a welfare setback to him, if P is socially or conventionally linked to the group even if P himself does not identify with the group.¹⁹³

Nonetheless, it is problematic to claim that a high conventional probability of premature death is itself a welfare setback. First, that claim runs afoul of the Experiential component of welfare, for just the same reason that the parallel claims with respect to statistical and physical probability do.¹⁹⁴ Second, I am not persuaded that O_R —specified in terms of conventional probability—is worse than O_S as a matter of Preference or Value. P might prefer O_S to O_R , but would he prefer O_S under conditions of full information? Physical probability is a better gauge of P's security from premature

Cass R. Sunstein, *Reinventing the Regulatory State*, 62 U. CHI. L. REV. 1, 48-64 (1995) (same).

193. See *supra* text accompanying note 187.

194. See *supra* text accompanying notes 170-71, 185.

death than either statistical or conventional probability, so why think that fully informed P would care about a high conventional probability of a premature death? Similarly, why think that agents under ideal conditions (including full information) would converge on the judgment that it is better to be P in O_s than P in O_r ?

The answer could be that fully informed P, and ideal agents more generally, would care about a high conventional probability of P's premature death under some heading other than "security from premature death." Consider this analogy. An outcome in which P is conventionally thought to be stupid is worse for P, with respect to Value and Preference, than an outcome in which P is not conventionally thought to be stupid. But this is not because conventional judgments of stupidity are the best gauge of stupidity. Rather, it is intrinsically bad to be thought stupid; and thus fully informed P, and ideal agents, would disprefer an outcome in which the ordinary member of P's society takes him to be stupid, even though there are much better indicators of P's real intelligence than this ordinary judgment.¹⁹⁵

I am not persuaded that the analogy works. Certain beliefs that others might hold about P—in particular, the belief that P is less than a full person, or lacks some characteristic ordinarily held by full persons, or has acted in a way that full persons do not or ought not act—surely *do* satisfy the Value and Preference components of well-being. Fully informed P, and ideal agents considering the prospect of being P, would disprefer outcomes in which others have such beliefs about P—quite apart from other, negative consequences with which such beliefs might be linked. Yet beliefs about P's security are not, I think, like this. P's stupidity detracts from her full personhood; her vulnerability to premature death does not, since even the fullest of (human) persons must die, and indeed will die prematurely as a result of *some* event, relative to the time of death had the event not occurred. Imagine two outcomes differentiated solely by the fact that, in one outcome, P's fellow

195. Cf. Matthew D. Adler, *Expressive Theories of Law: A Skeptical Overview*, 148 U. PA. L. REV. 1363, 1422-25, 1432-33, 1445-47 (2000) (describing the widespread argument, within constitutional and criminal-law scholarship, that draws a link between the lowered social status for a given category of individuals constituted by prevalent beliefs that they lack full personhood, and the harm of lowered self-respect suffered by those individuals).

citizens believe him to be at high risk of an early heart attack, while in the second outcome they do not possess this belief. The citizens' remaining beliefs, and their actions, are identical in both outcomes. Would P under full information disprefer the first outcome, even if (as he knows under full information) the physical probability of an early heart attack is low? It is hard to see why he would.

Finally, it is an open question whether conventional probabilities actually exist. Given some event E, E has a conventional probability (high or low) of some attribute only if there are existing norms, practices, or understandings, within the relevant society, for characterizing E.¹⁹⁶ When we hypothesize how an "ordinary" person would envision the threat that E poses to P's security, we mean a person who follows the prevailing norms for characterizing risky events. Absent some such norms, there would be no reason to expect any consensus in the characterization of E—no single reference class that all or most members of the society would deploy in judging E's riskiness.

Much work in sociology suggests that collective norms for characterizing risky events (if they ever existed) have broken down, at least in the United States and other developed countries. Mary Douglas and Aaron Wildavsky, in their so-called "cultural" theory of risk, famously argue that three divergent "cultures," or socially constructed forms of life and world-views—a market culture, a hierarchical culture, and a sectarian culture—coexist in the United States and substantially shape risk perception.

[P]ublic perception of risk and its acceptable levels are collective constructs, a bit like language and a bit like aesthetic judgment

. . . .

. . . Analysis of this kind recognizes different types of cultures resulting from their members' sustained attention to feasible social goals. Three have seemed sufficient for the thesis being argued. Two of these derive from the established stock of political thought in the West, market individualism and hierarchy. Neither encompasses the structure and goals of the environmental movement. [That movement exemplifies] the [cultural] type we have called sectarian. . . . [T]he risks of war are not acceptable to the hierarchist because he is

196. On social facts in general, see JOHN R. SEARLE, *THE CONSTRUCTION OF SOCIAL REALITY* 1-126 (1995). Social norms, in particular, have garnered much recent attention by legal scholars. See Adler, *supra* note 195, at 1373 n.44. Helping to define a legally operative concept (in this case, risk) is certainly the kind of legal role that social norms could play.

focused on dangers of foreign relations Risks of economic collapse are ever present to the mind of the market individualist Risks from technology are uppermost in the sectarian mind¹⁹⁷

In short, whether an individual tends to perceive a high level of "risks of human violence," as opposed to "risks from technology," or "risks from economic failure," depends on her "cultural" type (hierarchical, sectarian, market).¹⁹⁸ To be sure, Douglas and Wildavsky's reduction of risk perception and evaluation to three cultural types, thus specified, is quite controversial within the sociological literature on risk.¹⁹⁹ But their work advances a broader point: that in a modern society there exist multiple and competing socially constructed entities that guide risk perception and evaluation. These socially constructed entities are what I referred to above as norms for characterizing risky events. They might be specified as "ways of life," perspectives, world-views, or cultural patterns. In any event, the broader Douglas/Wildavsky thesis that in a developed society divergent norms will produce divergent perceptions of the riskiness, and appropriate response to, a given event is *not* controversial among sociologists. To quote one recent review article,

By going beyond purely psychological analyses, the social approaches to risk perception highlight . . . that the notion of a 'lay' public as an undifferentiated risk-perceiving entity is a misnomer: a society is always composed of many groups with very different attitudes towards, and appraisals of, what risk is, which risks should be run, and what values are relevant to making acceptability decisions.²⁰⁰

E. HYBRID ACCOUNTS

I have repeatedly noted that risk, in the frequentist sense, is not harmful with respect to the Experience component of welfare.²⁰¹ The reader might wonder why the complaint is not answered by *hybridizing* a frequentist probability with a

197. MARY DOUGLAS & AARON WILDAVSKY, *RISK AND CULTURE: AN ESSAY ON THE SELECTION OF TECHNICAL AND ENVIRONMENTAL DANGERS* 186-88 (1982).

198. *Id.* at 187.

199. See, e.g., SHRADER-FRECHETTE, *supra* note 5, at 29-39; Rayner, *supra* note 69, at 83.

200. Nick F. Pidgeon & Jane Beattie, *The Psychology of Risk and Uncertainty*, in *HANDBOOK OF ENVIRONMENTAL RISK ASSESSMENT AND MANAGEMENT* 289, 302 (Peter Calow ed., 1998).

201. See *supra* text accompanying notes 170-71, 185, 194.

mental state, such as a belief.²⁰² In other words, construe O_R and O_S as follows:

Riskier and Safer Outcomes: A Hybrid Account, with Frequentist Probabilities

- O_S is a world in which E does not occur.
- O_R is a world in which (1) E occurs and (2) the frequentist probability (in some sense) of P dying prematurely as a result of E is high, and in addition (3) P “experiences” the difference between the outcomes, i.e., believes he is subject to a high frequentist probability of dying prematurely as a result of E.
- O_R and O_S are otherwise identical with respect to P’s welfare.

Isn’t it *now* the case that O_R is worse for P than O_S ? After all, the Experientialist demand for supervenience is satisfied, since P in O_R has a mental state (a belief) which he lacks in O_S . Note further that the Integration component of welfare is also satisfied, since (as I shall argue at greater length below)²⁰³ differences in someone’s mental states are plausibly sufficient if not necessary to satisfy Integration.

Still, it is not clear that O_R is really worse for O_S . Here again is a point that shall be developed at greater length below, in the course of my analysis of the Bayesian view of risk: it is open to question whether mere differences in what someone believes can change the value of outcomes for him.²⁰⁴ To begin, assume that mere beliefs do not change the value of outcomes. If so, and if the high frequentist “probability” in O_R is either a high *statistical* probability or a high *conventional* probability, the upshot will be that O_R is *not* worse than O_S with respect to the Value and Preference components of welfare. As I have already argued, neither a high statistical probability of premature death, nor a high conventional probability, is a

202. It is a standard view, within the philosophical literature on value, that things can be valuable as parts of larger wholes. See, e.g., NOAH M. LEMOS, *supra* note 178, at 32 (articulating the “principle of organic unities,” namely “that the value of some wholes is not the same as the sum of the values of their parts”); RAMON M. LEMOS, *supra* note 178, at 41-52 (discussing “contributory value”). By extension (or equivalently), it is plausible that the welfare impact of a hybrid entity is not always separable into the welfare value of its parts. In the economic tradition, this problem goes under the heading of “separability.” See JOHN BROOME, *WEIGHING GOODS* 60-89 (1991) (discussing separability).

203. See *infra* text accompanying notes 208-10.

204. See *infra* text accompanying notes 214-22.

setback with respect to either Value or Preference.²⁰⁵ If, alternatively, the high frequentist probability in O_R is a high *physical* probability, the problem of determinism emerges—a problem that besets any physicalist defense of the “risk is harm” claim, be it a hybrid or simple defense.

Next, let us assume that mere differences in what someone believes *can* change the value of outcomes from him. Now, O_R —construed as a hybrid of a high frequentist probability of premature death for P, plus P’s belief that this probability obtains—is worse for P than O_S . Experience and Integration are satisfied and, given our altered assumption about value, so are Value and Preference. But by altering the assumption about Value (and Preference), we have made frequentist probability *superfluous* to any harm that P suffers. Imagine a world O_{R^*} , in which (1) E occurs, and (2) P is subject to a trivially low frequentist probability of premature death as a result of E, but (3) P has the same belief as in O_R , namely that he is subject to a high frequentist probability of premature death. Given the altered assumption about the power of mere belief-differences to affect value, it follows that O_{R^*} is worse for P than O_S . Further, O_R is *not* worse for P than O_{R^*} ; the two outcomes differ only in the frequentist probability, and that difference (without more) does not constitute a welfare difference.²⁰⁶

205. See *supra* text accompanying notes 172-80, 194-95.

206. It might be objected that my analysis fails to consider (1) the possibility that frequentist risk plus a belief is harmful even though neither is harmful alone, and (2) the possibility that, although beliefs alone can be harmful, frequentist risk compounds the harm. My response is that neither possibility is warranted by the nonepistemic values I have considered or will consider, e.g., the value of security or freedom from fear. See *supra* text accompanying notes 174-75, 185, 195; *infra* Part IV.B. As for epistemic values, such as the putative values of knowledge, justified belief, or true belief, these values do allow the welfare importance of beliefs to change depending on other factors (truth and/or justifiability), but they would not warrant the conclusion that O_R is worse for P than O_S in virtue (in part) of the frequentist risk. If, for example, the relevant epistemic value is avoiding false beliefs, then O_R may not be worse than O_S (because in O_R P truly believes that a high frequentist probability obtains, although falsely believes he will die prematurely), and in any event O_{R^*} is even worse (because in O_{R^*} P both falsely believes he will die prematurely and falsely believes there is a high frequentist probability). If the relevant epistemic value is avoiding unjustified beliefs, then O_R is no worse than O_S (because P’s belief in his premature death is justified by the frequentist probability) and again O_{R^*} is worse than O_R (because there the belief in premature death is not justified). In short, the strategy of adding a belief about frequentist probabilities to the outcome in which P is (without awareness) subject to a high frequentist probability of

In sum, the flaws in the frequentist version of the risk-is-harm claim, elaborated in sections B, C, and D above, are not cured by hybridizing relative frequencies with beliefs. More generally, these flaws are not cured by hybridizing relative frequencies with *any* mental states, be they beliefs, feelings, affects, urges, preferences, or something else.

IV. IS THE RISK OF DEATH A HARM? BAYESIAN THEORIES OF RISK

The prior Part argued that risk in the frequentist sense is not harmful. More precisely, there are a plurality of frequentist accounts of risk, corresponding to different types of reference classes: statistical, physical, and conventional. The claim advanced above was that no such account makes risk a welfare setback.

This Part examines the Bayesian view of risk. Here too, as we shall see, there are a plurality of accounts. The Bayesian views risk as a “degree of belief,” a partial belief or state of uncertainty that can be represented by numbers conforming to the probability calculus.²⁰⁷ But whose beliefs are at issue? Are they actual or hypothetical? Are they separated from other mental states, or rather packaged together with feelings, volitions, desires, or other non-belief states to create a psychological hybrid? Different answers to these questions generate different Bayesian accounts of risk.

Since we are concerned with the risk of premature death, the probabilistic beliefs underlying the various accounts are (actual or hypothetical) probabilistic beliefs about premature death—more precisely, beliefs that a particular event E will result in the premature death of some particular person P. I will argue that P’s own probabilistic belief that E will result in P’s premature death, taken alone, may not be a welfare setback for P. However, this belief, as part of a complex psychological hybrid that amounts to a kind of *fear* on P’s part—a hybrid including not only P’s probabilistic belief about his premature death, but also the *distress* or other such negative affective state occasioned by P’s belief—does constitute a welfare setback for P. Finally, I will suggest that Bayesian accounts that

premature death tends to improve, not degrade, the outcome with respect to plausible epistemic values because the frequentist probability functions as a truth maker or justification maker with respect to the belief.

207. See *supra* text accompanying notes 73-76.

prescind from P's actual beliefs, and look instead to his hypothetical beliefs, or to the actual or hypothetical beliefs of other persons, face more substantial difficulties in explaining why risk is harmful.

A. FIRST-PERSON ACTUAL BELIEFS

Throughout this Part, we will be comparing two outcomes, O_s and O_r , that differ only in the occurrence of an event (E) such that someone has a high degree of belief that E will result in P's premature death. That person might be P himself. If so, and if this belief is taken to be actual not hypothetical, the Bayesian account will specify the difference between O_s and O_r as follows:

Riskier and Safer Outcomes: A Bayesian Account, with Actual First-Person Beliefs

- O_s is a world in which E does not occur.
- O_r is a world in which (1) E occurs and (2) P actually believes, to a high numerical degree, that E will result in P's premature death.
- O_r and O_s are otherwise identical with respect to P's welfare.

Is O_r , thus specified, worse for P's welfare than O_s ?

Consider once more the four plausible components of a welfare account: Experience, Integration, Preference, and Value. Experientialists (at least those who specify the Experience component in terms of a supervenience requirement) will readily agree that O_r can be worse for P than O_s . There is a psychological difference between O_r and O_s . P possesses a particular probabilistic belief in the first outcome, but not the second. Because belief is a mental state, and because the Bayesian equates risk with a certain kind of belief, a mental-state supervenience requirement poses no substantial obstacle to a Bayesian's claim that risk is harmful.

The Integration component of welfare stipulates that the difference between O_r and O_s must concern P's own life. This is very vague, but still I think we can see that O_r and O_s (as here specified) do thus differ. More generally, I suggest that two outcomes differing in P's psychological properties are different "for P." A psychological difference is sufficient, if not necessary, for Integration.²⁰⁸ What could be more intimately part of P

208. See Shelly Kagan, *The Limits of Well-Being*, 9 SOC. PHIL. & POL'Y 169,

than his own mental states? Mark Overvold's version of Integration, which remains the leading version within the philosophical literature, bears out my suggestion that psychological differences will always satisfy the Integration demand.²⁰⁹ Overvold suggests that O^* is better or worse for P than O only if some proposition is true in O^* but not O , where this proposition concerns P because it entails P 's existence.²¹⁰ The fact that P possesses a given psychological state in O^* —such as a probabilistic belief—clearly entails that P exists in O^* . Many propositions would not have this entailment. For example, the pains and pleasures of other persons cannot, on Overvold's view, change P 's own welfare, because these pains and pleasures might obtain without P ever existing. So the Integration requirement, in Overvold's hands, is quite demanding—but the requirement is structured in such a way that psychological changes in a person, including the cognitive changes highlighted by Bayesians, will always satisfy it.

Value and Preference pose more serious obstacles to the claim that O_R is worse for P than O_S . Consider Value first. Values, as already discussed, plausibly satisfy a condition of convergence under idealized conditions.²¹¹ O_S is more welfare valuable for P than O_R only if everyone, deliberating under these conditions, would prefer to be P in O_S rather than P in O_R . In general, there is nothing puzzling in the claim that mental states are valuable or disvaluable. Take the case of physical pain. Assume that P is in a state of pain in O^* , but not O . Presumably our idealized agent will not himself be in a state of pain. Still, it is straightforward that the agent could prefer to be P in O , rather than P in O^* . Pain feels bad; idealized agents, contemplating the prospect of being various persons in various non-ideal scenarios, could readily prefer non-ideal

182-87 (1992) (suggesting that changes in our mental states paradigmatically satisfy the Integration demand because they are changes in our intrinsic rather than relational properties).

209. See sources cited *supra* note 61.

210. This is a rough formulation of Overvold's view. More precisely, he proposes,

[T]he only desires and aversions that are logically relevant to the determination of an individual's self-interest are those in which (1) it is logically necessary that the individual exist at t for the object of one's desire or aversion to obtain at t , and (2) the reason for this desire is due to one's essential involvement in the state of affairs.

Overvold, *Self-Interest*, *supra* note 61, at 190.

211. See *supra* text accompanying notes 177-79.

scenarios that are less painful.²¹²

Could the same be said about P's beliefs, specifically his probabilistic beliefs concerning his own premature death? In O_s , E does not occur; in O_r , E occurs and P believes to a high degree that this event will result in his premature death. Let us assume that our idealized agent is perfectly informed and thus knows the truth about E and P, namely, that E will not result in P's premature death. Still, the idealized agent *could* prefer being P in O_s to being P in O_r . The idealized agent's knowledge does not erase the psychological difference between the two non-ideal states under comparison: the fact that P possesses a certain belief in O_r , and not in O_s .²¹³ Perfectly informed agents, choosing between life histories where the subjects of these histories will be imperfectly informed, could prefer life histories in which those subjects lack certain beliefs—for example, beliefs that the choosers, in their position of perfect information, know to be false. There is no conceptual incoherence in the claim that false beliefs, or some other class of beliefs, are *per se* disvaluable.

But is the claim correct? Is the possession of false beliefs really an intrinsic setback to welfare value? Or are such beliefs only instrumentally disvaluable, in that the choices of ignorant actors produce consequences which are typically worse for those actors than the consequences that would have resulted from better-informed choices?²¹⁴ I see no clear, intuitive answer to this question. Beliefs, unlike pains, are affectively neutral.²¹⁵ False beliefs do not feel worse than true beliefs. Indeed, beliefs (unlike pains) need not be occurrent. I need not be conscious of my beliefs, be they full beliefs or partial

212. While the Benthamite reduction of happiness or welfare to pleasure and the avoidance of pain is implausible, those quite plausibly "belong instead in our inventory of the (typical or standard) sources of happiness or misery." SUMNER, *supra* note 8, at 142.

213. See SMITH, *supra* note 179, at 151-52 (distinguishing between the "evaluating possible world," where the evaluator has full information, and the "evaluated possible world," where the person whose choices or life history are being evaluated does not).

214. See, e.g., NOAH M. LEMOS, *supra* note 178, at 3-31 (analyzing the concept of intrinsic value); RAMON M. LEMOS, *supra* note 178, at 34-71 (same).

215. See Ramsey, *supra* note 76, at 71 (considering and rejecting the notion of equating a person's degree of belief with the intensity of his feeling of conviction in the proposition believed—"the beliefs which we hold most strongly are often accompanied by practically no feeling at all; no one feels strongly about things he takes for granted").

(probabilistic) beliefs.²¹⁶

A number of prominent value theorists have argued that knowledge is an intrinsic welfare value. For example, George Sher writes,

[D]espite our manifest differences, each of us has both a native capacity to understand the world and an inescapable tendency to try to exercise that capacity. We seek such understanding not only when we propose a scientific hypothesis or try to predict the stock market, but when we idly read the newspaper, engage in conversation, or simply look about us. . . . In each case, our mental activity centers on some proposition, or some structured set of propositions, that has engaged our attention; and in each case, our guiding question is whether these propositions are true. Even when we try to avoid all truth seeking, our success is temporary at best. . . . Thus, if what has inherent value is the successful exercise of fundamental capacities, it would be very surprising if reason-based true belief—or, in other words, knowledge—were not inherently good.²¹⁷

John Finnis, James Griffin, and Derek Parfit also include “knowledge” or “understanding” on their lists of objective values.²¹⁸ On the other hand, Martha Nussbaum suggests that value inheres in the process of thinking, in intellectual activity, not in the sheer possession of knowledge.²¹⁹ Thomas Hurka has shown, in a systematic way, that a value theory predicated (in part) on the intellectual capacities of persons need not end up valuing knowledge. Instead, such a theory could focus on the justifiability or reasonableness of beliefs, or even on the simple

216. See JAEGWON KIM, *PHILOSOPHY OF MIND* 13-14, 158 (1996) (explaining that beliefs and desires are “propositional attitudes,” which have no distinctive “feel” and can be unconscious); see also Wayne A. Davis, *The Varieties of Fear*, 51 *PHIL. STUD.* 287, 287 (1987) (noting that “propositional fear”—meaning a desire that some harm not occur plus a belief that it may occur—“need not be occurrent, and may have no physiological effect”).

217. SHER, *supra* note 49, at 203.

218. See FINNIS, *supra* note 49, at 59-80 (arguing for intrinsic value of knowledge); GRIFFIN, *supra* note 23, at 29-30 (including “understanding” on a list of basic welfare values, and stating that “[s]imply knowing about oneself and one’s place in the world—certain important anthropocentric knowledge—is part of a good life”); PARFIT, *supra* note 18, at 499 (proposing a list of objective values that includes “knowledge”).

219. See NUSSBAUM, *supra* note 23, at 78-79. Nussbaum includes “Senses, Imagination, and Thought” on a list of basic values and specifies this series of values as follows:

Being able to use the senses, to imagine, think, and reason—and to do these things in a “truly human” way, a way informed and cultivated by an adequate education Being able to use imagination and thought in connection with experiencing and producing self-expressive works and events of one’s own choice

Id.

number of beliefs (true or false), rather than on the truth of belief.

[We have a] choice between four views that I call *attempt*, *deserving attempt*, *success*, and *deserved success*. . . . [These] views would say a person has one state of theoretical perfection [and thereby produces intrinsic value] for every:

1. belief he has (attempt);
2. justified belief he has (deserving attempt);
3. true belief he has (success); [or]
4. justified true belief, that is, item of knowledge, he has (deserved success).²²⁰

In defense of Hurka's second option, "deserving attempt," note the following. It may be essentially human, and unavoidable for humans, to aim at true belief, but humans (except in unusual cases) lack direct access to truth, and regulate their intellectual activities by criteria to which they do have access, namely "justifiability" or "reasonableness."²²¹ Arguably, then, it is the possession of justified beliefs, and the activity of testing beliefs for justification, not true belief (or true justified belief) that constitutes a kind of welfare value.

In sum, both sides of the Value issue are plausible, given the current state of philosophical writing about Value. Neither the claim that O_R is worse than O_S with respect to welfare value, nor the claim that the two outcomes are just the same, is clearly correct.²²² Therefore it is also unclear whether O_R is worse than O_S with respect to the Preference component of welfare—given the tight linkage between Value and fully informed preferences.

220. HURKA, *supra* note 49, at 103. *But see id.* at 112 (suggesting that the "deserved success" view is intuitively most appealing).

221. See generally JONATHAN DANCY, AN INTRODUCTION TO CONTEMPORARY EPISTEMOLOGY (1985) (discussing concepts central to epistemology, including knowledge and justification).

222. Note that, strictly, the view that knowledge (justified true belief) is beneficial would not entail a welfare difference between O_S and O_R as here described. In O_S , P lacks a belief about E's connection to her premature death; in O_R , she falsely believes that E will cause her premature death. In neither outcome does she possess knowledge about E. Still, the view that knowledge is beneficial sits comfortably with the further claim that false beliefs are harmful relative to a no-belief state. On this view, persons at risk in the Bayesian sense are thereby harmed, except in the unusual case where their beliefs are correct—the risky event will in fact kill them—in which case, assuming the belief is justified, the beliefs actually benefit them! On Hurka's "deserving attempt" view, some persons merely at risk are also actually benefited, not harmed, by their beliefs (insofar as the beliefs are justified); only those whose beliefs are unjustified are not.

I should emphasize that I have not definitively concluded that O_R and O_S are equal with respect to Value or fully informed Preference. P's probabilistic belief that E will result in his premature death is *arguably* disvaluable, without more. Readers who think it is—because false beliefs are intrinsically disvaluable, or perhaps on other grounds—should conclude that O_R is worse for P than O_S . But readers who doubt that false probabilistic beliefs, or other categories of probabilistic beliefs, are intrinsically disvaluable, taken alone, should not conclude that Bayesian risk is harmless. Bayesian risk, although harmless on its own, could be harmful as part of a larger psychological hybrid such as *fear* or cognate emotional states. This is the possibility I shall now consider.

B. FEAR

What is fear? It is a standard view within experimental psychology and philosophy that fear and other emotions have cognitive, attitudinal, and affective components.²²³ Wayne Davis offers a crisp and paradigmatic analysis along these lines.²²⁴ He distinguishes between “propositional fear” and “experiential fear.” As for propositional fear,

[p]ropositional fears have attitudinal and cognitive components. If I am afraid that it will rain, then I must not want it to rain; indeed, I must desire (on balance) that it not rain. Moreover, I must be neither certain that it will rain, nor certain that it will not. Fear, in other words, entails *aversion* and *uncertainty*.²²⁵

223. The literature on emotions is vast. Important works in the philosophical literature include: RONALD DE SOUSA, *THE RATIONALITY OF EMOTION* (1987); PETER GOLDIE, *THE EMOTIONS: A PHILOSOPHICAL EXPLORATION* (2000); ROBERT M. GORDON, *THE STRUCTURE OF EMOTIONS: INVESTIGATIONS IN COGNITIVE PHILOSOPHY* (1987); O.H. GREEN, *THE EMOTIONS* (1992); PATRICIA GREENSPAN, *EMOTIONS & REASONS: AN INQUIRY INTO EMOTIONAL JUSTIFICATION* (1988); MARTHA NUSSBAUM, *UPHEAVALS OF THOUGHT: THE INTELLIGENCE OF EMOTIONS* (2001); RICHARD WOLLHEIM, *ON THE EMOTIONS* (1999); John Deigh, *Cognitivism in the Theory of Emotions*, 104 *ETHICS* 824 (1994). This literature is, in turn, dwarfed by the psychological literature. For an overview, see RANDOLPH R. CORNELIUS, *THE SCIENCE OF EMOTION: RESEARCH AND TRADITION IN THE PSYCHOLOGY OF EMOTION* (1996). Economists have also begun to study emotions, see Eric A. Posner, *Law and the Emotions*, 89 *GEO. L.J.* 1977, 1978 n.2 (2001) (citing scholarship), as have legal scholars, see Peter H. Huang, *Trust, Guilt and Securities Regulation*, 151 *U. PA. L. REV.* (forthcoming 2003); Posner, *supra*, at 1978 n.1; cf. George F. Loewenstein et al., *Risks as Feelings*, 127 *PSYCH. BULL.* 267, 267-71 (2001) (focusing on role of affect in explaining behavior under uncertainty).

224. See Davis, *supra* note 216; see also Wayne Davis, *A Causal Theory of Experiential Fear*, 18 *CANADIAN J. PHIL.* 459 (1988).

225. Davis, *supra* note 216, at 289.

There is an obvious connection between what Davis calls "propositional fear," and risk in the Bayesian sense. Risk in the Bayesian sense is a kind of uncertainty—an uncertainty whose degree can be expressed in terms of numbers conforming to the probability calculus. If a person (1) has a probabilistic belief that he will die prematurely, and (2) prefers not to die prematurely, then—according to Davis—that person has a "propositional fear" of premature death. "Propositional fear" does not entail bad feelings, or other affective states, on the subject's part. A subject may not even be aware of his "propositional fears," since he need not be aware of his beliefs or his preferences. By contrast, "experiential fear" necessarily includes an affective component. Davis defines "experiential fear" as a combination of occurrent propositional fear *plus* "involuntary arousal" and "unhappiness":

A subject experiencing fear must have some propositional fears. These fears must, furthermore, be *occurrent*: they must be before the subject's mind; he must be thinking about them. . . .

An occurrent propositional fear is not sufficient for experiential fear, however. If Alan is having the time of his life with Monique, and the thought that he might fail his exam just pops into his mind for an instant, it does not follow that he is in a state of fear. What else is required? For one thing, the propositional fear must cause a particular condition characterized by symptoms like rapid heartbeat, increased respiration rate, perspiration, abdominal distress, pallor, restless sleep, channeled and disorderly cognition, muscular tension, and impulsive or reflexive motor action. I refer to this condition as *involuntary arousal*. . . . It must be stressed that the symptoms of involuntary arousal are not perfectly correlated with each other, and that no one symptom is necessary for fear. A subject with a pounding heart but no abdominal distress, for example, could still be afraid.²²⁶

Davis stipulates that "experiential fear" involves propositional fear plus both (1) involuntary arousal and (2) "unhappiness" (that is, a sense of distress), so as to distinguish the genuine experience of fear from cases where people enjoy scary thrills.

People seek out scary thrills to enjoy themselves. The involuntary arousal is part of the fun. Similarly, a performer suffering from stage-fright may be ecstatic because this is his "big break." Fear, however, is a type of mental pain or anguish. Subjects must be miserable or distressed because they fear that some harm will occur.²²⁷

In short, according to Davis, "S is experiencing fear iff [if and only if] S is in a state of involuntary arousal and unhappiness

226. *Id.* at 299.

227. *Id.* at 301.

as a direct result of the occurrent propositional fear that something will be harmed."²²⁸

Davis's analysis, and the numerous similar accounts of fear within the psychological and philosophical literatures,²²⁹ suggests a modification to the Bayesian account that will make risk an essential component of harm, if not harmful on its own. Combine the cognitive state of probabilistic belief with a negative affective state (paradigmatically, one that involves both "involuntary arousal" and "unhappiness"), and then add suitable preferences.²³⁰ In other words, replace outcome O_R with O_{R+A} (risk plus affect), as follows:

Riskier and Safer Outcomes: A Bayesian Account, with Actual First-Person Beliefs Plus Negative Affect

- O_S is a world in which E does not occur.
- O_{R+A} is a world in which (1) E occurs and (2) P experientially fears that E will result in P's premature death. Specifically, P believes, to a high numerical degree, that E will result in his premature death; this probabilistic belief is occurrent; and it causes P to experience involuntary arousal plus a sense of distress or unhappiness.
- O_S and O_{R+A} are otherwise identical with respect to P's welfare.
- In both outcomes, P disprefers premature death and the experiential fear of premature death.

I suggest that O_{R+A} is worse for P than O_S .

In the preceding section, I argued that O_R (an outcome involving P's probabilistic belief about premature death, but no negative affect) is worse than O_S with respect to Experience and Integration, but might not be worse than O_S with respect to Value and Preference. What do these arguments imply for the comparison of O_{R+A} (an outcome involving P's experiential fear of premature death, including both probabilistic belief and negative affect) and O_S ? First, the arguments imply that O_{R+A} , like O_R , is worse than O_S with respect to Experience and

228. *Id.* at 302.

229. See, e.g., AHARON BEN-ZE'EV, THE SUBTLETY OF EMOTIONS 473-89 (2000); GREENSPAN, *supra* note 223, at 15-36; Barrie Falk, *What Are We Frightened of?*, 25 INQUIRY 165, 186-95 (1982); Robert Gordon, *Fear*, 89 PHIL. REV. 560 (1980); O.H. Green, *Fear of Death*, 43 PHIL. & PHENOMENOLOGICAL RES. 99, 105 (1982); Mohan Matthen, *Biological Universals and the Nature of Fear*, 95 J. PHIL. 105, 121-32 (1998).

230. On hybrids, see *supra* note 202.

Integration. The difference between O_R and O_S satisfies the Experientialist demand for supervenience; a fortiori, the difference between O_{R+A} and O_S does, since these outcomes now differ with respect to *two* aspects of P's psychology, his affect and his cognitions, not just with respect to his cognitions. Similarly, if the difference between O_R and O_S concerns P's own life, then it follows a fortiori that the difference between O_{R+A} and O_S does.

How about the Value component of welfare? It is here that the shift from O_R to O_{R+A} becomes important. Experiential fear involves a physically "aroused" sense of unhappiness or distress. That affective state is, intuitively, a bad state to be in: not just bad in some impersonal sense, but bad for the unhappy subject, bad for her well-being. Many (although perhaps not all) value theorists would recognize experiential fear as a disvaluable state. In her list of welfare constituents, Nussbaum explicitly singles out "fear" and "anxiety" as welfare setbacks.

[The] Central Human . . . Capabilities [include] . . . *Emotions*.—Being able to have attachments to things and people outside ourselves; . . . in general, to love, to grieve, to experience longing, gratitude, and justified anger. Not having one's emotional development blighted by overwhelming fear and anxiety . . .²³¹

Griffin's list includes "enjoyment"; he writes under this rubric that "[w]e value pleasures, the perception of beauty, the enjoyment of the day-to-day textures of life,"²³² and presumably would agree that pain or suffering (including the physically aroused type of distress that is partly constitutive of fear) are bad things. Sher argues,

[T]he elements of a good life include knowledge, rational action, close relationships, and various other forms of contact with the world. This means that whether someone lives a good life cannot depend exclusively on the quality of his experience. But neither is the quality of his experience irrelevant. We can hardly deny that happiness, pleasure, and enjoyment are among life's goods, so any satisfactory unifying theory must appeal to a property or relation that is capable of belonging to experiential as well as nonexperiential states.²³³

Sher's own theory of value—one that sees value in those activities, experiences, etc., closely connected to "near-

231. NUSSBAUM, *supra* note 23, at 79.

232. GRIFFIN, *supra* note 23, at 30; *see also id.* (including under the value of human agency "the basic capabilities that enable one to act," including "freedom from great pain and anxiety").

233. SHER, *supra* note 49, at 229.

universal, near-unavoidable goals"²³⁴—supports the view that fear and other unpleasant experiences are harmful, since “no one can altogether suppress his impulses to pursue pleasant sensations and escape unpleasant ones.”²³⁵ By contrast, value theorists such as Hurka who ground value in the human essence²³⁶ might not recognize fear as a welfare setback, since the capability to experience pain or negative feelings, as opposed to the capability to think and to plan, might not be essentially human. A human person who could think and plan, but couldn’t experience fear or other emotions, would still be a human person²³⁷—or perhaps she wouldn’t be, in which case an essentialist view of Value might be able to recognize the disvalue intrinsic to fear. In any event, powerful arguments have been advanced that essentialism is an overly narrow view of Value.²³⁸

As for Preference: a person could, in principle, be indifferent to (or even prefer) a state of experiential fear. It is certainly possible for P to be indifferent to or even prefer painful, unpleasant, distressing, or unhappy experiences. Thus the theorist who builds actual preferences into her account of welfare will only agree that experiential fear is harmful on the contingent condition that it is actually dispreferred.²³⁹ I have specified O_{R+A} and O_S accordingly. What of the theorist who looks to fully informed or otherwise idealized preferences, in lieu of or in addition to actual preferences? Given the conceptual links between Value and idealized preference, the arguments just advanced to show that experiential fear is disvaluable also demonstrate that this complex of mental states is ideally dispreferred.²⁴⁰ A stipulation that P not only disprefers but ideally disprefers O_{R+A} would be otiose.

In sum, the experiential fear of premature death (including as components both an occurrent probabilistic belief that premature death is likely, and a physically powerful sense of distress occasioned by this belief) is harmful, at least if actually

234. *Id.*

235. *Id.* at 231.

236. HURKA, *supra* note 49, at 9-51.

237. *Cf. id.* at 44, 139-40 (suggesting that emotions are valuable in virtue of the beliefs and desires but not the feelings that they incorporate).

238. *See, e.g.,* SHER, *supra* note 49, at 225-26; SUMNER, *supra* note 8, at 194-95.

239. On the need to incorporate actual preferences within welfare accounts, see Adler & Posner, *supra* note 2, at 200-02.

240. *See supra* text accompanying notes 177-79.

dispreferred. O_{R+A} is worse than O_S with respect to all four of the now-familiar dimensions of welfare.²⁴¹ Nor should this conclusion be surprising to lawyers and legal scholars, since the law has long recognized fear as a welfare setback. Consider the tort of assault. At common law, the infliction of emotional distress without physical injury was not generally tortious,²⁴² but the infliction of fear was—at least if that infliction was intentional and was produced by the threat of immediate physical injury, thereby constituting an assault.²⁴³ Richard Epstein goes so far as to suggest that the fright/shock paradigm—“A frightened B”—is one of four basic paradigms of tortious conduct, along with paradigms that express the causation of harm through force, compulsion, and the creation of a dangerous condition.²⁴⁴ Modern tort doctrine recognizes a tort of “intentional infliction of emotional distress,” including although not limited to the infliction of fear.²⁴⁵ If fear is not harmful, then assaults and fear inflictions ought not be tortious, unless the widely held view that harm is a precondition for tort liability is rejected. A similar point could be made about the compensability of fear as an element of tort damages, under the rubric of “pain and suffering.” If A causes physical injury to B, causing him pecuniary loss, physical pain, and emotional distress such as fear, then B is entitled to compensation from A for all three of these setbacks, not just the pecuniary loss and physical pain.²⁴⁶

What objections might be raised to my claim that risk is harmful as a component of a hybrid psychological state, namely experiential fear? One objection runs as follows. Fear, as defined by Davis and by other philosophers who offer similar

241. See Lisa Heinzerling, *Environmental Law and the Present Future*, 87 GEO. L.J. 2025, 2030-41 (1999) (describing the dread-reduction benefits of environmental regulation). Heinzerling argues that “reducing risk is itself a benefit, separate and apart from the prevention of illness and death,” in part because “a reduction in risk carries with it a decrease in dread and its debilitating effects on individuals and communities.” *Id.* at 2030. But see JUDITH JARVIS THOMSON, *THE REALM OF RIGHTS* 249-71 (1990) (arguing that belief-mediated distress is not harmful).

242. See 2 FOWLER HARPER ET AL., *THE LAW OF TORTS* § 9.1 (2d ed. 1986) (noting “the older rule that there could be no recovery for mere mental distress,” with the exception of assault).

243. See 1 *id.* §§ 3.4-.5 (3d ed. 1996).

244. See RICHARD A. EPSTEIN, *A THEORY OF STRICT LIABILITY: TOWARD A REFORMULATION OF TORT LAW* 22-46 (1980).

245. See 2 HARPER ET AL., *supra* note 242, §§ 9.1-.7.

246. See 4 *id.* § 25.10.

accounts, is a hybrid of a cognitive state, namely *uncertainty*—the subject's uncertainty whether he, or something else he cares about, will be harmed—and an affective state.²⁴⁷ But people can be uncertain without being at risk.²⁴⁸ Risk, in the Bayesian sense, is a state of uncertainty, or partial belief, that can be assigned a precise number satisfying the axioms of the probability calculus.²⁴⁹ For various reasons, it might be impossible to assign P's partial belief this sort of number.²⁵⁰ Therefore, risk is not an essential component of fear. If P is *uncertain* whether E will cause his premature death, and this uncertainty causes P involuntary arousal and unhappiness, then P is fearful of dying prematurely whether or not his uncertainty is measurable on a numerical scale.

My response to this objection is as follows. According to Davis's definition, fear has an essential, cognitive component. *Some kind* of uncertainty is a necessary part of fear;²⁵¹ mere arousal and unhappiness is not sufficient. The cognitive component of fear might take different forms. It might take the form of P's probabilistic belief, or it might take the form of a partial belief of P's that is not a probabilistic belief, where (for whatever reason) P's belief cannot be assigned a precise number satisfying the probability calculus. By claiming that P's risk of premature death is harmful as a component of P's fear of premature death, I do *not* mean to suggest that the fearful P who possesses a probabilistic belief concerning premature death is worse off than the fearful P whose uncertainty about premature death cannot be assigned a

247. See *supra* text accompanying notes 223-28.

248. See generally MICHAEL SMITHSON, *IGNORANCE AND UNCERTAINTY: EMERGING PARADIGMS* 41-151 (1989) (discussing both Bayesian probability and other approaches to conceptualizing uncertainty).

249. See *supra* text accompanying notes 73-76.

250. If probability numbers are assigned along with the assignment of a "utility function," it might emerge that P's preferences are not sufficiently well behaved to be represented by such a function. If probability numbers are assigned by asking what odds P would demand for various sorts of bets, it might emerge that P is just not the betting type—or that he is, but that his betting odds are radically incoherent, and cannot be easily adjusted to satisfy the probability calculus. If probability numbers are assigned by asking P to express his degree of uncertainty on a zero-one scale, he might be unwilling or unable to do that, or able only to generate numbers that, once more, are radically inconsistent with the axioms of probability. See *supra* text accompanying notes 74-76 (discussing different methods for assigning Bayesian probabilities).

251. See Davis, *supra* note 216, at 289-90.

number. O_{R+A} may well be just as bad for P as O_{R+A} , where O_{R+A} and O_{R+A} are identical but for the fact that in O_{R+A} P is nonprobabilistically uncertain whether E will result in his premature death. However, it remains true that fear (including risk as a component) is harmful, since O_{R+A} is worse for P than O_S . Perhaps a more precise formulation of my "risk is harmful as part of fear" claim would be this: The fear of premature death is harmful; fear essentially includes a cognitive component, namely uncertainty; and this cognitive component is realized, inter alia, by a state of probabilistic belief, i.e., by risk in the Bayesian sense.

A second objection to my "risk is harmful as part of fear" claim is more troubling: Whatever might be the cognitive component of fear, that component is irrelevant as far as welfare is concerned. In O_{R+A} , P possesses a partial belief that E will result in his premature death, and he also experiences a negative affective state (analyzed by Davis as "involuntary arousal and unhappiness"²⁵²). Imagine O_A , where P experiences that very same affective state, but lacks the partial belief he possesses in O_{R+A} . O_A is just as bad for P as O_{R+A} —or so the critic of my account might argue. What makes fear bad is its affective component, not its cognitive component. Although O_{R+A} is worse for P than O_S , it is worse than O_S *solely in virtue of* the affective states ("involuntary arousal and unhappiness") that differentiate the two outcomes. Risk, or uncertainty more generally, is no essential part of the welfare story.

A number of philosophers have pointed out that a person can indeed experience the affective state characteristic of fear without possessing the characteristic cognitive state, namely uncertainty.²⁵³ Consider the phobic person: she is aroused and made unhappy by the object of her phobia, even though she believes quite definitely that the object will not harm her. Or consider the person who is "scared" by a scary movie. She might experience all the feelings of the truly frightened, yet be completely certain that nothing depicted in the movie will actually occur.

A heavy smoker may believe an article in a medical journal giving statistics on the correlation of smoking with cancer, but that belief may not scare her in the way watching an autopsy on a smoker with

252. *Id.* at 302.

253. The seminal article here is Kendall L. Walton, *Fearing Fictions*, 75 J. PHIL. 5, 6-10 (1978). For more recent discussions, see Deigh, *supra* note 223, at 835-42; and John Morreall, *Fear Without Belief*, 90 J. PHIL. 359 (1993).

lung cancer would. Indeed, watching a *fictional* scene in a movie of such an autopsy might well be able to scare her more than *believing* the statistics would. Some "mere thoughts" about danger . . . are better able to cause fear than some beliefs about danger. . . . David Hume gave the example of how, as we look over a cliff, we can become afraid by imagining falling to our death: "[W]e tremble on the brink of a precipice, tho' we know ourselves to be in perfect security, and have it in our choice whether we will advance a step farther."²⁵⁴

Although I agree that the distress and arousal characteristic of fear might occur without a concomitant belief, I would deny that fear is harmful solely in virtue of its affective component. Phobia about premature death may be a bad thing, but full-blooded fear of premature death is *worse*. O_s is our baseline outcome, where P is neither distressed and aroused about the prospect of premature death, nor has a probabilistic belief that premature death will occur. O_A is our affectively charged outcome, where P is distressed and aroused at that prospect, even though he is quite certain it will not occur. O_{R+A} is the outcome where P fully fears premature death, since he experiences distress and arousal as a result of the risk, i.e., his probabilistic belief that death will occur. Then, O_A might be worse for P than O_s , but in any event O_{R+A} is worse for P than O_A . The harmfulness of fear is partly grounded in its characteristic cognitions; it is not solely grounded in fear's characteristic affects.

What warrants my claim that O_{R+A} is worse for P's welfare than O_A ? Consider that P's distress in O_{R+A} is a warranted response to his belief state; while in O_A P recognizes, or would recognize if he considered the matter, that he lacks any epistemic basis for these feelings. Wholehearted distress is worse, *ceteris paribus*, than ambivalent distress. It is clear that the wholehearted distress/ambivalent distress distinction satisfies Experience, Integration, and (if wholehearted distress is dispreferred) actual Preference, so the only real sticking point is Value and its close cousin, fully informed Preference. Here I would suggest that our inability to discount the frightening situation as riskless—the fact that there is genuine uncertainty about whether harm will occur—is one essential part of fear's characteristic disvalue.

Consider the parallel with pain. P's feeling of pain, coupled with a belief that the pain results or may result from some permanent damage to his body, is worse for P than the

254. Morreall, *supra* note 253, at 365 (quoting DAVID HUME, A TREATISE OF HUMAN NATURE 445 (L.A. Selby-Brigge ed., Oxford 1888) (1739)).

very same pain coupled with assurance that his bodily integrity is intact. Pain plus a belief that bodily integrity has been damaged amounts to a kind of *suffering*;²⁵⁵ pain itself does not, except in special cases where pain is chronic or very severe. Similarly, P's distress and arousal plus a belief that harm may occur amounts to *fear*, and fear is worse, *ceteris paribus*, than P's distress and arousal coupled with a firm belief that everything is safe.²⁵⁶

In closing, I should note that the hybrid account of risk's harmfulness, presented in this section, might be broadened to include other emotions besides fear. Arguably, there are a range of negative emotions that are welfare setbacks, that essentially include uncertainty as a cognitive component, and that P's probabilistic belief concerning premature death will give rise to, when combined with a suitable negative affect. For example, that belief might make P anxious about premature death, rather than fearful of premature death, assuming anxiety and fear have the same cognitive component but differ in the way they feel.²⁵⁷ If so, the arguments presented in this section carry over to the case of anxiety. P's anxiety about premature death will be a welfare setback for him if

255. See Eric J. Cassel, *The Nature of Suffering and the Goals of Medicine*, 306 NEW ENG. J. MED. 639, 640-41 (1982) (distinguishing between pain and suffering).

256. Crucial to my claim here is the "*ceteris paribus*" clause: Wholehearted distress is worse, *ceteris paribus*, than ambivalent distress. We may be tempted to think that phobia is just as bad as fear, or even worse, by imagining the *intense* distress and arousal that phobics often feel. We may also be tempted to think that phobia or other epistemically ungrounded affective states are worse than fear states by imagining that these states are accompanied by *frustration*—that is, by higher-order feelings of distress occasioned by the fact that the subject is feeling first-order distress without any reason. Clearly, intense phobia may be worse than low-level fear; and phobia plus second-order distress may be worse than fear without second-order distress. But these comparisons are irrelevant to the claim I am advancing that the cognitive component of fear contributes to its harmfulness, in other words that O_{R+A} is both worse than O_S and worse than O_A . O_{R+A} is an outcome in which P's negative feelings are just as intense as in O_A , and which is otherwise identical to O_A except for the belief state which characterizes P in O_{R+A} but not in O_A . Hold fixed the intensity of feeling, and the emotional background, and it becomes quite plausible that fear is worse than phobia, the feelings one experiences during scary movies, and other such epistemically ungrounded affective states.

257. Cf. Davis, *supra* note 216, at 300, 302 (noting that "[f]ear is closely related to *anxiety*" and delineating affective and cognitive distinctions between these two emotions: "anxiety requires arousal but not unhappiness; and the arousal may result from any fear, not just from the fear of harm").

dispreferred, and P's probabilistic belief concerning premature death will contribute to the harmfulness of anxiety (that is, anxiety will be worse than a state in which P lacks both the belief and anxious affect *and* worse than a state in which P is certain that premature death will not ensue, but nonetheless experiences the negative feelings associated with anxiety). Risk in the Bayesian sense is harmful as part of any welfare-reducing emotional complex, including but not necessarily limited to fear, that is predicated on the subject's uncertainty about harmful occurrences such as her own death.

C. OTHER BAYESIAN ACCOUNTS: THIRD-PERSON BELIEFS AND HYPOTHETICAL FIRST-PERSON BELIEFS

The Bayesian account just considered conceptualizes the "risk" of premature death as the actual, probabilistic belief of the very person whose death is at issue. Arguably, that person's probabilistic belief is itself a harm, without more; and in any event that belief may constitute the cognitive component of fear, which is a harmful package of mental states. This Bayesian account is an actual, first-person account—by contrast with hypothetical, first-person accounts and third-person accounts. Hypothetical, first-person accounts are Bayesian accounts that conceptualize the "risk" of premature death as the hypothetical, probabilistic belief of the very person whose death is at issue. Third-person accounts conceptualize the "risk" of premature death as the actual or hypothetical probabilistic belief of *other persons* concerning the premature death of the person "at risk" of that misfortune. Neither hypothetical, first-person accounts, nor third-person accounts, construe "risk" in a way that makes the risk of premature death a welfare setback.

Consider, to begin, the hypothetical, first-person account of risk. O_R , the risky world, is here construed as a world identical to O_S but for the fact that E occurs and P hypothetically would possess—under appropriate conditions—the numerically high probabilistic belief that E will result in his premature death.²⁵⁸ Is O_R worse for P than O_S ? O_S and O_R do not differ with respect to P's actual mental states; thus, the supervenience requirement, standardly invoked by Experientialists, is not satisfied here. Integration is doubtful. For *any* event, P could

258. See Ramsey, *supra* note 76, at 67 (distinguishing between actual and hypothetical degrees of belief).

be brought to have all manner of hypothetical beliefs or other mental states concerning the event, given the right specification of the hypothetical condition. If the occurrence of some event that hypothetically would prompt P to have a particular reaction, under some condition, is taken as a welfare setback for P, then *every* event harms P—which is absurd, and erases the Integration requirement. If the occurrence of some event that hypothetically would prompt P to have a particular reaction, under a specified condition, is taken as a welfare setback for P, then we will need to understand why the specified condition is special—more intimately connected to P's welfare than other hypothetical conditions. As far as I am aware, the only hypothetical condition that philosophers have singled out as specially connected to P's welfare is the condition of full information.²⁵⁹ At this juncture, the problem of determinism reemerges. If the connection between macroscopic events and human death is deterministic, only those persons who will actually die as a result of such events are "put at risk" by them in the hypothetical full-information Bayesian sense.

As for Value, Value theorists who argue that knowledge possession is a welfare value, and false beliefs a welfare setback, look to *actual* beliefs, not hypothetical beliefs.²⁶⁰ And value theorists who argue that fear is a welfare setback look to actual fears—actual beliefs packaged with actual affect—not hypothetical fears.

The arguments just advanced strongly suggest that Bayesian risk in the hypothetical, first-person, sense is not a welfare setback. They also carry over, in part, to the case of Bayesian risk in the third-person sense. Third-person accounts divide into two subcategories: accounts that look to the *hypothetical* probabilistic beliefs of some third person concerning P's premature death, and accounts that look to the *actual* probabilistic beliefs of some third person concerning P's premature death. If P's own hypothetical probabilistic beliefs concerning his premature death are not welfare-reducing for him, then a fortiori the hypothetical, probabilistic beliefs of some third party concerning P's premature death are not welfare reducing for P.

We are left with Bayesian accounts that construe P's "risk"

259. See, e.g., SMITH, *supra* note 179, at 155-61; sources cited *supra* note 45.

260. See sources cited *supra* notes 217-18.

of premature death as the *actual* probabilistic beliefs of some third person concerning that eventuality. Is risk, in this sense, harmful for P? The answer depends upon who the third person is. Risk might be construed as the actual probabilistic beliefs of ordinary members of P's society; but "risk" on this construal is vulnerable to all the objections leveled against the frequentist view of risk as conventional relative frequencies plus more. The ordinary members of P's society may not have convergent probabilistic beliefs about the harmfulness of a given event; even if they do, such beliefs are not value reducing for P, because they neither constitute the best gauge of his security, nor detract from his full personhood. Further, P would not prefer under full information that ordinary members of his society lack such beliefs; and Integration fails because the beliefs and fears of the ordinary members of P's society do not, without more, fall inside the boundary of P's own life. Shifting from the actual probabilistic beliefs of "ordinary" third persons, to the actual probabilistic beliefs of those persons who are expert in some specified domain, does not solve any of these problems—certainly not the last three.²⁶¹

A somewhat more promising approach is to identify the relevant third parties as those persons with whom P is linked by ties of love and affection—his friends and family—or a subcategory of this group.²⁶² However, the "friends and family" account, like the others just considered, runs up against the problem that the third persons whose probabilistic beliefs are defined as P's "risk" may not have convergent beliefs. Consider that one friend of P's might ascribe a high probability to the prospect of E causing P's premature death, while another friend might ascribe that a low probability. There are also conceptual difficulties with the "friends and family" account. Assume P has a single, best friend Q, such that Q's beliefs or at least fears are harms for P himself. Let us say that P faces a

261. Cf. William Meadow & Cass R. Sunstein, *Statistics, Not Experts*, 51 DUKE L.J. 629, 630-31 (2001) (arguing that the legal system should be cautious in relying on expert judgments of risk).

262. In short, P's "risk" of premature death could be equated with the actual, probabilistic beliefs of his circle of intimates. Some Integrationists might insist that *no* third party's beliefs or fears can constitute a welfare setback for P, but this seems too strong. The fact that P's good friend, spouse, or child is scared that P will die prematurely can—intuitively—constitute a harm for P. Experientialists, too, may balk here, but we might resolve this objection by stipulating that P becomes aware of the beliefs and fears of his loved ones.

high, low, or intermediate risk of premature death from event E depending on whether Q has a numerically high, low, or intermediate probabilistic belief that E will result in P's death. The conceptual problem here is that we have lost our placeholder for P's own beliefs concerning his premature demise. We will need to construct some artificial term for those first-person beliefs—since it is those beliefs (not Q's) that explain P's actions, and since those beliefs (not just Q's) constitute a welfare setback for P. Another conceptual oddity is that P ends up being harmed by his "risk" of premature death (Q's beliefs) only if P is an altruist. If P does not care about Q's beliefs or fears, they do not harm P. But surely, the risk of premature death—like death itself—should, if harmful, be a harm that can befall altruists and egoists alike.

The conceptually simpler course, and one that does much less violence to our pretheoretical understanding of "risk," is to equate the Bayesian "risk" to P with *his own* actual probabilistic beliefs. The actual beliefs of certain other persons (friends and family) may be harmful to P, but those beliefs are not, properly speaking, "risks." They constitute a third and additional kind of harm that can result if fatal activities are left unregulated—call this harm flowing from third-party beliefs—where the first is the harm to those who die prematurely, and the second is the harm to those at risk in the Bayesian sense of dying prematurely, i.e., those who actually possess a probabilistic belief (and resultant fear) that they will die prematurely.

V. IMPLICATIONS

Death is a harm to the person who dies. The risk of death, on a Bayesian account that equates risk with actual, first-party probabilistic beliefs, is a harm to the person at risk—at least where those beliefs are packaged together with appropriate affective states and preferences and create a condition of fear, anxiety, or some such negative emotion. The risk of death, on a Bayesian account that equates risk with hypothetical, first-party probabilistic beliefs, or with third-party probabilistic beliefs, is *not* a harm to the person at risk. The risk of death, on a frequentist account that equates risk with statistical, conventional, physical, or other such relative frequencies, is *not* a harm to the person at risk.²⁶³

263. By "harm," as already explained, I mean to denote a setback to

These propositions are the fruits of the analytic labors undertaken in Parts II, III and IV of this Article. To summarize this analysis even more compactly: risk regulation properly rests on what might be called a Hybrid (Bayesian) View of risk, death, and harm, not a Risk View, a Death View, a Null View, or a Hybrid (Frequentist) View.

Is the Risk of Death a Harm?

		<i>No</i>	<i>Yes</i>
Is Death a Harm?	<i>No</i>	Null View	Risk View (Frequentist and/or Bayesian)
	<i>Yes</i>	Death View	Hybrid View (Frequentist and/or Bayesian)

What, concretely, are the implications for risk regulation of the Hybrid (Bayesian) View? Answering this question in detail would take another article, or more. But the reader is entitled to wonder whether the philosophical claims defended in the main body of this Article have genuine relevance for the decisions of real risk regulatory agencies, such as the Consumer Product Safety Commission (CPSC), the Environmental Protection Agency (EPA), the Food and Drug Administration (FDA), the National Highway Traffic Safety Administration (NHTSA), the Occupational Safety and Health Administration (OSHA), and others. I submit that the Hybrid (Bayesian) View has very large implications for the actual practices of risk regulators, and in this Part I will sketch what those implications are.

A. COST-BENEFIT ANALYSIS

Decisions by risk regulators are, increasingly, a product of cost-benefit analysis (CBA).²⁶⁴ Federal risk regulators now

lifetime well-being in the ex post sense: a feature of a possible life history that makes that life history worse for the subject than alternatives. *See supra* text accompanying notes 101-09, 147-48.

264. On CBA, see generally Adler & Posner, *supra* note 2. For a brief overview of the use of CBA in the federal government, see *id.* at 169-76. For more detailed empirical studies on how federal agencies use CBA and what its effects are—both econometric studies and qualitative studies—see Matthew D.

standardly engage in CBA when faced with large-scale regulatory choices, those with substantial economic impacts.²⁶⁵ Systematic federal use of CBA dates to 1981 when the Reagan administration issued Executive Order 12,291, requiring that most regulations issued by federal executive agencies satisfy a cost-benefit test (in contexts where that test was statutorily permissible) and directing agencies to submit written cost-benefit analyses of "major rules" to the President's Office of Management and Budget (OMB).²⁶⁶ The Bush administration kept Order 12,291 in force. The Clinton administration replaced it with Executive Order 12,866, which retains the central features of the earlier order, namely a general cost-benefit test for regulation by executive agencies and OMB review of major rules.²⁶⁷ Specifically, Order 12,866 requires that executive agencies adhere to the following principle, "to the extent permitted by law":

Each agency shall assess both the costs and the benefits of [a proposed] regulation and, recognizing that some costs and benefits are difficult to quantify, propose or adopt a regulation only upon a reasoned determination that the benefits of the . . . regulation justify its costs.²⁶⁸

Where an agency is considering a "significant regulatory action," centrally a rule whose annual costs exceed \$100 million, the agency must prepare a written cost-benefit analysis of the proposed rule, for review by OMB.²⁶⁹

Adler, *The Positive Political Theory of Cost-Benefit Analysis: A Comment on Johnston*, 150 U. PA. L. REV. 1429, 1450 n.46 (2002).

265. See, e.g., Richard D. Morgenstern, *Conducting an Economic Analysis: Rationale, Issues, and Requirements*, in *ECONOMIC ANALYSES AT EPA: ASSESSING REGULATORY IMPACT* 25, 25 (Richard D. Morgenstern ed., 1997) [hereinafter *ECONOMIC ANALYSES AT EPA*] (noting that the EPA has conducted more than 100 formal cost-benefit analyses since the issuance of Executive Order 12,291); Robert W. Hahn et al., *Assessing Regulatory Impact Analyses: The Failure of Agencies to Comply with Executive Order 12,866*, 23 HARV. J.L. & PUB. POL'Y 859, 861 (2000) (studying 48 formal cost-benefit analyses issued for environmental, health, and safety regulations between 1996 and 1999).

266. Exec. Order No. 12,291, 3 C.F.R. 127 (1982), *reprinted in* 5 U.S.C. § 601 (1988).

267. Exec. Order No. 12,866, 3 C.F.R. 638 (1994), *reprinted in* 5 U.S.C. § 601 (2000). For a comparison of the two executive orders, see Peter M. Shane, *Political Accountability in a System of Checks and Balances: The Case of Presidential Review of Rulemaking*, 48 ARK. L. REV. 161, 174-92 (1995).

268. Exec. Order No. 12,866 § 1(b)(6), 3 C.F.R. 638, 639 (1993), *reprinted in* 5 U.S.C. § 601.

269. *Id.* §§ 3(f), 6(a)(3); see also Unfunded Mandates Reform Act of 1995, 2 U.S.C. § 1532 (2000) (requiring the preparation of cost-benefit analyses for

Despite the recent change of administration, Executive Order 12,866 remains governing law.²⁷⁰ The current presidency has reaffirmed the general commitment to CBA expressed first by Order 12,291 and now by 12,866. To be sure, these orders apply a cost-benefit test to agency rules only where that test is statutorily permissible. Some federal statutes, in the area of risk regulation, *do* clearly preclude CBA.²⁷¹ Instead, these statutes give priority to health and safety benefits over economic and other costs. For example, the Occupational Safety and Health Act stipulates that OSHA shall regulate toxic substances in the workplace so as to ensure, to the extent feasible, that no worker suffers "material impairment of health or functional capacity."²⁷² The pre-1990 version of section 112 of the Clean Air Act, which covers "hazardous air pollutants," instructed the EPA to set permissible emissions of each such pollutant at a level that "provides an ample margin of safety to protect the public health."²⁷³

On the other hand, a substantial number of federal statutes permit or require risk regulators to use CBA. One or two do so quite explicitly;²⁷⁴ others do so by using balancing language that courts and agencies quite straightforwardly have interpreted as permitting or requiring CBA.²⁷⁵ For example, the Toxic Substances Control Act allows the EPA to ban or restrict chemicals that the agency has "a reasonable basis to conclude . . . will present an unreasonable risk of injury to health or the environment";²⁷⁶ the Consumer Product Safety

significant regulatory actions).

270. See Exec. Order No. 13,258, 67 Fed. Reg. 9385 (Feb. 26, 2002) (making amendments to Executive Order 12,866, mainly with respect to the Vice President's role).

271. For a general discussion of the legal status of CBA under current environmental and other social regulatory statutes, see Cass R. Sunstein, *Cost-Benefit Default Principles*, 99 MICH. L. REV. 1651, 1663-67 (2001).

272. 29 U.S.C. § 655(b)(5) (2000).

273. 42 U.S.C. § 7412(b)(1)(B) (1988), amended by Pub. L. No. 101-549, § 301 (1990); see Alon Rosenthal et al., *Legislating Acceptable Cancer Risk from Exposure to Toxic Chemicals*, 19 ECOLOGY L.Q. 269, 300-04 (1992) (describing the original version of § 112 of the Clean Air Act).

274. See, e.g., Safe Drinking Water Act of 1974, 42 U.S.C. § 300g-1(b)(6) (2000).

275. See Sunstein, *supra* note 271, at 1666-67 (discussing statutes that clearly or arguably permit or require CBA).

276. 15 U.S.C. § 2605(a) (2000).

Act employs a similar “unreasonable risk of injury” standard;²⁷⁷ and the Federal Insecticide, Fungicide, and Rodenticide Act instructs the EPA to license pesticides that will not cause “unreasonable adverse effects on the environment.”²⁷⁸ Finally, numerous federal statutes in the area of risk regulation are ambiguous with respect to CBA; they neither clearly permit it, nor clearly preclude it.²⁷⁹ Consider the Food, Drug, and Cosmetic Act provision governing pharmaceuticals, which requires them to be “safe” and “effect[ive],”²⁸⁰ or the provision of the same statute applicable to foods, which prohibits foods containing “poisonous or deleterious” substances “unless the quantity of such substance . . . does not ordinarily render it injurious to health.”²⁸¹ The emerging trend in the case law is to read ambiguous statutes as permitting or even requiring CBA.²⁸² Cass Sunstein has gone so far as to suggest that there is now a canon of statutory interpretation favoring the use of CBA unless Congress’s intent to prohibit CBA is clearly expressed.²⁸³

It bears emphasis that CBA is not confined to the federal government. A recent survey of state regulatory review structures notes that some states “strongly encourage the use of cost-benefit analysis and cost-effectiveness analysis” by state administrative agencies, and that “[o]ver half of the states require agencies to assess the economic impact of all proposed rules.”²⁸⁴

277. *Id.* § 2056(a).

278. 7 U.S.C. § 136a(c)(5) (2000).

279. See Sunstein, *supra* note 271, at 1666 (“A large number of statutes ask agencies to ‘take into consideration’ various factors, including cost, in addition to the principal factor to which the statute draws the agency’s attention . . .”).

280. 21 U.S.C. § 355(d) (2000).

281. *Id.* § 342(a)(1).

282. See Sunstein, *supra* note 271, at 1668-82; see also Cass R. Sunstein, *Regulating Risks After ATA*, 2001 SUP. CT. REV. 1, 27-30 (reading *Whitman v. American Trucking Ass’n*, 531 U.S. 457 (2001), to embrace the principle that agencies are allowed to consider costs absent a clear statutory prohibition). But see Richard J. Pierce, Jr., *The Appropriate Role of Costs in Environmental Regulation* 19-22 (Geo. Wash. U.L. Sch., Pub. L. Res. Paper No. 36) available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=301479 (last visited Feb. 11, 2003) (reading *American Trucking* to suggest that agency consideration of costs is disfavored, and concluding that the opinion gives no useful guidance as to the applicability of this “anti-cost” canon versus the “pro-cost” canon previously recognized by lower courts).

283. Sunstein, *supra* note 271, at 1668.

284. Robert W. Hahn, *State and Federal Regulatory Reform: A Comparative Analysis*, in COST-BENEFIT ANALYSIS: LEGAL, ECONOMIC, AND

Why *should* federal and state agencies, including risk regulatory agencies, rely on CBA to assess their options? CBA is traditionally linked to economic efficiency, more precisely to the concept of Kaldor-Hicks efficiency.²⁸⁵ Eric Posner and I have argued that this linkage is mistaken, since Kaldor-Hicks efficiency lacks moral significance.²⁸⁶ Rather, CBA is a workable proxy for the criterion of overall well-being.²⁸⁷ The regulatory option favored by CBA will typically (if not always) be the option that maximizes overall well-being. And overall well-being *does* have moral significance, if not conclusively so. It is one moral factor, among others, that bears upon what agencies morally ought to do, all things considered.

This view helps explain what “costs” and “benefits” are, for purposes of CBA. A regulatory outcome involves a “cost,” relative to some alternative, if it involves a setback to some aspect of well-being. A regulatory outcome involves a “benefit,” relative to some alternative, if it improves human well-being in some way. And distinct categories of “costs” or “benefits” represent different aspects or *dimensions* of welfare—different ways in which human welfare can be advanced or impeded.

CBA comes in two versions: monetized and nonmonetized.²⁸⁸ Both versions start by describing the costs and benefits of possible regulatory outcomes along various dimensions of welfare, with different, easily measurable units used for each dimension. For example, it might be calculated that an air pollution regulation will prevent 360 deaths, 6800 cases of chronic bronchitis, 1300 hospital admissions for congestive heart failure, and 106,000 lost days of work; will improve visibility, measured in kilometers of visual range, by 11%, for an affected population of 50 million; will impose

PHILOSOPHICAL PERSPECTIVES 37, 38 (Matthew D. Adler & Eric A. Posner eds., 2001) [hereinafter COST-BENEFIT ANALYSIS].

285. See Adler & Posner, *supra* note 2, at 187-94 (discussing the conventional defenses of CBA).

286. *Id.* at 190-91.

287. This view of CBA is presented and defended in Matthew D. Adler & Eric A. Posner, *Implementing Cost-Benefit Analysis When Preferences Are Distorted*, in COST-BENEFIT ANALYSIS, *supra* note 284, at 269, 272-80; Adler, *Incommensurability*, *supra* note 96, at 1371-83; Adler & Posner, *supra* note 2, at 194-238.

288. See, e.g., LESTER B. LAVE, *THE STRATEGY OF SOCIAL REGULATION: DECISION FRAMEWORKS FOR POLICY* 17-19, 23-25 (1981) (describing these two variants of CBA); Adler & Posner, *supra* note 2, at 177-87, 233-36 (describing “CBA,” meaning monetized cost-benefit analysis, and “qualitative direct multidimensional assessment,” meaning nonmonetized cost-benefit analysis).

financial costs of \$100 million on the regulated companies; and will result in 20,000 fewer jobs. The monetized version of CBA then reduces each cost and benefit of the regulatory proposal to a dollar figure, and calculates a net monetary cost or benefit for the proposal. The nonmonetized version of CBA relies on intuition and judgment rather than formal commensuration to balance the divergent costs and benefits of the proposal. No attempt is made to reduce all these welfare impacts to a common scale, be it a dollar scale or some other.²⁸⁹

The analysis presented in the earlier parts of this Article has large and quite obvious implications for CBA, in both the monetized and nonmonetized versions. The analysis clarifies what the *appropriate* dimensions of CBA are—what *properly* counts as a cost or benefit.

Imagine that some risky activity or substance E produces (1) X premature deaths; (2) Y cases in which persons are at substantial risk, in the actual first-party Bayesian sense, of dying prematurely as a result of E, i.e., they actually believe to a high probabilistic degree that they will die prematurely as a result of E; and (3) Z cases in which persons have a high risk, in some frequentist sense, of dying prematurely as a result of E.²⁹⁰ An agency is choosing between the status quo, with the

289. For a series of case studies that nicely illustrate how agencies actually do both monetized and nonmonetized CBA, see ECONOMIC ANALYSES AT EPA, *supra* note 265.

290. This is a simplified case, similar to those presented in welfare economics and CBA textbooks, where the outcome of the proposal is certain (in this case with respect to the number of adverse events, if not the identities of the persons involved). See, e.g., ROBIN BOADWAY & NEIL BRUCE, WELFARE ECONOMICS 226-34 (1984) (discussing monetary measurement of individual welfare effects under uncertainty after discussing valuation under certainty); E.J. MISHAN, COST-BENEFIT ANALYSIS: AN INFORMAL INTRODUCTION 375-419 (4th ed. 1988) (discussing uncertainty at the end of the text). This simplified case is designed to show how the views defended in this Article concerning the harmfulness of death, Bayesian risk, and frequentist risk bear upon the CBA valuation of outcomes. More realistically, a regulatory agency will be considering various possible choices or "options," the possible outcomes of which are uncertain. (For example, the agency might be uncertain about the number of premature deaths prevented by each regulatory option. Perhaps a given option might prevent between zero and 100 deaths, with the agency ascribing frequentist or Bayesian probabilities to each such possible outcome.) In this more realistic case, applying CBA will involve a view about how one "translates" the CBA valuation of possible outcomes into a CBA valuation of the possible choices facing the decision maker—a view about how one moves from the ex post to the ex ante. See *supra* text accompanying notes 95-100. This is an issue which I have deliberately avoided discussing in this Article, given space constraints. See *id.* Still, it seems clear that, on any plausible

activity or substance left unregulated, and a regulatory scheme that will control it—and thereby prevent the X premature deaths, Y cases of first-person Bayesian risk, and Z cases of frequentist risk. According to the Null View of risk regulation, neither the X premature deaths, nor the Y cases of Bayesian risk, nor the Z cases of frequentist risk count as a cost of leaving the status quo in place (or, reciprocally, as a benefit of the regulation). According to the Death View, the X premature deaths count as a cost of leaving the status quo in place (and, reciprocally, as a benefit of the regulation), but the Y cases of Bayesian risk and the Z cases of frequentist risk do not. According to a Risk (Bayesian) View, only the Y cases of Bayesian risk count as a cost of the status quo; according to a Risk (frequentist) View, only the Z cases of frequentist risk count as a cost.

Finally, under the view defended here—the Hybrid (Bayesian) View—both the X premature deaths and the Y cases of first-person, Bayesian risk count as costs of the status quo, for purposes of CBA. Reciprocally, the regulation has *two* benefits: eliminating the deaths and eliminating the Bayesian risk. But the Z cases of frequentist risk are irrelevant to the CBA analysis.

Clearly, the choice between the Null View, Risk Views, Death View, and Hybrid Views has large implications for the proper practice of CBA. How revisionary are these implications? To what extent does CBA, as actually practiced by risk regulatory agencies, deviate from the Hybrid (Bayesian) View? Let us consider first the monetized and then the nonmonetized versions of CBA.

1. CBA: Monetized Version

Faced with a risky substance or activity E that produces X cases of premature death, Y cases in which persons actually believe to a high degree that E will result in their premature deaths, and Z cases in which persons have a high frequentist risk of dying prematurely, the monetized cost of that substance or activity is properly calculated as follows: $\sum_i V_{\text{death},i} + \sum_j V_{\text{Bayesian}}$.

account of the ex post/ex ante “translation,” the CBA valuation of outcomes will be a crucial ingredient in determining the CBA ranking of possible actions facing an uncertain regulator. Thus, I feel quite confident that the choice between the Death, Null, Risk, and Hybrid Views has much relevance for the real-world regulatory application of CBA, and more generally for real-world regulatory choice.

$V_{\text{death},i}$ represents the monetized cost of death for person i , where person i is one of the X persons who will die prematurely as a result of E . $V_{\text{Bayesian-risk},j}$ represents the monetized cost for person j of actually believing to a high degree that E will result in his premature death, where person j is one of the Y persons who possess this belief. $V_{\text{Bayesian-risk},j}$ will presumably depend on what probability person j attributes to premature death. Note that there is no third term, in this formula, for the putative "cost" of the Z cases of high frequentist risk, since on the proper view of welfare frequentist risk is no cost (welfare setback) at all.

How do we calculate $V_{\text{Bayesian-risk},j}$ and $V_{\text{death},i}$? CBA standardly monetizes costs and benefits by asking what the individuals affected would be willing to pay, or accept, for these welfare changes.²⁹¹ This technique is plausibly applied to monetize the cost of first-person, Bayesian risk. If person j believes that the risky event or substance creates a given probability of her premature death, then $V_{\text{Bayesian-risk},j}$ is plausibly the amount that person j would pay to prevent an event with that probability (and with no other welfare impact on her), or would accept in compensation for the occurrence of such an event. Indeed there is now a large econometric and survey literature documenting the amounts that persons are willing to pay or accept in return for small probabilities of premature death.²⁹²

We might think that first-person Bayesian risk is not a harm, taken alone, but is only harmful as a component of fear, anxiety, or cognate affectively laden states. On that view, $V_{\text{Bayesian-risk},j}$ is zero if person j is not fearful of the risky event. If she is fearful, $V_{\text{Bayesian-risk},j}$ is the amount she would be willing to pay to eliminate an event causing her to experience that level of fear (and with no other welfare impact on her), or to accept in compensation for the occurrence of such an event.²⁹³

291. A standard textbook on monetized CBA is MISHAN, *supra* note 290. Numerous other such works are cited in Adler, *supra* note 58, at 249-51 nn.29, 34 & 36-37.

292. Much of this literature is summarized in W. KIP VISCUSI, *FATAL TRADEOFFS* 34-74 (1992). I say "plausibly" because it might be argued that what is being measured here is the ex ante cost of death, not the ex post cost of being under a risk (assuming the latter can be a cost without associated affective states). Whether the two kinds of monetary valuation indeed differ, and if so how to measure the latter rather than the former, cannot be addressed here.

293. There are surely many subtleties in the monetization of fear, which unfortunately I do not have space to discuss here. The issue merits much

Calculating $V_{\text{death},i}$ is much trickier.²⁹⁴ We cannot ask what person i would be willing to pay to avoid premature death or to accept in return for premature death, since in many cases these numbers will be infinite, or bounded only by the person's total wealth. Nor are the financial consequences of person i 's death, e.g., lost wages or income, an appropriate measure of $V_{\text{death},i}$. It is tempting to think that $V_{\text{death},i}$ is simply incalculable²⁹⁵—that death cannot be priced in dollars, for CBA purposes. This temptation should be resisted. CBA is a rough proxy for judgments of overall well-being. The proper method for calculating $V_{\text{death},i}$ takes account of this construal of CBA. $V_{\text{death},i}$ is the amount of money that tracks the effect of person i 's death on overall welfare. It is the amount of money that—used as an input into a CBA analysis—minimizes the deviation between the output of that analysis, and the underlying criterion of overall well-being.

This definition of $V_{\text{death},i}$ is vague, but it can be made more precise. Assume that each person's lifetime well-being, in various outcomes, can be represented by a "utility" function U_i . This is a standard assumption of welfare economists. Further, assume that these "utilities" are interpersonally comparable. In other words, assume that outcome O_1 is better for overall well-being than outcome O_2 if and only if the cumulative utility of O_1 , across the population, is larger than the cumulative utility of O_2 . For welfare effects other than death—welfare effects that are monetized, for purposes of CBA, using the willingness to pay/accept methodology—we can calculate a ratio between the monetized value of the welfare effect and its utility value.

For example, if a new car has a monetary value for some person of \$10,000, and improves his utility by 200 units, then the ratio of monetary value and utility value here is 50. If eliminating annoying noise has a monetary value for some

more scholarly attention than it has received. On the closely related issue of valuing anxiety, see Mordechai Shechter, *Incorporating Anxiety Induced by Environmental Episodes in Life Valuation*, in 2 APPLIED BEHAVIOURAL ECONOMICS 529 (Shlomo Maital ed., 1988); and Bruce H. Smith, *Anxiety as a Cost of Commuting to Work*, 29 J. URB. ECON. 260 (1991). Cf. Loewenstein et al., *supra* note 223, at 267-71 (focusing on the role of affect in explaining behavior under uncertainty).

294. For an overview of the problem of monetizing the value of life, see JOHN KLEINIG, *VALUING LIFE* 145-63 (1991).

295. See John Broome, *Trying to Value a Life*, 9 J. PUB. ECON. 91, 91-95 (1978).

person of \$600, and improves his utility by 10 units, then the ratio of monetary value and utility value for *this* welfare effect is 60. Now, calculate the *average* ratio between monetary value and utility value, for all welfare effects (other than death) priced by CBA. Then $V_{\text{death},i}$ is the loss of utility incurred by person *i*, multiplied by this average ratio.²⁹⁶

Of course, it is very difficult to know what person *i*'s particular loss of utility will be—especially since risk regulators almost never know, in advance, who in particular will die prematurely as a consequence of some risky substance or event. Yet economists and public policy analysts might try to calculate what the typical or average loss of utility is when people die prematurely as a result of the kinds of substances or activities generally targeted by risk regulators, and then multiply this typical utility loss by the ratio described in the preceding paragraph to arrive at a typical or average cost of death, V_{death} . Using this average or typical value, the cost for CBA purposes of an activity or substance *E* that results in *X* premature deaths is: $X \cdot V_{\text{death}}$. And our formula for calculating the monetized cost of an activity that results in *X* premature deaths, *Y* cases of first-person Bayesian risk, and *Z* cases of frequentist risk, becomes:

$$X \cdot V_{\text{death}} + \sum_j V_{\text{Bayesian-risk},j}$$

We are finally in a position to assess the actual practices of risk regulators. The standard approach to monetizing the costs of a risky substance or activity such as *E*—as explained in the economics literature on CBA and now generally employed by regulatory agencies—is to multiply the number of premature deaths resulting from *E* by the so-called value of statistical life (VOSL).²⁹⁷ The current best estimate for VOSL is \$6 million.²⁹⁸

296. See Matthew D. Adler, *Risk, Death, and Time: A Comment on Judge Williams' Defense of Cost-Benefit Analysis*, 53 ADMIN. L. REV. 271, 285-86 (2001) (proposing such an approach to valuing life).

297. This approach derives from the work of E.J. Mishan and Thomas Schelling. See E.J. Mishan, *Evaluation of Life and Limb: A Theoretical Approach*, 79 J. POL. ECON. 687 (1971); T.C. Schelling, *The Life You Save May Be Your Own*, in PROBLEMS IN PUBLIC EXPENDITURE ANALYSIS 127 (Samuel Chase, Jr. ed., 1968). The current literature on VOSL is vast. For good surveys, see, for example, A. MYRICK FREEMAN III, THE MEASUREMENT OF ENVIRONMENTAL AND RESOURCE VALUES: THEORY AND METHODS 314-66 (1993); VISCUSI, *supra* note 292, at 34-74; W. KIP VISCUSI, RATIONAL RISK POLICY 45-68 (1998); M.W. Jones-Lee, *Safety and the Saving of Life: The Economics of Safety and Physical Risk*, in COST-BENEFIT ANALYSIS 290, 290-318 (Richard Layard & Stephen Glaister eds., 2d ed. 1994); David Pearce, *Valuing Risks*, in HANDBOOK OF ENVIRONMENTAL RISK ASSESSMENT AND

Further, as a matter of standard CBA practice, the Bayesian or frequentist "risks" resulting from a hazardous substance or activity are not monetized separately from the premature deaths that the activity or substance causes.²⁹⁹ In other words, the monetized cost of E causing X premature deaths, Y cases of first-person Bayesian risk, and Z cases of frequentist risk would be calculated as: $X \cdot \text{VOSL}$. No separate terms for the Y belief states or Z instances of frequentist risk appear in this formula.

VOSL, in turn, is calculated as follows. A premature death resulting from some substance or activity is seen as imposing a small "risk" of death on each person who is "exposed" (in some sense) to the substance or activity. For example, if one person dies out of an exposed population of 1 million, then the "risk" of death for each person is seen to be 1 in 1 million. If the exposed population is 100,000, then the "risk" of death for each person is seen to be 1 in 100,000. The VOSL is the dollar amount that the members of the population would be willing to pay, in the aggregate, to avoid these "risks." As Kip Viscusi, a leading proponent of the VOSL method, explains,

[I]magine a group of 10,000 people, such as an audience in a stadium. . . . If there was one expected statistical death from this large audience of 10,000 participants, and if each of the participants is willing to pay \$500 to eliminate the risk, then it would be possible to raise \$5 million collectively from this audience to eliminate the one statistical death to the group, thus establishing the value of life.³⁰⁰

On the assumption that (1) the exposed population is large, and that (2) the willingness to pay/accept for risk increases in a linear fashion with the risk level, given small risks, this aggregate amount—the VOSL—will be the same number regardless of the precise population size. "[T]he statistical value of life is the total amount of compensation n workers would require to face one expected death from their group, where n is a large number."³⁰¹

How does the standard formula for monetizing the costs of a risky substance or activity ($X \cdot \text{VOSL}$) compare with the

MANAGEMENT 345, *supra* note 200, at 345-75; and W. Kip Viscusi, *The Value of Risks to Life and Health*, 31 J. ECON. LIT. 1912 (1993).

298. See Ackerman & Heinzerling, *supra* note 33, at 1-2.

299. See ECONOMIC ANALYSES AT EPA, *supra* note 265 (providing case studies illustrating how the EPA performs monetized CBA); sources cited *supra* note 297 (explaining VOSL methodology).

300. VISCUSI, *supra* note 297, at 46.

301. Viscusi, *supra* note 297, at 1930.

approach recommended here, namely, $X \cdot V_{\text{death}} + \sum_j V_{\text{Bayesian-risk},j}$? First, $X \cdot \text{VOSL}$ should *not* be taken as an approximation for the second term in my equation, $\sum_j V_{\text{Bayesian-risk},j}$. This second term, $\sum_j V_{\text{Bayesian-risk},j}$, asks how much those who *actually* have certain probabilistic beliefs about E would pay to prevent it. For example, if E is a toxic release which occurs in a population of 1 million and will kill 1 person, but only 3000 are actually aware of E, then at most 3000 people can be at risk in the actual, first-person Bayesian sense. Only a person in this small group of 3000 can actually possess a high probabilistic belief that his or her death will occur prematurely as a result of E. Y, in this example, is no greater than 3000. But the VOSL calculation, as explained above, assumes that *everyone* in the population of 1 million is "at risk" of dying prematurely. The term "risk" is—in effect—used by VOSL proponents to mean a *hypothetical* first-person Bayesian risk, not an actual first-person Bayesian risk. We calculate VOSL by asking what each person in the population would, hypothetically, pay to prevent E, were she to believe that her probability of dying prematurely as a result of E were the number of premature deaths (1 in the example here) divided by the population size (1 million)—and then aggregating the amounts. This produces a dollar sum very different from the sum of willingness to pay/accept among those actually aware of, or afraid of, the hazard. For instance, while $X \cdot \text{VOSL}$ in the example under consideration is \$6 million, $\sum_j V_{\text{Bayesian-risk},j}$ might well be as low as \$18,000!³⁰²

Even if $X \cdot \text{VOSL}$ were a decent approximation for $\sum_j V_{\text{Bayesian-risk},j}$, this would hardly vindicate CBA as currently practiced. The appropriate formula for calculating the costs of a risky substance or activity E is: $X \cdot V_{\text{death}} + \sum_j V_{\text{Bayesian-risk},j}$. If $X \cdot \text{VOSL}$ does approximate $\sum_j V_{\text{Bayesian-risk},j}$, the upshot would be that CBA as currently practiced dramatically underestimates the costs of risky substances and activities, and dramatically understates the benefits of risk regulation. Why? Consider that, if the approximation just described were to obtain, then CBA as currently practiced would include a cost term roughly equaling $\sum_j V_{\text{Bayesian-risk},j}$ (namely $X \cdot \text{VOSL}$), but would lack any cost term at all for $X \cdot V_{\text{death}}$.

A better argument in defense of current practice runs as

302. A \$6 million VOSL implies a willingness to pay \$6 to avoid a 1 in 1 million risk. If this amount is used to calculate $\sum_j V_{\text{Bayesian-risk},j}$ in the example at hand, where 3000 people are actually aware of the risky event, that equals \$18,000.

follows. For complicated reasons that cannot be elaborated here, it turns out that VOSL may be a decent approximation for V_{death} ³⁰³—on the premise that people price small risks of death in accordance with the standard economic theory of individual decision making known as “expected utility” theory. Further, it seems reasonable to assume that $\sum_j V_{\text{Bayesian-risk},j}$ will often be small. After all, many individuals are unaware of (let alone fear) the dangerous substances and activities that regulatory agencies target. So $X \cdot \text{VOSL}$ could be a decent approximation for the formula advocated here: $X \cdot V_{\text{death}} + \sum_j V_{\text{Bayesian-risk},j}$.

Does this line of argument succeed? It is beyond the scope of this Article to evaluate whether VOSL really is a decent approximation for V_{death} . Let me note, however, that the premise on which the argument rests—that expected utility theory successfully predicts how individuals price small risks of death—is highly contestable. Much literature in experimental psychology suggests that expected utility theory is a poor predictor of individual behavior, including pricing behavior.³⁰⁴ VOSL is predicated on econometric studies of the incremental wages that workers demand to work in hazardous jobs, and on other such evidence of what individuals are willing to pay or accept in return for fatality risks. If this pricing behavior is not driven by expected utility calculations, but instead by social norms, by altruistic considerations, or by other psychological mechanisms inconsistent with expected utility theory, then the approximate equivalence of VOSL and V_{death} is undermined.³⁰⁵

To sum up, there is considerable reason to think that the current approach to monetizing the costs of death and risk, in the area of risk regulation, is incorrect. The correct formula ($X \cdot V_{\text{death}} + \sum_j V_{\text{Bayesian-risk},j}$) is conceptually quite different from the currently employed formula ($X \cdot \text{VOSL}$), and it is far from clear whether the two formulae are even approximately equal.

303. See Adler, *supra* note 296, at 285-86.

304. For an accessible introduction to this literature, see PLOUS, *supra* note 68, at 79-161.

305. The point that individuals do not actually price fatality risks in accordance with expected utility theory is developed in PETER DORMAN, *MARKETS AND MORTALITY: ECONOMICS, DANGEROUS WORK, AND THE VALUE OF HUMAN LIFE* 142-50 (1996). It should be noted that expected utility theory has different variants. One variant—the variant required to make VOSL a decent approximation for V_{death} —presumes self-interested behavior, and is therefore inconsistent with altruistic behavior.

2. CBA: Nonmonetized Version

Agencies engaged in CBA pursuant to Executive Order 12,866, or in other contexts, do not invariably reduce the "costs" and "benefits" of proposed rules to monetary amounts. Indeed, a recent study found that a majority of cost-benefit analyses submitted by health, safety, and environmental agencies to OMB failed to monetize at least some of the welfare impacts delineated therein.³⁰⁶ Because there is no scholarly work documenting in any systematic way how agencies actually perform nonmonetized CBA—how they actually identify and commensurate costs and benefits that are measured on diverse scales and not rescaled in terms of dollars or some other common metric—it would be overreaching for me to make claims about the "normal" or "average" practices of agencies here. It appears, however, that at least some risk-regulatory agencies engaged in nonmonetized CBA count the elimination of risk in the *frequentist* sense as a benefit of regulatory proposals, distinct from the prevention of premature death.³⁰⁷ This represents an important divergence between nonmonetized and monetized CBA, since as explained above the standard methodology for monetized CBA is to calculate and monetize the premature deaths prevented by a regulatory proposal, and to ignore frequentist risk.

Why might frequentist risk become a factor in agency decision making? A good answer requires some attention to the role of so-called "risk assessment" within the modern administrative state.³⁰⁸ A risk assessment is, in effect, a systematic and carefully structured description of the fatalities, fatality risks, and other health or safety setbacks caused by some hazard.³⁰⁹ Risk-regulatory agencies that use CBA (either

306. See Hahn et al., *supra* note 265, at 866-70.

307. See *infra* text accompanying notes 319-22.

308. On that role, see CARL F. CRANOR, *REGULATING TOXIC SUBSTANCES* 103-51 (1993); *QUANTITATIVE RISK ASSESSMENT IN REGULATION* (Lester Lave ed., 1982); Howard Latin, *Good Science, Bad Regulation, and Toxic Risk Assessment*, 5 YALE J. REG. 89, 95-126 (1988); Rosenthal et al., *supra* note 273, at 295-322; Mark Eliot Shere, *The Myth of Meaningful Environmental Risk Assessment*, 19 HARV. ENVTL. L. REV. 409, 417-68 (1995); Symposium, *Risk*, 63 U. CIN. L. REV. 1531 (1995); and Symposium, *Risk Assessment in the Federal Government*, 3 N.Y.U. ENVTL. L.J. 251 (1995) [hereinafter Symposium, *Risk Assessment*].

309. See generally VINCENT T. COVELLO & MILEY W. MERKHOFFER, *RISK ASSESSMENT METHODS: APPROACHES FOR ASSESSING HEALTH AND ENVIRONMENTAL RISKS* (1993); *HANDBOOK OF ENVIRONMENTAL RISK*

the monetized or nonmonetized version) to evaluate policy alternatives frequently employ risk assessment as an initial step in that analysis. Indeed, the OMB guidance document issued pursuant to Executive Order 12,866 states that the documentation for major rulemakings should include a risk assessment.

The risk assessment should generate a credible, objective, realistic, and scientifically balanced analysis [of risks to health, safety, and the environment]; present information on hazard, dose-response, and exposure . . . ; and explain the confidence in each assessment by clearly delineating strengths, uncertainties, and assumptions, along with the impacts of these factors on the overall assessment.³¹⁰

Risk assessments are used in other contexts as well, for example when federal agencies implement statutes that preclude CBA but still require or permit a systematic agency review of the fatalities, illnesses, etc., caused by hazardous substances or activities.³¹¹

The techniques of risk assessment have become quite standardized—in part because risk analysis has become a distinct academic discipline with its own norms and practices, including a conventional framework for risk assessment,³¹² and in part because that framework was articulated and codified by the National Research Council in a highly influential 1983 study, *Risk Assessment in the Federal Government: Managing the Process*.³¹³ This study, known as the Red Book, divided risk assessment into the following steps: hazard identification, dose-response assessment, exposure assessment, and risk characterization.³¹⁴ “Hazard identification,” crudely, means

ASSESSMENT AND MANAGEMENT, *supra* note 200, at 9-245; HUMAN AND ECOLOGICAL RISK ASSESSMENT: THEORY AND PRACTICE 1-366 (Dennis J. Paustenbach ed., 2002); THE RISK ASSESSMENT OF ENVIRONMENTAL AND HUMAN HEALTH HAZARDS: A TEXTBOOK OF CASE STUDIES 1-328 (Dennis J. Paustenbach ed., 1989).

310. U.S. OFFICE OF MGMT. & BUDGET, ECONOMIC ANALYSIS OF FEDERAL REGULATIONS UNDER EXECUTIVE ORDER 12866 at § III.A.4, available at <http://www.whitehouse.gov/omb/inforeg/riaguide.html> (Jan. 11, 1996).

311. See Rosenthal et al., *supra* note 273, at 295-322 (describing the EPA's use of risk assessment in a variety of statutory contexts).

312. On the history of risk assessment, see Dennis J. Paustenbach, *Retrospective on U.S. Health Risk Assessment: How Others Can Benefit*, 6 RISK: HEALTH, SAFETY & ENV'T 283, 285-87 (1995).

313. NAT'L RESEARCH COUNCIL, RISK ASSESSMENT IN THE FEDERAL GOVERNMENT: MANAGING THE PROCESS (1983).

314. See *id.* at 19-20; see also NAT'L RESEARCH COUNCIL, SCIENCE AND JUDGMENT IN RISK ASSESSMENT 4-6 (1995) (summarizing the Red Book framework).

identifying some substance or activity for which there is scientific evidence of a causal link to fatalities and other health or safety setbacks. "Dose-response assessment" means quantifying that causal link. Paradigmatically, risk assessment focuses on toxic substances (such as air or water pollutants, pesticides, toxic food additives or naturally occurring food constituents, workplace carcinogens or other workplace toxins, or substances used in consumer products) and the dose-response assessment takes the form of a graph or mathematical equation showing the relationship between the dose of the toxic substance and the percentage of individuals exposed to that dose who develop an adverse effect. "Exposure assessment" means characterizing the pattern of exposures to the toxic substance or other hazard that would occur if the status quo were left unregulated and if various regulatory alternatives were put in place. Finally, "risk characterization" means integrating the dose-response assessment and the exposure assessment so as to generate a *prediction* of the fatalities, illnesses, etc., that would occur in the status quo outcome and in the various outcomes of the regulatory alternatives under consideration.³¹⁵

Risk assessments typically generate two kinds of predictions, insofar as fatalities are concerned: (1) a prediction of the total number of premature deaths that would occur as a result of different regulatory choices; and (2) a prediction of the distribution of the *frequentist risk* of premature death that would occur as a result of these choices. This frequentist risk is what risk assessments typically call "individual risk." As one textbook on risk assessment explains,

Most [risk assessments] include several common measures of individual and societal risk, in particular:

- *Individual risk*, which is the probability of a specified individual dying prematurely as a result of exposure to the risk agents. . . .
- *Individual risk contours* show the geographical distribution of individual risk
- *Maximum individual risk* is the individual risk to the person experiencing the highest risk in the exposed population. . . .
- Various measures of *societal risk*, such as . . . the expected number

315. On the meaning of "hazard identification," "dose-response assessment," "exposure assessment," and "risk characterization," see, for example, COVELLO & MERKHOFFER, *supra* note 309, at 27-29; and NAT'L RESEARCH COUNCIL, *supra* note 313, at 19-20.

of fatalities as a function of location or population subgroup³¹⁶

To see why the “individual risk” numbers generated by traditional risk assessment are frequentist, not Bayesian risks, and to understand why risk assessment readily lends itself to this kind of prediction, consider the following simple example. Imagine that our “hazard identification” has identified some chemical compound that has accumulated in a waste dump and that is causally linked to an inevitably fatal cancer. One million people drink water that is contaminated, to some degree, by the dump. Our “dose-response assessment” determines that the relationship between the amount of the compound ingested (over a lifetime) and the percentage of a group ingesting this particular dose that would incur fatal cancer as a result can be represented by the following graph.

Figure 1: Dose-Response Curve for Carcinogen in Waste Dump

Finally, our “exposure assessment” determines that, absent regulation, the one million people drinking the contaminated water would in fact be exposed to the following lifetime doses.

316. COVELLO & MERKHOFFER, *supra* note 309, at 231.

Exposure Level

	Very Low (1 gram)	Low (5 grams)	Moderate (15 grams)	High (25 grams)
Number of Individuals	500,000	400,000	80,000	20,000

Combining the "dose-response assessment" and the "exposure assessment," we can predict that absent regulation (1) the total number of premature deaths caused by the compound would be 12; and (2) the Very Low Exposure, Low Exposure, Moderate Exposure, and High Exposure segments of the population would face the following "individual risks."

Exposure Level

	Very Low	Low	Moderate	High
Individual Risk	1 in 250,000	1 in 100,000	1 in 20,000	1 in 10,000

These "individual risks" are *frequentist* risks. They are a function of the individual's dosage of the toxic substance concerned—indeed, can be read directly off the dose-response curve—and equal the frequency with which persons (or persons of a particular type) exposed to that dose die prematurely as a result. In this example, the dose-response curve is quite crude and shows the percentage of a group consisting of *all persons exposed to a particular dose of the toxin* that would die prematurely as a result of the toxin. A more refined dose-response curve might display the percentage of a group consisting of all persons (1) with characteristics $C_1 \dots C_n$, and (2) exposed to a particular dose of the toxin that would die prematurely as a result. (For example, a risk assessment could have one dose-response curve for adults, and another for children.) Whether the dose-response curve is simple or refined, the "individual risks" generated by risk assessments have nothing to do with degrees of belief (actual or hypothetical), but instead are calculated by subsuming the given individual, P, within some class of persons exposed to the toxic substance, and determining what percentage of that class

would die prematurely because of it.³¹⁷

How are the dual predictions of the typical risk assessment—a prediction concerning total premature deaths, and a prediction concerning the *frequentist risk* of premature death incurred by various members of the population—employed by agencies conducting monetized or nonmonetized CBA? In the case of monetized CBA, the prediction concerning frequentist risk is typically ignored. As I have already explained, monetized CBA takes the number of premature deaths prevented by a regulatory proposal, multiplies that number by VOSL (\$6 million) to determine the monetized value of lifesaving, and moves on; there is no separate term for the Bayesian or frequentist risks that would be eliminated by the proposal.³¹⁸ However, at least some agencies conducting nonmonetized CBA *do* give weight to the risk assessor's prediction concerning frequentist risk. Alon Rosenthal, George Gray, and John Graham surveyed the use of risk assessment by the EPA with respect to carcinogenic chemicals, and identified several instances in which that agency, in the course of implementing "balancing" statutes, took account of "individual risk" in making regulatory choices.³¹⁹ For example, the authors found that the Toxic Substances Control Act (TSCA) was administered as follows:

The OTS [the office within EPA that administers the TSCA] has not formalized its risk management criteria, in part because TSCA requires a discretionary, judgmental balancing of numerous factors. Generally, the OTS believes that situations in which lifetime individual risk is less than 10^{-6} , or in which population risk (i.e., the number of expected cases of cancer per year) is less than one, do not warrant the attention of the agency. . . . [T]he OTS considers both individual risk and population risk [total deaths] in deciding whether to regulate. For example, in deciding to phase out the use of asbestos in almost all products over a seven-year period, the OTS concluded that the rule would prevent 200 cases of cancer and would also relieve certain highly exposed individuals of a 10^{-3} lifetime excess cancer risk. It cited both of these factors as justifying the costs caused by the suspension.³²⁰

Similarly, the authors found that the EPA, in applying the Federal Insecticide, Fungicide and Rodenticide Act (FIFRA)—another balancing statute—placed some emphasis on

317. See, e.g., COVELLO & MERKHOFFER, *supra* note 309, at 151-68 (discussing dose-response models).

318. See *supra* text accompanying notes 297-99.

319. Rosenthal et al., *supra* note 273, at 295-322.

320. *Id.* at 308-09 (citations omitted).

"individual risk." "The [EPA under FIFRA] tends to set acceptable risk levels for the food-consuming population within or below the range of 10^{-5} to 10^{-6} , while it tends to accept occupational risks that are less than 10^{-4} to 10^{-5} ."³²¹ Several further examples where the EPA has counted the elimination of frequentist risk as a benefit of regulation, for purposes of nonmonetized CBA, are provided in a recent study describing the use of economic analysis by that agency.³²²

The cases I have just described are not cases in which (1) the frequentist risk information simply recapitulates information about the number of premature deaths in different outcomes, and is not given additional weight by the agency; or (2) frequentist risks are the "bridging" probabilities that an agency, uncertain about which outcomes will result from its various possible choices, employs to move from the "ex post" ranking of outcomes to the "ex ante" ranking of choices.³²³ Rather, it appears, these are cases in which agencies take the level of frequentist risk to be an independent dimension of welfare—a welfare-relevant feature of a possible outcome that has independent negative weight, quite apart from the occurrence of premature death, in determining the ranking of that outcome.

This practice should be ended. Agencies that use nonmonetized CBA to evaluate regulatory proposals should *not* give independent weight to frequentist risk in determining the CBA ranking of outcomes. CBA is a proxy for overall well-being—genuine "costs" and "benefits" track genuine dimensions of welfare. Frequent risk, however, is *not* a genuine welfare setback, and the elimination of frequentist risk is not a genuine welfare improvement. Part III of this Article argued that neither a high *statistical* probability of some person's premature death, nor a high *physical* probability, nor a high *conventional* probability, diminishes that person's well-being.³²⁴

The frequentist risk numbers that are predicted by risk

321. *Id.* at 306 (citations omitted).

322. See Peter Caulkins & Stuart Sessions, *Water Pollution and the Organic Chemicals Industry*, in ECONOMIC ANALYSES AT EPA 87, *supra* note 265, at 108, 110-11, 115; Ronnie Levin, *Lead in Drinking Water*, in ECONOMIC ANALYSES AT EPA 205, *supra* note 265, at 225; Mahesh Podar et al., *Municipal Sewage Sludge Management*, in ECONOMIC ANALYSES AT EPA 365, *supra* note 265, at 371.

323. See *supra* text accompanying notes 95-100.

324. See *supra* Part III.B-D.

assessment documents and sometimes figure within agency decision making—so-called “individual risks”—are in fact either *statistical* or *conventional* probabilities, not physical probabilities. “Individual risk” is not calculated by subsuming the individual within a reference class sharing all of his causally relevant characteristics. Rather, as I explained above, the reference class is much cruder, often consisting simply of all persons exposed to a particular dose of the toxic substance under assessment,³²⁵ and in any event will be heterogeneous with respect to some characteristics that affect individual responses to that dose. Consider the crude calculation that P, who consumes a lifetime dose of 1 gram of some carcinogen C, has an “individual risk” of 1 in 20,000 of dying as a result, because the frequency with which members of the reference class of “all individuals who consume a lifetime dose of 1 gram of C” die prematurely is 1/20,000. This 1 in 20,000 probability will be P’s statistical probability of dying prematurely from C if and only if the narrowest P-containing reference class for which reliable frequency information is available is the crude class defined in terms of lifetime dosage. That might not be the case; better information could be available. Even so, the 1 in 20,000 probability is a kind of conventional probability. P’s lifetime dosage represents a conventionally salient casual feature, albeit not the only causal feature or the only one about which statisticians might be informed.

Whether statistical, conventional, or both, the frequentist probabilities of premature death that are predicted by risk assessments should not be used—within the context of monetized or nonmonetized CBA—as a partial determinant of the agency’s ranking of outcomes and therewith its ranking of the regulatory choices possibly resulting in those outcomes. A change in someone’s level of frequentist risk, be it statistical, conventional, or even physical, is neither a welfare setback nor an improvement and therefore is not a cost or benefit for CBA purposes.³²⁶

325. See *supra* text accompanying note 317.

326. Nor can the frequentist risks predicted by risk assessments be reinterpreted as hypothetical Bayesian risks or as actual third-party Bayesian risks (for example, the probabilistic beliefs of experts concerning the likely death of persons exposed to toxic substances), and thereby given a legitimate outcome-ranking role within CBA. As I demonstrated in Part IV.C, hypothetical Bayesian risk and third-party Bayesian risk, like frequentist risk, are not welfare setbacks.

B. RISK-RISK ANALYSIS

I have shown that the Hybrid (Bayesian) View of risk has large implications for cost-benefit analysis. What are the implications if we shift from CBA to competing frameworks for regulatory choice? I have argued at length, elsewhere, that CBA is the appropriate framework,³²⁷ as have numerous other scholars,³²⁸ and CBA is now widely practiced by administrative agencies.³²⁹ But it is still premature to suggest, as does Cass Sunstein in a recent article, that we have moved from "first generation" to "second generation" debate about CBA—where "second generation" debate concerns "how (not whether) to engage in cost-benefit analysis."³³⁰ CBA remains controversial, particularly among academics.³³¹ One competing framework, for purposes of risk regulation, is the framework known as "risk-risk" analysis (RRA).³³² RRA focuses on the health and safety effects of regulatory options, rather than on other types of costs or benefits.

One way of stating the objective [of RRA] is that society desires to minimize the adverse health effects associated with a given food such as bacon [or with other toxic substances or activities]. Thus society would permit nitrite [a carcinogen] in bacon if the improvement in the health of consumers from botulism protection exceeded the decrement in health from the risk of cancer. Yet it is evident that the direct risk-risk framework takes only the first step of considering the health of the person consuming the food. People are also associated with the production and distribution of food; society [under RRA] desires to minimize the adverse health effects associated with producing as well as consuming bacon³³³

RRA has a robust version and a weak version. The robust version of RRA ignores all nonhealth effects of regulatory options: financial costs, employment, aesthetic effects,

327. See sources cited *supra* note 287.

328. See, e.g., COST-BENEFIT ANALYSIS, *supra* note 284. This book is an anthology of papers on CBA, many of which defend the practice.

329. See *supra* text accompanying notes 264-84.

330. Sunstein, *supra* note 271, at 1655-56.

331. See Adler & Posner, *supra* note 2, at 170-72 nn.12-18 (citing sources).

332. See, e.g., LAVE, *supra* note 288, at 15-17; Frank B. Cross, *When Environmental Regulations Kill: The Role of Health/Health Analysis*, 22 *ECOLOGY L.Q.* 729 (1995); John D. Graham & Jonathan Baert Wiener, *Confronting Risk Tradeoffs*, in *RISK VERSUS RISK 1*, *supra* note 5, at 1-41; Cass R. Sunstein, *Health-Health Tradeoffs*, 63 *U. CHI. L. REV.* 1533, 1538-52 (1996); W. Kip Viscusi, *Risk-Risk Analysis*, 8 *J. RISK & UNCERTAINTY* 5, 5-7 (1994).

333. LAVE, *supra* note 288, at 16.

convenience, and so on.³³⁴ The weak version of RRA accords priority to health and safety—it gives these dimensions of well-being more weight than they would have within CBA—but does not wholly ignore nonhealth effects.³³⁵ The case for RRA is either a moral case (namely, that agencies are morally obliged *not* to maximize well-being, but instead to give special weight to health and safety)³³⁶ or an institutional case (that RRA turns out to be a more administrable proxy for overall well-being than CBA, given the limited resources, information, and capabilities of agencies).³³⁷

What is the legal status of RRA? Some statutes governing risk-regulatory agencies *are* best interpreted as requiring those agencies to use RRA rather than CBA.³³⁸ Other statutes use general balancing language without clearly choosing between RRA and CBA.³³⁹ Proponents of RRA would argue that these statutes, too, should be read as setting forth RRA (not CBA) as the framework for agency choice.

The view of risk, death, and harm defended in this Article has substantial relevance for RRA. Consider the choice between the Null View, the Death View, the Hybrid (Frequentist) View, and the Hybrid (Bayesian) View. On the Null View, RRA should focus on illness or injury, not on death itself or the risk of premature death. The fact that a policy reduces the incidence of, say, emphysema, headaches, or paraplegia, would legitimately count in favor of the policy.³⁴⁰ However, the fact that a policy reduces fatalities or fatality

334. See, e.g., *id.* at 15-18 (contrasting risk-risk framework, which on Lave's construal "precludes consideration of non-health effects," and "risk-benefit" framework, which does consider non-health effects).

335. See Mark Geistfeld, *Reconciling Cost-Benefit Analysis with the Principle that Safety Matters More than Money*, 76 N.Y.U. L. REV. 114, 116-28 (2001) (arguing for a modified version of CBA that gives extra weight to safety interests but does not ignore other interests).

336. See *id.*

337. See Adler & Posner, *supra* note 2, at 229-33.

338. See Sunstein, *supra* note 271, at 1664-65.

339. See *id.* at 1666-67.

340. For discussions of valuation of morbidity, see generally Mark Berger et al., *Framework for Valuing Health Risks*, in VALUING HEALTH FOR POLICY 23, *supra* note 29; Richard Clemmer et al., *Household Health Production, Property Values and the Value of Health*, in VALUING HEALTH FOR POLICY 105, *supra* note 29; Robert Fabian, *The Quality Approach*, in VALUING HEALTH FOR POLICY 118, *supra* note 29; Donald Kenkel et al., *Contingent Valuation of Health*, in VALUING HEALTH FOR POLICY 72, *supra* note 29; and Donald Kenkel, *Cost of Illness Approach*, in VALUING HEALTH FOR POLICY 42, *supra* note 29.

risks would not. *Morbidity* is policy relevant, but *mortality* is not. On the Death View, mortality as well as morbidity is policy relevant, but the risk of mortality is not. On the Hybrid (Frequentist) View, mortality, morbidity, and, arguably, the frequentist risk of mortality or morbidity would all be legitimate factors for RRA.

Finally, and interestingly, the Hybrid (Bayesian) View has the same implications for RRA as the Death View. Death is a welfare setback, *pace* the Null View. If so, it is exceedingly difficult to see why RRA should be structured as a technique that aims at, or gives priority to, the reduction of morbidity, but does not aim at or give priority to the reduction of mortality. Death is, typically, a more serious welfare setback than injury and illness. Death is the more extreme result of the processes that produce injury and illness, namely, injurious events and disease. Death, like injury and illness, reduces the subject's well-being in (broadly speaking) the same way, namely, by intruding on her physical integrity. All these considerations suggest that RRA should focus on reducing death as well as reducing injury and illness.

Fear and Bayesian risk are different. These *are* welfare setbacks, according to the Hybrid (Bayesian) View, but are typically less serious setbacks than injury, illness, and death; and fear and Bayesian risk do not, except in very unusual cases, represent a bodily dysfunction or failure of physical integrity. Bayesian risk is merely a kind of belief; fear is a kind of emotion. Well-functioning humans have beliefs and experience emotions. It is hard to see what moral or institutional considerations could counsel agencies to (1) ignore or give lower priority to financial costs, unemployment, convenience, ecosystem disruption, and other serious nonhealth effects, but (2) aim at or give high priority to emotional and epistemic well-being. Finally, since risk in the frequentist sense is not a welfare setback, frequentist risk should be no more relevant to RRA than it is to CBA.

In short, assuming that there are good moral or institutional grounds for agencies to engage in RRA rather than CBA, agencies should aim at or give priority to the reduction of premature death (along with illness and injury) but should *not* aim at or give priority to the reduction of risk in the frequentist or Bayesian sense.³⁴¹

341. By "aim at," here, I mean to maintain the focus on outcomes that has

I do not know how this recommendation squares with the actual practices of agencies engaged in RRA. However, it is inconsistent with the view taken by John Graham and Jonathan Wiener in a major scholarly work on RRA, *Risk Versus Risk*. Wiener and Graham argue that all of the following considerations should be relevant to risk-risk tradeoffs: "magnitude [of risk], degree of population exposure, certainty, type of adverse outcome, distribution, and timing."³⁴² They elaborate as follows:

Magnitude (Probability) of Risk. Recall that risk is defined as the probability of an adverse outcome. . . . If the adverse outcomes are similar or identical (for example, fatal cancer vs. fatal cancer) . . . the relative probabilities of the adverse outcomes become critical.

. . . [I]t is important for decisionmakers to recognize the difference between a lifetime probability of, say, 1 in 100 and a lifetime probability of 1 in 100,000. . . .

. . . .

Size of Population. If we hold constant the probability of adverse outcome to a specified individual . . . , the relative size of the exposed populations is an important consideration. Statistics such as aggregate cancer incidence (sometimes called "population risk" in contrast to "individual risk") are useful in capturing the importance of the size of an exposed population.³⁴³

In short, Graham and Wiener claim that *both* "individual risk" (in the frequentist sense) *and* "population risk" (the total number of deaths, illnesses, or injuries) should figure in RRA.³⁴⁴ Imagine that in outcome O_1 one cancer death occurs, out of a population of 100 unknowingly exposed to the carcinogen; and in outcome O_2 two cancer deaths occur, out of a population of 10,000,000 unknowingly exposed to the carcinogen. According to Graham and Wiener, as I read them, the policy maker here has *prima facie* reason to favor the outcome, O_2 , with more deaths, since the other outcome exposes individuals to a much higher frequentist risk. Indeed, given the stark disparity between the two outcomes on the risk

characterized this entire Article. Neither frequentist risk nor Bayesian risk should bear on the RRA ranking of outcomes. Death should bear on the RRA ranking of outcomes; and so the Bayesian or frequentist probabilities of different regulatory choices resulting in outcomes characterized by death are potentially relevant—relevant "ex ante"—to the RRA analyst. See *supra* text accompanying notes 95-100.

342. Graham & Wiener, *supra* note 332, at 30.

343. *Id.* at 30-31.

344. For a similar claim, see Adam M. Finkel, *Comparing Risks Thoughtfully*, 7 RISK 325, 342-44 (1996).

dimension, and the smaller disparity on the death dimension, it seems likely that—under the Graham and Wiener approach—it would be best, all things considered, for outcome O_2 to obtain. O_1 has one premature death, and 100 cases of a very high (1 in 100) frequentist risk of premature death, while O_2 has two premature deaths, and ten million cases of a trivially low (2 in 10 million) frequentist risk of premature death.

One upshot of the Hybrid (Bayesian) View of risk defended in this Article is that the Graham and Wiener conception of RRA—the leading conception within the scholarly literature—is mistaken. In applying RRA to the hypothetical outcomes just described, the policy maker should unequivocally rank O_1 , the outcome with fewer premature deaths, over outcome O_2 ; he has no reason whatsoever to give a higher ranking to O_2 . More generally, although “distribution,” “timing,” “type of adverse outcome,” and perhaps other considerations could properly figure in RRA along with “degree of population exposure” (the number of actual deaths, illnesses and injuries), risk in the frequentist sense should not.

C. HEALTH AND SAFETY THRESHOLDS

Numerous regulatory statutes set health or safety *thresholds*; they direct agencies to determine whether a particular level of some substance, or amount of some activity, would be “safe” or “unsafe,” “healthy” or “unhealthy,” or cognate terms. These thresholds arise in a variety of statutory contexts. Sometimes, the health or safety thresholds function as statutory *goals* and seem entirely to displace a balancing analysis. In such instances, the agency is seemingly enjoined by statute to eliminate health or safety threats above the threshold, regardless of countervailing costs, financial or other.³⁴⁵ For example, the Food, Drug and Cosmetic Act requires food additives to be “safe.”³⁴⁶ Section 112 of the Clean Air Act, until its amendment in 1990, instructed the EPA to set

345. See Sunstein, *supra* note 271, at 1663-64. I say “seemingly” because, arguably, even statutes that facially require agencies to maximize health or safety should be interpreted to incorporate some cost considerations, absent very clear language to the contrary. See *Natural Res. Def. Council, Inc. v. EPA*, 824 F.2d 1146, 1163-66 (D.C. Cir. 1987) (en banc) (interpreting the pre-1990 version of section 112 of the Clean Air Act to authorize the EPA to give limited consideration to cost and technological feasibility). But see *Whitman v. Am. Trucking Ass'ns, Inc.*, 531 U.S. 457, 464-71 (2001) (reading section 109(b) of the Clean Air Act to preclude the EPA from considering costs).

346. 21 U.S.C. § 348(c)(3)(A) (2000).

emissions standards for hazardous air pollutants at a level that would "protect the public health" with an "ample margin of safety."³⁴⁷

In other cases, statutes create health or safety thresholds that function as goals and preclude CBA- or RRA-type balancing, but place a constraint of technological or economic "feasibility" on regulation; since regulations that are so costly as to cause widespread bankruptcy among regulated firms would be economically "infeasible," and since the elimination of safety risks where that necessitates a complete cessation in the production of some good or service is often seen as technologically "infeasible," this "feasibility" constraint amounts to a crude type of balancing.³⁴⁸ Section 6(b)(5) of the Occupational Safety and Health Act is a good example of this type of statute. OSHA is directed to set a standard for each toxic substance in the workplace that "most adequately assures, to the extent feasible, . . . that no employee will suffer material impairment of health or functional capacity."³⁴⁹ Another example is the pre-1996 version of the Safe Drinking Water Act. The EPA was required first to set nonenforceable "maximum contaminant level goals" (MCLGs) for water pollutants, which are concentrations at which no adverse human health effects are believed to occur, and then to promulgate enforceable standards, "maximum contaminant levels," which were to be set as close to the MCLGs as is "feasible with the use of the best technology, treatment techniques, and other [available] means."³⁵⁰

Finally, statutory health or safety thresholds are sometimes integrated with a CBA- or RRA-type balancing analysis, in more or less complex ways. A statute might set a threshold, preclude regulation of substances or activities below the threshold,³⁵¹ and instruct the agency to engage in CBA- or RRA-type balancing to determine whether to regulate

347. 42 U.S.C. § 7412(b)(1)(B) (1988), *amended by* Pub. L. No. 101-549, § 301 (1990); *see* Rosenthal et al., *supra* note 273, at 300-04 (describing the original version of section 112 of the Clean Air Act).

348. *See* Sunstein, *supra* note 271, at 1665-66.

349. 29 U.S.C. § 655(b)(5) (2000).

350. *See* Rosenthal et al., *supra* note 273, at 309-13 (describing the pre-1996 version of the Safe Drinking Water Act) (citing 42 U.S.C. § 300g-1(b)(4), (5) (1988)). The Act has subsequently been amended to include a cost-benefit provision. *See supra* note 274.

351. *See* Sunstein, *supra* note 271, at 1664 (discussing statutes that require "significant" risk as a prerequisite for regulation).

substances or activities above the threshold. The threshold, in this context, might be seen as the outcome of a legislative balancing sensitive to deliberation costs: The benefits of regulating activities or substances below the threshold are likely to be so small that they do not justify the (substantial) administrative deliberation costs involved in full-fledged CBA or RRA.³⁵² Consider section 3(8) of the Occupational Safety and Health Act, which generally authorizes OSHA to issue standards that are "reasonably necessary or appropriate to provide safe or healthful employment and places of employment."³⁵³ This statutory language arguably requires or permits OSHA to employ CBA or RRA in regulating dangerous activities or conditions—as opposed to toxic substances, which are covered not only by section 3(8) but also by the anti-balancing language of section 6(b)(5)³⁵⁴—and indeed has been so interpreted by several circuit courts.³⁵⁵ However, section 3(8) has also been read by a plurality of the Supreme Court as setting a health/safety threshold. In *Industrial Union Department v. American Petroleum Institute* (the *Benzene* case), Justice Stevens famously determined that OSHA lacked the power to regulate "insignificant" health or safety risks.³⁵⁶

By empowering the Secretary to promulgate standards that are "reasonably necessary or appropriate to provide safe or healthful employment and places of employment," the Act implies that, before promulgating any standard, the Secretary must make a finding that the workplaces in question are not safe. But "safe" is not the equivalent of "risk-free." There are many activities that we engage in every day—such as driving a car or even breathing city air—that entail some risk of accident or material health impairment; nevertheless, few people would consider these activities "unsafe." Similarly, a workplace can hardly be considered "unsafe" unless it threatens the workers with a significant risk of harm.

Therefore, before he can promulgate *any* permanent health or safety standard, the Secretary is required to make a threshold finding that a place of employment is unsafe—in the sense that significant

352. On the deliberation costs involved in CBA, see Richard D. Morgenstern & Marc K. Landy, *Economic Analysis: Benefits, Costs, Implications*, in ECONOMIC ANALYSES AT EPA 455, *supra* note 265, at 460-62; and Paul R. Portney, *The Benefits and Costs of Regulatory Analysis*, in ENVIRONMENTAL POLICY UNDER REAGAN'S EXECUTIVE ORDER 226 (V. Kerry Smith ed., 1984).

353. 29 U.S.C. § 652(8).

354. *Id.* § 655(b)(5).

355. See Cross, *supra* note 332, at 769-70 (discussing section 3(8) of the Occupational Safety and Health Act and relevant case law).

356. 448 U.S. 607 (1980) (plurality opinion).

risks are present and can be eliminated or lessened by a change in practices.³⁵⁷

In short, the statutory regime governing OSHA regulation of dangerous activities and conditions is arguably the following: no regulation of activities or conditions posing "insignificant" safety risks, and cost-justified regulation of activities or conditions posing "significant" safety risks.

A more complicated statutory structure, combining a health/safety threshold with a balancing test, is found in the current version of section 112 of the Clean Air Act, as amended in 1990.³⁵⁸ The Act lists various air pollutants. For new sources of each pollutant, the EPA is required to set the most restrictive emission standard that is achievable under the best, currently utilized technology. If this technology-based standard would still allow a level of pollution posing a non-minimal health risk—if that level exceeds a health/safety threshold—and if a balancing analysis supports lowering the pollution level, then EPA must enact a more restrictive standard.

In sum, health/safety thresholds appear in numerous federal statutes, either as regulatory goals that displace a broader balancing test (perhaps hedged by a feasibility constraint), or as tests of significance that precede balancing, or in some more complicated way. How are these thresholds implemented? What does it mean for some substance or activity to be "unsafe," or to pose a "substantial" or "significant" threat to health or safety, or to rise above a "de minimis" level of danger? The answer: *Agencies often (if not always) construe statutory safety or health thresholds as levels of frequentist risk.* This practice was fueled by the *Benzene* case, where Justice Stevens not only interpreted section 3(8) of the Occupational Safety and Health Act as precluding the agency from regulating "insignificant" workplace hazards, but also suggested that OSHA look at *individual risk* to determine whether this safety/health threshold was exceeded. In a crucial passage, Justice Stevens wrote,

[T]he requirement [under section 3(8)] that a "significant" risk be identified is not a mathematical straitjacket. It is the Agency's responsibility to determine, in the first instance, what it considers to be a "significant" risk. Some risks are plainly acceptable and others

357. *Id.* at 642.

358. 42 U.S.C. § 7412 (2000); see Cross, *supra* note 332, at 770-72 (discussing the structure of the current version of section 112).

are plainly unacceptable. If, for example, the odds are one in a billion that a person will die from cancer by taking a drink of chlorinated water, the risk clearly could not be considered significant. On the other hand, if the odds are one in a thousand that regular inhalation of gasoline vapors that are 2% benzene will be fatal, a reasonable person might well consider the risk significant and take appropriate steps to decrease or eliminate it.³⁵⁹

OSHA now generally uses a 1 in 1000 level of "individual risk" as its test of significance for carcinogens: If some such substance in the workplace creates more than a 1 in 1000 risk of premature death to any worker, that substance exceeds the health/safety threshold of section 3(8) and can be regulated.³⁶⁰ The EPA interpreted the original version of section 112 of the Clean Air Act—which required an "ample margin of safety" for the public—as setting a 1 in 10,000 risk threshold: The permissible amount of a covered air pollutant would need to be set sufficiently low that the "maximally exposed individual" would face less than a 1 in 10,000 risk of premature death from the pollutant.³⁶¹ The current version of section 112 explicitly states that its health/safety threshold should be specified in terms of "individual risk." The EPA must "reduce lifetime excess cancer risks to the individual most exposed to emissions . . . to less than one in one million."³⁶² EPA also uses "individual risk" (either risk to the maximally exposed individual, or some average of individual risk across the population) in specifying safety thresholds under the Clean Water Act, the Safe Drinking Water Act, the "Superfund" statute (CERCLA), and the Resource Conservation and Recovery Act (RCRA).³⁶³ The FDA has at various junctures interpreted the Food, Drug and Cosmetic Act as including a "de minimis" safety threshold and defined that threshold as a 1 in 1 million individual risk of premature death.³⁶⁴

359. *Indus. Union Dep't*, 448 U.S. at 655.

360. See March Sadowitz & John D. Graham, *A Survey of Residual Cancer Risks Permitted by Health, Safety and Environmental Policy*, 6 RISK 17, 24-25 (1995). This valuable survey of agency threshold setting was published more than a half-decade ago; I am not aware of a more recent one. It is possible that some of the agency practices which are described in the Sadowitz and Graham survey, and which I note here and immediately below, have changed since its publication.

361. See *id.* at 25-26; Rosenthal et al., *supra* note 273, at 300-04.

362. 42 U.S.C. § 7412(f)(2)(A); see Cross, *supra* note 332, at 771 (discussing this provision).

363. See Rosenthal et al., *supra* note 273, at 309-22; Sadowitz & Graham, *supra* note 360, at 26-30.

364. See Sadowitz & Graham, *supra* note 360, at 21-24. For a discussion of

The FDA, EPA, OSHA, and other agencies that specify health/safety thresholds in terms of “individual risk” (from 1 in 1000 at the high end, to 1 in 1 million at the low end) implicitly or explicitly understand “individual risk” in frequentist terms. A 1 in x risk of cancer, for these purposes, means a 1 in x risk as predicted by the methodology of risk assessment; but, as I have already explained, that methodology equates “individual risk” with the percentage of some exposed population that dies or becomes ill or injured.

The Hybrid (Bayesian) View of risk, defended in this Article, therefore has quite straightforward and dramatic implications for the pervasive agency practice of setting health/safety thresholds in terms of individual (frequentist) risk. *The practice cannot be justified, and ought not continue.* The fact that an individual’s frequentist risk of premature death lies above or below a given level makes no difference to her well-being. Therefore, it makes no sense for agencies to aim at health/safety thresholds defined in frequentist terms, or to specify the thresholds that trigger balancing in frequentist terms, or to equate the thresholds that play a more complicated statutory role with a frequentist risk of premature death. Consider the case of a health/safety threshold that functions as a regulatory goal and displaces balancing, for example, the original version of section 112 of the Clean Air Act. Assume the EPA sets the health/safety goal for each pollutant at some level of frequentist risk L , say, 1 in 10,000. This means the following: Given two outcomes O_1 and O_2 , such that some individuals in O_1 would have an above- L risk of premature death from the pollutant, while no individuals in O_2 will have an above- L risk of premature death from the pollutant, EPA will prefer O_2 to O_1 . Unless, however, O_1 and O_2 differ with regard to some metric other than frequentist risk (for example, the aggregate number of deaths, or the level of fear), the Hybrid (Bayesian) View of risk defended in this Article implies that the two outcomes are *identical* with respect to well-being. No one is worse off in O_1 than in O_2 , absent some divergence between the outcomes other than the divergence captured by the EPA’s risk threshold L .³⁶⁵

the use of “individual risk” by the Nuclear Regulatory Commission and state agencies to set safety thresholds, see *id.* at 18-21, 31-33.

365. How can O_1 and O_2 differ with respect to the risk-threshold L but not the number of deaths? As a simple example, imagine that in O_1 n persons die out of a small exposed population, while in O_2 n persons die out of a large

In the case of a health/safety threshold that serves as a trigger for balancing, or that plays a yet more complicated statutory role, the specification of that threshold in frequentist-risk terms is vulnerable to the very same objection. Why should the fact that one outcome lies below the risk threshold, while another lies above it, affect the regulatory status of the outcomes? That difference between the outcomes does not, in itself, change anyone's welfare.

One possible response is that while the level of frequentist risk lacks *welfare* significance, it remains morally significant in some other way. For example, safety and health might be understood as "primary goods" or "resources"; a "primary good" or "resource" is a *precondition* for welfare or a means to welfare rather than a part of welfare itself.³⁶⁶ Money is the classic example: Having more money does not itself improve someone's welfare, but rather enables her to realize welfare improvements that would otherwise not be feasible for her. However, given the arguments set forth in Part III of this Article, it is very hard to see how a lower frequentist risk of premature death could be a "resource" analogous to money or, for that matter, actual health and longevity.³⁶⁷ A high frequentist risk of premature death does not (at least if the individual is unaware of that risk) prevent her from living a perfectly good life, by contrast with ill health, hunger, fear, or death; nor can a low risk be traded for welfare improvements, by contrast with money.

Alternatively, a high frequentist risk of premature death for some person could be seen as an infringement of his moral *rights*, rather than as a diminution of his welfare or resources.³⁶⁸ Arguably, rights infringements do not entail

exposed population, with the frequentist risk equaling n divided by the population size. A somewhat more complicated example is as follows: For one toxic substance the "maximally exposed individual" has an above- L risk, but most of the population has a small risk; while for another toxic substance most of the population has a larger risk, still below L , and the "maximally exposed individual" now also has a below- L risk. Further, the probability distributions are such that the regulators predict the same number of deaths to occur as a result of each substance. Examples like these help clarify that frequentist risk and the number of deaths are analytically distinct features of outcomes.

366. For a discussion of "primary goods" or "resources" and citations to the philosophical literature, see Adler & Posner, *supra* note 2, at 211-12.

367. See *supra* Part III.B-D.

368. For an argument that risk imposition is indeed a kind of *wrongdoing* (rights violation), see McCarthy, *supra* note 64, at 208-15. But see Stephen Perry, *Imposing Risk* 48-67 (unpublished paper, on file with author)

welfare or resource reductions. For example, if I trespass on your property without damaging anything, and without harming you, I may still have violated your moral rights.³⁶⁹ Might not I do the same by, say, playing a game of Russian Roulette with you as the unwitting target, or more generally by imposing a high frequentist death risk on you? Perhaps so. The arguments presented in Part III focus on welfare, not rights, and thus do not settle the issue whether risk impositions (in the frequentist sense) can be rights violations. That important issue is beyond the scope of this Article. It seems intuitively clear, however, that rights-violating risk impositions (if they exist) would involve a high frequentist risk plus other elements, for example a lack of voluntariness on the part of the victim. If I sell you a dangerous substance, and tell you that, and you are a competent adult who understands my warning, then your rights have not been violated even though I may have imposed a high frequentist risk of premature death on you. Yet agencies that specify health and safety thresholds in terms of frequentist risk do so without reference to the further elements that would (seemingly) be required to establish a rights violation;³⁷⁰ and they do so in contexts where the statute is focused on welfare not rights.

How should statutory health/safety thresholds be interpreted by regulatory agencies, if not as levels of frequentist risk? This is a complicated question, whose answer may vary depending on the statutory role of the threshold. In the case of a statute that instructs the agency to regulate a particular product or activity by aiming at safety or health, and ignoring cost considerations, one plausible answer is that the agency should set regulatory standards such that *no one* dies prematurely as a result of the product or activity. In other words, the health/safety threshold could be specified as a maximum permissible number of deaths resulting from the product or activity, rather than a maximum "individual risk" level; since even one premature death is "unsafe" for the person who dies, the maximum permissible number should, intuitively, be zero. One objection to this approach is that

(criticizing McCarthy's account).

369. For an in-depth, critical discussion of the possibility of harmless wrongdoing, focused on criminal law, see 4 JOEL FEINBERG, *THE MORAL LIMITS OF THE CRIMINAL LAW: HARMLESS WRONGDOING* (1988).

370. See Sadowitz & Graham, *supra* note 360 (surveying agency use of frequentist risk to set health and safety thresholds).

achieving a zero-death goal might be practically impossible. The objection becomes more powerful, the more broadly the regulated activity or product is defined. Regulating "food additives" so that no premature deaths result from food additives is eminently possible. If food additives are prohibited (and the prohibition is complied with), then they cause no deaths. Regulating "foods" so that no premature deaths result from food consumption is not possible, short of banning all foods. In any world in which people do consume food, there will be acts of food consumption such that the consumer dies earlier than she would have had she refrained. Another objection is that a zero-death goal is misconceived, even if health and safety are given priority over other aspects of welfare. Consider that, by banning a food additive which would cause deaths in some cases, the regulator might prevent the additive from blocking fatal processes that will result in many more deaths if the additive is banned.³⁷¹ Imagine that Additive A is a carcinogen which, if permitted for use in foods, would cause one death but also prevent twenty botulism deaths. Another way of phrasing this objection is that regulators instructed to prioritize health or safety should aim to maximize aggregate health or safety (concretely, to maximize aggregate longevity or health-adjusted longevity), rather than to eliminate deaths from particular causes (cancer versus botulism). A plausible response to the objection is that some statutes not only displace cost considerations but quite explicitly focus on the deaths that result from particular products or activities, rather than on global longevity, and that minimizing the number of deaths of this type—regardless of countervailing costs, including deaths of other types—is what legislative supremacy requires of regulators operating under these statutes.³⁷²

Further complications arise when a statute articulates a health/safety threshold but permits some regulatory consideration of cost. Consider the case where regulators are instructed to employ CBA in setting standards, but are precluded from regulating activities or products that are already "safe." In effect, "safe" activities or products are those that a full cost-benefit analysis likely would recommend leaving unregulated. Here, too, the health/safety threshold could be specified in terms of a maximum permissible number of deaths

371. See, e.g., LAVE, *supra* note 288, at 15-16.

372. See, e.g., Sunstein, *supra* note 271, at 1663-68 (describing the varying structures of safety-oriented statutes).

resulting from the activity or product. Note, however, that the maximum permissible number, in *this* statutory context, could well be greater than zero. Assume that an agency, to date, has regulated 100 different dangerous chemicals. The average cost of regulation has been, say, \$55 million. If V_{death} is \$6 million, and if it is assumed that the main benefit of regulation is the avoidance of premature death, then the agency or a court reviewing its decisions could reasonably determine that substances causing fewer than ten deaths are "safe."

In sum, agencies need not apply health/safety thresholds by looking to the levels of frequentist risk that occur in different regulatory outcomes. The thresholds might be construed in other ways, as tracking the number of premature deaths that occur in different outcomes, or the aggregate adjusted or unadjusted longevity. The current, risk-focused regulatory practice can and should change, given the welfare irrelevance of frequentist risk. What the revised practice ought to look like is a subtle and context-sensitive matter that deserves much greater attention than I have been able to give it here.

D. DISTRIBUTIVE JUSTICE

Arguably, risk regulators should not simply aim to maximize aggregate welfare (CBA), or to maximize aggregate health and safety (RRA), or to prevent health and safety setbacks of a particular kind, but should also take account of distributive considerations in shaping policy and setting regulatory and enforcement priorities. This might be disputed, either on the basis of a moral view, such as utilitarianism, that makes distributive justice irrelevant,³⁷³ or on the institutional grounds that distributive justice is best achieved through the tax-and-transfer system, not regulation.³⁷⁴ These objections to redistributive regulation cannot be evaluated here. Suffice it to say a wide range of moral theories, including but hardly limited to John Rawls's, incorporate distributive criteria;³⁷⁵ and that our tax-and-transfer system is seemingly far too insensitive to

373. See KAGAN, *supra* note 88, at 48-54 (discussing the possible moral role of distributive considerations).

374. See Louis Kaplow & Steven Shavell, *Why the Legal System Is Less Efficient than the Income Tax in Redistributing Income*, 23 J. LEGAL STUD. 667, 667-69 (1994).

375. See Adler & Posner, *supra* note 2, at 211 nn.126-30 (citing philosophical defenses of egalitarianism).

individual circumstance to warrant risk regulators in believing that inequities flowing from their policy choices will be corrected elsewhere.³⁷⁶

The idea of incorporating distributive justice into risk regulation is hardly novel. Activists in the so-called "environmental justice" movement, as well as numerous scholars writing about "environmental justice," have argued quite vigorously that the geographic distribution of air and water pollutants, toxic waste sites, and other health and safety hazards is unfair to racial and low-income groups, and that risk regulation should be structured to eliminate this inequity.³⁷⁷ The Clinton administration was responsive to these arguments. In 1994, President Clinton issued Executive Order 12,898, entitled "Federal Actions to Address Environmental Justice in Minority Populations and Low-Income Populations." The order states,

To the greatest extent practicable and permitted by law, . . . each Federal agency shall make achieving environmental justice part of its mission by identifying and addressing, as appropriate, disproportionately high and adverse human health or environmental effects of its programs, policies, and activities on minority populations and low-income populations in the United States³⁷⁸

This directive is reinforced by the broader mandate, set forth in Executive Order 12,866—the order generally requiring

376. See Chris William Sanchirico, *Deconstructing the New Efficiency Rationale*, 86 CORNELL L. REV. 1003, 1013 (2001) (noting that the existing income tax is quite imperfect with respect to redistribution).

377. The literature on "environmental justice" is vast. See, e.g., CHRISTOPHER H. FOREMAN JR., *THE PROMISE AND PERIL OF ENVIRONMENTAL JUSTICE* (1998); FENG LIU, *ENVIRONMENTAL JUSTICE ANALYSIS: THEORIES, METHODS AND PRACTICE* (2001); Vicki Been, *What's Fairness Got to Do with It: Environmental Justice and the Siting of Locally Undesirable Land Uses*, 78 CORNELL L. REV. 1001 (1993); Robert D. Bullard, *Leveling the Playing Field Through Environmental Justice*, 23 VT. L. REV. 453 (1999); Sheila R. Foster, *Meeting the Environmental Justice Challenge: Evolving Norms in Environmental Decisionmaking*, 30 ENVTL. L. REP. 10992 (2000); Robert R. Kuehn, *A Taxonomy of Environmental Justice*, 30 ENVTL. L. REP. 10681 (2000) [hereinafter Kuehn, *Taxonomy*]; Robert R. Kuehn, *The Environmental Justice Implications of Quantitative Risk Assessment*, 1996 U. ILL. L. REV. 103 [hereinafter Kuehn, *Environmental Justice*]; Richard J. Lazarus, *Pursuing "Environmental Justice": The Distributional Effects of Environmental Protection*, 87 NW. U. L. REV. 787 (1993); Bradford C. Mank, *Environmental Justice and Discriminatory Siting: Risk-Based Representation and Equitable Compensation*, 56 OHIO ST. L.J. 329 (1995); W. Kip Viscusi, *Risk Equity*, in COST-BENEFIT ANALYSIS 7, *supra* note 284, at 7; Rae Zimmerman, *Social Equity and Environmental Risk*, 13 RISK ANALYSIS 649 (1993).

378. Exec. Order No. 12,898 § 1-101, 3 C.F.R. 859, 859 (1995), *reprinted in* 42 U.S.C. § 4321 (2000).

federal executive agencies to engage in cost-benefit analysis—that “distributive impacts” and “equity” should be considered as part of this analysis. “[I]n choosing among alternative regulatory approaches, agencies should select those approaches that maximize net benefits (including potential economic, environmental, public health and safety, and other advantages; *distributive impacts; and equity*), unless a statute requires another regulatory approach.”³⁷⁹

Executive Orders 12,898 and 12,866 could not trump the statutory framework that governs risk regulators, nor do these orders purport to do so. Instead, they enjoin regulators to weigh distributive considerations only where statutorily permissible. Yet the EPA, FDA, OSHA, CPSC, and other risk-regulatory agencies *do* have substantial statutory discretion to take account of the racial or socioeconomic pattern of death, illness, injury, and exposure that their policy choices would produce.³⁸⁰ Statutes that enjoin agencies to balance health and safety against economic costs and other factors are typically open-ended enough to permit the incorporation of distributive justice into the balancing. For example, the Toxic Substances Control Act, the Consumer Product Safety Act, and the Federal Insecticide, Fungicide and Rodenticide Act all articulate an “unreasonable risk” standard.³⁸¹ An activity or substance that produces a higher rate of illness, injury or death for blacks, as compared with whites, can be characterized as “unreasonable” by virtue of this distributional inequity. Some statutes that give priority to health and safety can also be read to license distributive sensitivity. Consider the pre-1990 version of section 112 of the Clean Air Act, which required EPA to set legal emissions of hazardous air pollutants at a level that “provide[s] an ample margin of safety to protect the public health.”³⁸² Assume that at the level proposed by industry for a given pollutant, virtually no premature deaths or serious illnesses will result but certain less serious impacts, such as mild illness or physical discomfort, will not be eliminated and

379. Exec. Order No. 12,866 § 1(a), 3 C.F.R. 638, 638-39 (1994) (emphasis added), reprinted in 5 U.S.C. § 601 (2000).

380. See Foster, *supra* note 377, at 10994-97; Richard J. Lazarus & Stephanie Tai, *Integrating Environmental Justice into EPA Permitting Authority*, 26 *ECOLOGY L.Q.* 617, 625-50 (1999); Lazarus, *supra* note 377, at 825-55.

381. See *supra* text accompanying notes 276-78.

382. See *supra* text accompanying note 273.

will be concentrated in areas populated by low-income groups and racial minorities. Arguably, the proposed level would not "provide an ample margin of safety" for Clean Air Act purposes because it would produce a racially and socioeconomically skewed pattern of health setbacks.

Even statutes that do not permit agencies to consider distributive justice in shaping legal directives—that is, in issuing legally binding rules or legally binding individualized orders, such as permits to particular facilities or licenses for particular products—may well allow agencies to inject distributive considerations into other choices. Regulatory priorities, or enforcement priorities, could be influenced by racial or socioeconomic skews. OSHA's rules governing toxic workplace substances are supposed to ensure that no worker suffers "material impairment of health or functional capacity."³⁸³ This may leave no space for distributive considerations at the rule-issuance phase. But could OSHA, in deciding which workplace toxins to regulate first, give priority to those substances the toxic effects of which are borne disproportionately by blacks? Or could the agency, in determining which firms to target for violations of existing rules, give priority to those firms whose illegalities cause disproportionate harm to low-income workers? Arguably so. Finally, it should be noted that distributive considerations may bear upon *state* risk regulation—both as a matter of state law, and as a matter of federal law, in particular Title VI of the Civil Rights Act of 1964, which precludes racial discrimination in "any program or activity receiving Federal financial assistance."³⁸⁴ Title VI has been widely understood by federal agencies to authorize regulations prohibiting recipients of federal funds, including state actors, from activities that have a disparate racial impact.³⁸⁵ The upshot is that state risk regulators may run afoul of federal regulations enacted under Title VI if their choices produce a racially skewed pattern of death, illness, injury, toxic exposure, or other setbacks.³⁸⁶

383. See *supra* text accompanying note 272.

384. 42 U.S.C. § 2000d (2000).

385. See *id.* § 2000d-1 (instructing federal agencies, in section 602 of Title VI, to issue regulations to effectuate the statutory nondiscrimination norm); John Arthur Laufer, Note, *Alexander v. Sandoval and Its Implication for Disparate Impact Regimes*, 102 COLUM. L. REV. 1613 (2002) (describing agency practice under section 602 and the relevant case law).

386. See Lazarus, *supra* note 377, at 834-39; Bradford C. Mank, *The Draft Title VI Recipient and Revised Investigation Guidances: Too Much Discretion*

A consensus methodology for implementing Executive Order 12,898, and more generally for measuring the degree of distributional skews with respect to health and safety, has yet to emerge.³⁸⁷ It is frequently suggested, however, that the racial or socioeconomic distribution of *frequentist risk* is an important part of the environmental justice inquiry. For example, Robert Kuehn, a leading scholarly proponent of "environmental justice," argues that agency risk assessments should not simply aim to determine the average, median, or maximal (frequentist) risk created by some health or safety threat, but should determine (frequentist) risk for racial and economic subgroups of the exposed population. "[T]he inclusion of distributional information [should be] a matter of course in all risk assessments. Risk assessors should be required to include information on the exposures and risks experienced by relevant subpopulations disaggregated by race, ethnicity, income, age, and other important variables."³⁸⁸ John Graham—a very prominent scholar in the area of risk assessment, who now chairs the regulatory review office in OMB—agrees:

[W]hen the same hazard poses more danger to some citizens than others, Congress should insist that agencies report that information through *distributional methods of variability analysis*. For example, some citizens are more sensitive to environmental agents than others Some citizens are also exposed more to hazards than others Agencies should present to decision makers a public document with information about the number of citizens exposed to various levels of risk. Since low-income and minority citizens often incur a disproportionate share of public health and environmental risks, agencies should make a special effort to investigate those citizens' degree of exposure and susceptibility to hazards.³⁸⁹

The approach sketched by Kuehn and Graham was endorsed in 1996 by the National Research Council,³⁹⁰ and in 1997 by a Presidential/Congressional Commission on Risk Assessment

for EPA and a More Difficult Standard for Complainants?, 30 ENVTL. L. REP. 11144, 11147-48 (2000).

387. See U.S. OFFICE OF MGMT. & BUDGET, *supra* note 310, § III.A.8 ("There are no generally accepted principles for determining when one distribution of net benefits is more equitable than another.").

388. Kuehn, *Environmental Justice*, *supra* note 377, at 151.

389. John D. Graham, *Making Sense of Risk: An Agenda for Congress*, in RISKS, COSTS, AND LIVES SAVED 183, *supra* note 31, at 190-91.

390. See COMM. ON RISK CHARACTERIZATION, NAT'L RESEARCH COUNCIL, UNDERSTANDING RISK: INFORMING DECISIONS IN A DEMOCRATIC SOCIETY 40-41, 157 (Paul C. Stern & Harvey V. Fineberg eds., 1996).

and Risk Management.³⁹¹ It has also been used in empirical scholarship seeking to determine whether risk regulation and risky activities indeed produce distributional skews along racial or economic lines.³⁹²

More recently, a risk-based methodology for identifying such skews was sanctioned by the EPA in its draft Title VI guidelines.³⁹³ The guidelines provide that a Title VI challenge to a state or local permitting decision should be evaluated by determining whether there is some group that is adversely affected by the decision and that has a higher proportion of racial minorities and suffers a greater adverse impact from the decision than an appropriate comparison group. In turn, the risk imposed by the decision on the two groups is one legitimate measure of the decision's impact.

[The Title VI analysis] involves a disparity analysis that compares the affected population to an appropriate comparison population to determine whether [a] disparity exists that may violate EPA's Title VI regulations

A disparity may be assessed using comparisons both of the different prevalence of race, color, or national origin of the two populations, *and of the level of risk of adverse impacts experienced by each population.*³⁹⁴

"Risk," for purposes of this disparate impact analysis, clearly means frequentist, not Bayesian risk.

[Disparate impact analysis under Title VI] involves prediction of potentially significant exposures and risks resulting from stressors created by the [challenged] activities or other sources. These predictions may be based on ambient levels of stressors derived from monitoring or modeling In estimating cancer risks, such *unit risk factors* estimate the probability of contracting a cancer case for a unit of exposure. For example, an area's predicted cancer risk could be based on the estimated ambient concentration times the unit risk

391. See 2 THE PRESIDENTIAL/CONG. COMM'N ON RISK ASSESSMENT & RISK MGMT., RISK ASSESSMENT AND RISK MANAGEMENT IN REGULATORY DECISION MAKING 72-77 (1997).

392. See JAMES T. HAMILTON & W. KIP VISCUSI, CALCULATING RISKS?: THE SPATIAL AND POLITICAL DIMENSIONS OF HAZARDOUS WASTE POLICY 157-88 (1999).

393. Draft Title VI Guidance for EPA Assistance Recipients Administering Environmental Permitting Programs (Draft Recipient Guidance) and Draft Revised Guidance for Investigating Title VI Administrative Complaints Challenging Permits (Draft Revised Investigation Guidance), 65 Fed. Reg. 39,650 (June 27, 2000) [hereinafter EPA Draft Title VI Guidance]; see Mank, *supra* note 386 (providing detailed analysis of EPA Draft Title VI Guidance).

394. EPA Draft Title VI Guidance, *supra* note 393, at 39,681 (emphasis added).

factor.³⁹⁵

To sum up, distributive justice is often a relevant consideration for state and federal risk regulators, under applicable statutes and executive orders, and frequentist risk is widely proposed and used as a metric for distributive injustice.³⁹⁶ But this approach is undermined by the Hybrid (Bayesian) View of risk—specifically by the proposition, defended above in Part III, that risk in the frequentist sense is not a welfare setback. Assume that E is an environmental hazard that imposes a higher frequentist risk on a group of persons G_1 that live close to E, as compared to a group of persons G_2 that live farther from E, where G_1 is disproportionately nonwhite as compared to G_2 or the total population ($G_1 + G_2$). The Hybrid (Bayesian) View of risk implies that the members of G_1 are not worse off in welfare terms than their counterparts in G_2 , qua the difference in frequentist risk, nor are they worse off with respect to welfare resources. Conceivably, the high level of frequentist risk among G_1 might constitute a violation of its members' *rights*—a violation not suffered by the members of G_2 —but distributive justice is standardly seen to depend on the distribution of welfare or welfare resources, not on the pattern of rights violations. Further, a high level of frequentist risk could constitute a rights violation only when conjoined with further

395. *Id.* at 39,679; see also *id.* at 39,680 (endorsing the use of risk levels to determine whether a permitting decision has an adverse impact). Other official EPA documents also endorse use of a risk-based methodology for identifying distributive skews. See, e.g., 2 U.S. ENVTL. PROT. AGENCY, ENVIRONMENTAL EQUITY: REDUCING RISK FOR ALL COMMUNITIES 31-36 (1992); U.S. ENVTL. PROT. AGENCY, EPA GUIDANCE FOR CONSIDERATION OF ENVIRONMENTAL JUSTICE IN CLEAN AIR ACT SECTION 309 REVIEWS 16 (1999).

396. At the risk of undue repetition, I should make clear that the kind of "metric" I am interested in is a metric for the distributive injustice of outcomes, which in turn will help determine what policies regulators ought to choose in light of distributive justice. The legal proposals and academic commentary just described, as I read them, do conceptualize frequentist risk as a feature of outcomes, independent of premature death, that bears upon the justice of the outcomes. This is different from saying that (1) in a given outcome, it is the differential death rate between two groups that make the outcome unequal, and (2) in evaluating choices, regulators should (*ceteris paribus*) avoid a choice with a higher risk, in some sense, of producing an outcome skewed with respect to its death rate. See *supra* text accompanying notes 95-100 (distinguishing between ex post and ex ante views). Although some advocates of incorporating risk in distributive justice analysis may merely be proposing the use of risk in this sort of "bridging" or ex ante role, the sources I have just cited seem to propose that it have a quite different role in the analysis, a role in ranking outcomes.

elements, such as a suitable lack of voluntariness and responsibility on the part of the person suffering the risk. A metric for distributive skews focused on rights violations would need to be sensitive to these elements; but the risk-based metrics now proposed and used are not (typically) thus sensitive.

What *would* be an appropriate welfare-based metric for determining whether the situation of G_1 and G_2 constitutes a distributive injustice? Imagine that a substantially higher proportion of G_1 dies as a result of E than does G_2 . Since death is a welfare setback, the difference in death rates between G_1 and G_2 *could* be seen as an unfair distribution of welfare or welfare resources, and thus as an instance of distributive injustice along racial lines. To be sure, differences in death rates, as between population subgroups, will often track differences in frequentist risks. But the two kinds of comparisons are analytically distinct. For example, if the frequentist-risk comparison between G_1 and G_2 focuses on the maximally exposed individual, then clearly the two groups could have dramatically different levels of frequentist risk but not dramatically different death rates. The same is true (a bit less obviously) if the frequentist-risk comparison focuses on the risk to the median members of G_1 and G_2 .

To be sure, death is not the only dimension of welfare relevant to "environmental justice." A full distributive justice analysis of regulatory outcomes would arguably consider illness and injury; wealth and income; unemployment; loss of natural resources; and many other welfare benefits or hindrances such as "[n]oise, odors, blowing trash, aesthetic concerns, increased traffic, termites, decreased property values and uses, fires, accidents, [and] psychological harm."³⁹⁷ The Hybrid (Bayesian) View of risk implies that the distribution of Bayesian risk—by contrast with the distribution of frequentist risk—is, or might be, a relevant outcome-ranking factor for regulators authorized to consider "distributive impacts" and promote "environmental justice." If a large percentage of nonwhites believe to a high degree that they will die from environmental toxins, while a smaller percentage of whites believe that, then a disproportionate number of nonwhites suffer a kind of welfare setback—on the view that Bayesian risks need not be affectively charged to be harmful. At a minimum, probabilistic

397. Kuehn, *Taxonomy*, *supra* note 377, at 10,687.

beliefs about premature death that *are* affectively charged are harmful; and in practice, of course, such beliefs will be affectively charged. In short, the skewed distribution of *fear* is a kind of distributive injustice, one that (in principle) should be a relevant consideration under Executive Orders 12,898 and 12,866 and similar regimes. As the philosopher Keith Burgess-Jackson has argued in a different context,

[W]hile fear, or a certain amount . . . of fear, may be useful in keeping individuals safe from harm, it is, as a mental state, burdensome to those who experience it. Fear of crime in particular is a socially created burden [Further it is] women [who] bear a disproportionate share of this burden Justice, I maintain, requires that the state, employing institutions such as law and education and using the power of the purse, both reduce the overall level of fear in society and, more particularly, redistribute fear so that women no longer bear the brunt of it.³⁹⁸

E. COMPARATIVE RISK ASSESSMENT

"Risk assessment," the influential methodology described above, is widely used by agencies at the rule-making stage. Paradigmatically, a risk assessment identifies the deaths and frequentist risks that would result were a particular toxic substance left unregulated, and that information is then used to determine the advisability of a proposed rule.³⁹⁹ So-called "comparative risk assessment" (CRA) is quite different.⁴⁰⁰ CRA is a *priority-setting* tool. It informs threshold legislative, presidential, or administrative choices about which hazardous activities, substances, or conditions deserve governmental attention. Threshold choices about the appropriate priority that risk regulators ought to give to various hazards are made

398. Keith Burgess-Jackson, *Justice and the Distribution of Fear*, 32 S. J. PHIL. 367, 368 (1994).

399. See *supra* text accompanying notes 308-17317.

400. The literature on CRA is large. Important contributions include: BREYER, *supra* note 4; COMPARING ENVIRONMENTAL RISKS: TOOLS FOR SETTING GOVERNMENT PRIORITIES (J. Clarence Davies ed., 1996); WORST THINGS FIRST? THE DEBATE OVER RISK-BASED NATIONAL ENVIRONMENTAL PRIORITIES (Adam M. Finkel & Dominic Golding eds., 1994); John S. Applegate, *A Beginning and Not an End in Itself: The Role of Risk Assessment in Environmental Decision Making*, 63 U. CIN. L. REV. 1643, 1656-67 (1995); Finkel, *supra* note 344; Baruch Fischhoff, *Ranking Risks*, 6 RISK 191 (1995); Donald T. Hornstein, *Reclaiming Environmental Law: A Normative Critique of Comparative Risk Analysis*, 92 COLUM. L. REV. 562 (1992); Symposium, *Risk Assessment*, *supra* note 308; and Symposium, *Risk in the Republic: Comparative Risk Analysis and Public Policy*, 8 DUKE ENVTL. L. & POL'Y F. 1 (1997).

by OMB and Congress in the budget-setting process and elsewhere.⁴⁰¹ As for priority setting at the administrative level, that occurs, for example, when an agency decides which rule making to initiate first, or where to target enforcement effort.

The EPA pioneered CRA in the late 1980s, when it studied the relative importance of thirty-one different environmental problems and published the results of the study in a now-famous report, *Unfinished Business: A Comparative Assessment of Environmental Problems*.⁴⁰² The problems were ranked with respect to four different concerns: the occurrence of cancer and cancer risks, the occurrence of non-cancer health setbacks and risks, ecological effects, and "welfare effects" (meaning quantifiable non-health harms such as property damage).⁴⁰³ The ranking of the thirty-one problems with respect to cancer, as follows, illustrate the typical scope and conclusions of a CRA.

Consensus Ranking of Environmental Problem Areas on the Basis of Population Cancer Risk⁴⁰⁴

<u>Problem</u>	<u>Rank (from most to least important)</u>
Worker Exposure to Chemicals	1 (tied)
Indoor Radon	1 (tied)
Pesticide Residues on Food	3
Indoor Air Pollutants other than Radon	4 (tied)
Consumer Exposure to Chemicals	4 (tied)
Hazardous Air Pollutants	6
Depletion of Stratospheric Ozone	7
Hazardous Waste Sites— Inactive	8
Drinking Water	9
Application of Pesticides	10

401. See U.S. OFFICE OF MGMT. & BUDGET, BUDGET OF THE UNITED STATES GOVERNMENT, FISCAL YEAR 1992 pt. 2 § IX.C, at 371-76 (proposing use of risk assessment in determining agency budgets).

402. U.S. ENVTL. PROT. AGENCY, UNFINISHED BUSINESS: A COMPARATIVE ASSESSMENT OF ENVIRONMENTAL PROBLEMS (1987) [hereinafter UNFINISHED BUSINESS].

403. See *id.* at 5-16 (describing study methodology).

404. See *id.* at 28-34 (listing all thirty-one problems).

Since *Unfinished Business*, CRA has become an important feature of risk regulation. The EPA has continued to conduct CRA studies and to rely upon these studies in its internal priority setting.⁴⁰⁵ Other federal agencies, too, have engaged in CRA, although not as visibly as the EPA.⁴⁰⁶ Executive Order 12,866 implicitly directs federal executive agencies to conduct CRA, where permissible by statute, by mandating that, "[i]n setting regulatory priorities, each agency shall consider, to the extent reasonable, the degree and nature of the risks posed by various substances or activities within its jurisdiction."⁴⁰⁷ In the early 1990s, then-Judge Stephen Breyer wrote a widely cited book which argued that the existing regime of risk regulation failed to direct governmental effort and social resources to the most pressing problems, and advocated the creation of an OMB-like oversight agency staffed by high-quality civil servants to coordinate risk-regulatory efforts—in particular, to perform CRA studies and to implement these studies in setting intra- and interagency priorities.⁴⁰⁸ Prodded by this book, and by general enthusiasm for CRA among advocates of "regulatory reform," various statutory proposals to impose a broad CRA requirement on agencies have been considered by Congress, although none has yet been enacted.⁴⁰⁹

CRA has also become popular at the state and local level. California performed a major CRA study in 1994,⁴¹⁰ and many

405. See Ken Sexton, *Setting Environmental Priorities: Is Comparative Risk Assessment the Answer?*, in BETTER ENVIRONMENTAL DECISIONS: STRATEGIES FOR GOVERNMENTS, BUSINESSES, AND COMMUNITIES 195, 199 (Ken Sexton et al. eds., 1999) (describing use of CRA at the EPA). For an important and often-cited follow-up study to *Unfinished Business*, see SCI. ADVISORY BD., U.S. ENVTL. PROT. AGENCY, REDUCING RISK: SETTING PRIORITIES AND STRATEGIES FOR ENVIRONMENTAL PROTECTION (1990).

406. See, e.g., Thomas P. Grumbly, *Comparative Risk Analysis in the Department of Energy*, 8 DUKE ENVTL. L. & POL'Y F. 23 (1997); see also CARNEGIE COMM'N ON SCI., TECH. & GOV'T, RISK AND THE ENVIRONMENT: IMPROVING REGULATORY DECISION MAKING 75-82 (1993) (recommending that all federal risk-regulatory agencies engage in CRA).

407. Exec. Order No. 12,866 § 1(b)(4), 3 C.F.R. 638, 639 (1994), reprinted in 5 U.S.C. § 601 (2000).

408. See BREYER, *supra* note 4, at 59-81.

409. See Fred Anderson et al., *Regulatory Improvement Legislation: Risk Assessment, Cost-Benefit Analysis, and Judicial Review*, 11 DUKE ENVTL. L. & POL'Y F. 89, 96-108 (2000) (describing recent regulatory reform proposals, including proposed CRA provisions).

410. CAL. COMPARATIVE RISK PROJECT, TOWARD THE 21ST CENTURY: PLANNING FOR THE PROTECTION OF CALIFORNIA'S ENVIRONMENT (1994).

other states and localities have also done so.⁴¹¹ Finally, CRA has become a central focus of the academic literature on risk regulation. The pros and cons of CRA and specific CRA methodologies has emerged as a major topic of academic debate.⁴¹²

The Hybrid (Bayesian) View advanced in this Article has important implications for CRA. *No welfare-focused CRA should rely upon frequentist risk as a kind of consequence relevant to regulatory priorities.* Frequentist risk is properly irrelevant to the ranking of outcomes within the context of CBA, RRA, the setting of health and safety thresholds, and distributive justice analysis, or so I have argued.⁴¹³ Similarly, frequentist risk is irrelevant (at the outcome-ranking stage) for purposes of CRA—more precisely, for purposes of welfare-focused CRA. In principle, CRA might evaluate hazardous activities, substances, and conditions in light of their non-welfare effects—for example, in order of the rights infringements they involve. Since I am (for purposes of this Article) agnostic on the question whether the imposition of a high level of frequentist risk amounts to a harmless wrong, I am agnostic on the proper inclusion of frequentist risk in nonwelfarist CRA. But CRA is typically conceptualized in welfarist terms: as a methodology for evaluating hazardous activities, etc., with a view to their impact on human welfare. Frequentist risk (whether risk to the average individual, risk to the maximally exposed individual, or frequentist risk in some other sense) should play no role in ranking outcomes within CRA thus conceptualized.

What *is* the proper form of a welfare-focused CRA? A generic answer is this: Hazards should be ranked with respect to one or more aspects of human welfare. Most simply, hazards could be ranked in light of the total number of deaths they produce, under the regulatory status quo. Indeed, this was the approach followed by the EPA in the seminal *Unfinished*

411. See David M. Konisky, *Over a Decade of Comparative Risk Analysis: A Review of the Human Health Rankings*, 12 RISK 41, 42 (2001) (noting that “over thirty states and municipalities have completed comparative risk projects”); Sexton, *supra* note 405, at 199 (describing state CRA activity). CRA has also become important internationally. See Richard D. Morgenstern et al., *Comparative Risk Assessment: An International Comparison of Methodologies and Results*, 78 J. HAZARDOUS MATERIALS 19, 24-34 (2000) (surveying the international use of CRA).

412. See *supra* note 400.

413. See *supra* Part V.A-D.

Business study.⁴¹⁴ This simple approach to ranking hazards ignores the costs of ameliorating the hazard, but developing that information is itself costly (particularly since the cost will depend on the particular regulatory approach used to ameliorate the hazard). Once multiple dimensions of welfare are included, a problem of commensuration arises.⁴¹⁵ If the analyst has available a dollar valuation of lifesaving that approximately tracks its welfare effect (what I have called V_{death}), then hazards might be ranked in terms of the predicted net dollar benefits of ameliorating or eliminating them. Alternatively, they might be ranked, without pricing life, in terms of their marginal or average cost per life saved—the cost per life saved of a small or large regulatory expenditure to address the hazard.⁴¹⁶

It is beyond the scope of this Article to discuss whether welfarist CRA should be single- or multidimensional; which particular dimension(s) of welfare should be included; and how multiple dimensions should be commensurated. The Hybrid (Bayesian) View does have direct implications, however, concerning the *permissible dimensions* of welfarist CRA. Since death but not frequentist risk is a welfare setback, the number of deaths but not the level of frequentist risk is a permissible dimension for welfarist CRA—a kind of consequence permissibly given weight in welfarist priority setting.

It appears that most CRA studies, at least by United States governmental entities, have emphasized the total deaths resulting from various hazards and have given less weight to the level of frequentist risk.⁴¹⁷ On the other hand, the EPA's *Guidebook to Comparing Risks and Setting Environmental Priorities* states that both "population risk" (the total number of deaths) and average or maximal "individual risk" (what I term

414. See UNFINISHED BUSINESS, *supra* note 402, at 26, 41.

415. On multidimensional CRA, see John Kadvany, *From Comparative Risk to Decision Analysis: Ranking Solutions to Multiple-Value Environmental Problems*, 6 RISK 333, 342-49 (1995).

416. See BREYER, *supra* note 4, at 24-27 (ranking regulations by "[c]ost per premature death averted"); cf. LAVE, *supra* note 288, at 19-25 (describing both cost-benefit and cost-effectiveness techniques for regulatory analysis).

417. See U.S. ENVTL. PROT. AGENCY, A GUIDEBOOK TO COMPARING RISKS AND SETTING ENVIRONMENTAL PRIORITIES 2.2-15 (1993) ("While many regional and state comparative risk studies have estimated both individual risks and population risks, most have relied primarily on the population risk estimates in the final ranking of problem areas for cancer risks."). I have not surveyed the studies since 1993 to determine whether this statement still holds true.

frequentist risk) are equally valid criteria for comparing cancer hazards.⁴¹⁸ According to the *Guidebook*, “[t]he most useful presentations of cancer risk estimates for comparative risk studies are excess individual lifetime risks and the excess numbers of annual cancer cases expected in the exposed populations.”⁴¹⁹ Adam Finkel, a prominent expert on CRA, similarly claims that “population-based measures” and “[i]ndividual-risk measures independent of the number of persons at risk” are both appropriate dimensions for CRA.⁴²⁰

Estimating the magnitude of risk either by population-wide consequence alone or by maximum (or average) individual risk alone necessarily discriminates against people who live in either densely-populated or sparsely-populated areas. The former type of measure regards a risk of 10^{-6} /person in a city of two million as more serious than one of 10^{-2} /person in a village of 100. . . . On the other hand, if the individual-risk criterion were used exclusively, urbanites would lose out⁴²¹

If, however, neither the urbanites nor the villagers in this hypothetical are aware of the hazards, and hazard-prevention costs are the same, then the city hazard *should* take regulatory priority over the village hazard, as a matter of overall welfare or welfare resources—since the city hazard will lead to more deaths than the village hazard.⁴²² Notwithstanding claims to the contrary by Finkel, the EPA *Guidebook*, and other CRA analysts who have advanced similar suggestions, the level of “individual risk” created by some hazard should not be used, for purposes of regulatory priority setting, as a metric for the importance of the hazard distinct from the aggregate number of deaths it causes.

F. TORT AND CRIMINAL LAW

My focus in this Article has been risk regulation, but the view developed here also bears upon tort law and criminal law,

418. *Id.* at 2.2-15 to 2.2-18.

419. *Id.* at 2.2-15.

420. Finkel, *supra* note 344, at 343.

421. *Id.*

422. More precisely, in the simple case where the regulator knows for certain that one death will result in the village and two in the city, generating a 10^{-2} frequentist risk for everyone in the village and 10^{-6} risk for everyone in the city, the city hazard should take priority. If (more realistically) the legislator uses his beliefs about the level of frequentist risk to generate a probability distribution, of some kind, over the number of deaths, then choosing which hazard to address raises issues of ex ante valuation—issues beyond the scope of this Article. See *supra* Part I.C.

specifically with respect to the status of risk imposition as a crime or tort.

Consider first tort law. In a provocative article, Christopher Schroeder suggests that the actual occurrence of death, bodily injury, or other tangible harm need not be a standard prerequisite for a tort judgment.⁴²³ He delineates two hypothetical cases.

The Toxic Dumper Case

Amalgamated Manufacturing is discovered to have been dumping toxic chemicals into an underground aquifer. Company records reveal fairly well the amounts that have been discharged. Studies suggest that individuals who have consumed well water from the aquifer have between a twenty and thirty-three percent chance of contracting cancer within their lifetimes because of this consumption.

The Speeding Motorist Case

A motorist, in a hurry to get from her house to an important engagement in the neighboring town, enters the connecting freeway and speeds up to ten miles per hour over the posted speed limit. She arrives at her destination safely, only a few minutes late.⁴²⁴

Schroeder argues that persons who impose risks of harms on others should be subject to tort liability—not merely in the special case of “toxic torts,” exemplified by Toxic Dumper, but also in the garden-variety case of risk imposition, exemplified by Speeding Motorist—if the administrative costs of doing so are low.⁴²⁵ Schroeder’s claim has been seconded, both by scholars (like him) who understand tort law as an institution properly designed to serve corrective justice,⁴²⁶ and by law-and-economics scholars who see economic efficiency, social welfare, or similar welfarist goals as the purpose of that institution.⁴²⁷

The view that tangible harm is not a precondition for tort liability is, to be sure, merely a proposal. Current doctrine typically does require the occurrence of death, bodily injury, property damage, or at least emotional harm, for a tort suit to go forward.⁴²⁸ Risk imposition per se is, typically, not enough.

423. Christopher H. Schroeder, *Corrective Justice and Liability for Increasing Risks*, 37 UCLA L. REV. 439 (1990); see also Kenneth W. Simons, *Corrective Justice and Liability for Risk-Creation: A Comment*, 38 UCLA L. REV. 113 (1990) (responding to Schroeder).

424. Schroeder, *supra* note 423, at 439-40.

425. See *id.* at 460-77.

426. See David McCarthy, *Liability and Risk*, 25 PHIL. & PUB. AFF. 238 (1996).

427. See WILLIAM M. LANDES & RICHARD A. POSNER, *THE ECONOMIC STRUCTURE OF TORT LAW* 263-69 (1987).

428. See Note, *Latent Harms and Risk-Based Damages*, 111 HARV. L. REV.

But there are a few important doctrinal exceptions from this demand for tangible harm. First, in so-called "loss-of-chance" cases, some courts have imposed liability for the mere creation of a risk.⁴²⁹ The standard "loss-of-chance" scenario involves a person threatened by tangible harm and a defendant in a special relationship with that person, such that the defendant more likely than not was not in a position to avert the harm, but acted or failed to act in a way that increased the risk (in some sense) of the person suffering the harm. For example, imagine that Pete now has advanced skin cancer, which doctor Dan negligently failed to diagnose. Had Dan diagnosed and properly treated the cancer, Pete's risk (in some sense) of dying from the cancer would have been sixty percent. Dan's failure to diagnose the disease increased Pete's chance of dying to ninety percent. Under current doctrine, in this hypothetical case, some courts would award Pete damages for the thirty percent "loss of chance" of avoiding death caused by Dan's negligence. Other courts would not, but *would* have Dan pay damages to Pete's survivors if he were to die—even though, on those facts, Dan would not be the likely cause of Pete's death.⁴³⁰

Second, even without the special relationship between the defendant and the risked or injured party characteristic of "loss-of-chance" cases, courts may impose risk liability on a defendant whose action has both caused present injury to some person and created a risk of future injury to that same person. For example, in a case where the defendant negligently caused head injury to the plaintiff, damages were increased to compensate for the small risk of epilepsy that might result from the head injury.⁴³¹

Finally, in the toxic tort context illustrated by Schroeder's Toxic Dumper hypothetical, some courts do award damages to plaintiffs who have been exposed to a toxic substance and

1505, 1505 (1998).

429. See DAN B. DOBBS, *THE LAW OF TORTS* 436-49 (2000); see also Joseph H. King, Jr., "Reduction of Likelihood" Reformulation and other Retrofitting of the Loss-of-a-Chance Doctrine, 28 U. MEM. L. REV. 491, 499-516 (1998) [hereinafter, *Reduction of Likelihood*]. For an early article that spurred development of this doctrine, see Joseph H. King, Jr., *Causation, Valuation, and Chance in Personal Injury Torts Involving Preexisting Conditions and Future Consequences*, 90 YALE L.J. 1353 (1981).

430. See *Reduction of Likelihood*, *supra* note 429, at 502-11 (distinguishing between past-harm and future-harm variants of the "loss-of-a-chance" scenario).

431. Note, *supra* note 428, at 1510.

whose risk of death or disease has been increased by that exposure, but have not died and are currently asymptomatic for the diseases possibly resulting from the exposure.⁴³² This occurs, typically, in class action suits, where the class is structured to include some plaintiffs who have already been injured by the exposure, and others merely at risk of injury.⁴³³

What kind of "risk" is involved in the "loss-of-chance" cases, toxic torts cases, and others where liability has been imposed or increased for risk imposition? The answer varies and is not always clear. It appears that risk imposition involves either (1) a frequentist risk of premature death, relative to some kind of reference class; (2) a hypothetical Bayesian risk, namely the probabilistic belief of some hypothetical bystander or "reasonable person" that the defendant's action will cause premature death or other injury; or (3) an actual Bayesian risk, specifically the actual probabilistic belief of the actor himself that his action will result in premature death or other injury. The frequentist construal of risk imposition is certainly employed in the toxic tort context.⁴³⁴ If, for example, a firm's toxic release has exposed 1000 persons to a carcinogenic substance, one of whom has already died, and experts predict that eight more will too, the court adjudicating a class action against the firm might well determine that the firm has imposed a 9-in-1000 or 8-in-999 risk of death on each class member. This is, of course, a frequentist risk relative to the reference group of all exposed persons (namely, the plaintiff class). On the other hand, the increased risks that trigger liability in the "loss-of-chance context" may (at least sometimes) be Bayesian, not frequentist.⁴³⁵

What does my Article imply for the practice of imposing tort liability for risk imposition: both the (fairly limited)

432. See *id.* at 1511; see also Tamsen Douglass Love, *Detering Irresponsible Use and Disposal of Toxic Substances: The Case for Legislative Recognition of Increased Risk Causes of Action*, 49 VAND. L. REV. 789, 809-10 (1996); David Rosenberg, *Individual Justice and Collectivizing Risk-Based Claims in Mass-Exposure Cases*, 71 N.Y.U. L. REV. 210, 216-24 (1996).

433. Note, *supra* note 428, at 1511.

434. Cf. Rosenberg, *supra* note 432, at 235. See generally Schroeder, *supra* note 423, at 477 (suggesting that risk imposition should be tortious only in cases where the risks are "actuarial").

435. Cf. *Reduction of Likelihood*, *supra* note 429, at 532-33 (arguing that both "individualized" and "statistical" evidence should be used to establish "loss of chance").

practice under current doctrine, and the much broader practice advocated by Schroeder and other scholars? The Article undercuts one important justification for that practice.⁴³⁶ One way to defend the practice, of course, is to argue that the risk of death, bodily injury, or other tangible harm is *itself* a harm to the person upon whom the risk has been imposed. The Hybrid (Bayesian) View of risk I have presented here vitiates that defense—at least if risk is construed in one of the three senses I have just described. Imagine that toxic dumper D_1 imposes a high frequentist risk of premature death on P_1 . Imagine, now, that doctor D_2 fails to diagnose P_2 's disease, such that the hypothetical observer's probabilistic belief that P_2 will die prematurely has increased substantially. Imagine, finally, that doctor D_3 fails to diagnose P_3 's disease, and the good doctor herself possesses a probabilistic belief that this failure will cause P_3 's death. In none of these cases has the plaintiff yet been harmed (at least consequentially) by the defendant's actions, absent some fear on the plaintiff's part or some other impact additional to the enhanced risk of death that he now suffers. That conclusion follows directly from the Hybrid (Bayesian) View of risk defended here. Note that P_2 's risk of death is what I earlier termed a "hypothetical" Bayesian risk, and that P_3 's risk is an actual, third-party Bayesian risk rather than an actual, first-party Bayesian risk, since it is D_3 's beliefs, not P_3 's, that constitute the "risk" here. Although my Hybrid (Bayesian) View recognizes that actual, first-party beliefs may be harmful, it denies that hypothetical or third-party beliefs are.⁴³⁷

To be sure, risk imposition construed in a different way—in actual, first-party terms—is harmful, at a minimum if the actual beliefs on P 's part produced by defendant D 's action are a component of fear, anxiety, or some other epistemic-affective hybrid. "Emotional harm" is a genuine harm.⁴³⁸ What courts

436. Stephen Perry, in his rigorous and important work on the tortiousness of risk imposition, concludes that risk imposition in the frequentist sense is not harmful, at least where the harm-causing processes are deterministic. Perry, *supra* note 25, at 336. As stated earlier I endorse Perry's conclusion and some (if not all) of his arguments. See *supra* text accompanying notes 160-65.

437. I do acknowledge that the beliefs or at least fears of a person's friends and family members may be harmful to her. See *supra* note 262 and accompanying text.

438. See Perry, *supra* note 25, at 338-39 (distinguishing risk imposition from the causation of psychological harm).

in the toxic tort and "loss-of-chance" contexts have done, however, is to recognize a tort cause of action for risk imposition *separate* from any cause of action for emotional distress.⁴³⁹ My Article suggests that *this* practice is justifiable, if at all, only on a revisionary understanding of tort law: only if tort liability is justifiably imposed without harm at all. Many law and economists would accept that revisionary view of tort law.⁴⁴⁰ So would some corrective justice theorists—for example, theorists who think that an action which is wrongful but harmless (say, an action that ought not have been performed in light of its probable harms and benefits, but may not cause those harms and in any event has not yet) may, consistent with corrective justice, be adjudged tortious.⁴⁴¹ On the other hand, corrective justice theorists, and others, who believe that the traditional doctrine that the tort plaintiff must show actual harm reflects a deep and important feature of tort law, should conclude that risk imposition ought not be tortious.⁴⁴² Cases like *Toxic Dumper* and *Speeding Motorist* cannot be assimilated to the standard paradigm that identifies duty, breach of duty, causation, and harm as the individually necessary and jointly sufficient elements of a tort suit.

This last claim is a bit too strong. It might be argued that "harm" for purposes of the standard tort paradigm is a welfare setback understood in deontological rather than consequentialist terms. On the consequentialist construal, P is harmed by D's action (relative to an alternative choice) just in case the outcome resulting from that action is worse for P's welfare than the outcome that would have resulted from the alternative choice. If P can be harmed deontologically without being harmed consequentially, and if deontological harming suffices for tortious harming, then the arguments of this Article

439. See *Reduction of Likelihood*, *supra* note 429, at 504 (distinguishing claim for lost chance from claim for mental distress); Love, *supra* note 432, at 809-10 (distinguishing claim for increased risk from claim for fear of future disease).

440. But see ARIEL PORAT & ALEX STEIN, TORT LIABILITY UNDER UNCERTAINTY 103-10 (2001) (arguing against liability for risk imposition).

441. See Schroeder, *supra* note 423, at 460-77 (arguing that corrective justice permits liability for expected harm). Indeed, Schroeder's position seems to be that corrective justice permits liability, in the amount of expected harm, for actions that do not cause harm *and* are not wrongful in any sense beyond their creating a risk of harm.

442. See Perry, *supra* note 25, at 338 (stating that he is "generally prepared to accept [the] claim" that "tort liability should only be imposed in order to compensate for damage caused").

would *not* preclude assimilating risk imposition to the standard tort paradigm. Still, the proponent of that approach to justifying tort liability for risk imposition has much philosophical work to do—the philosophical work of showing that the propositions contained in the two “if” clauses in the previous sentence are indeed true.

This Article has implications for the *criminal* status of risk imposition similar to those for tort law. Liability for risk imposition is less exceptional as a matter of current criminal law doctrine than it is as a matter of current tort doctrine. The Model Penal Code includes reckless endangerment as a criminal offense: “A person commits a misdemeanor if he recklessly engages in conduct which places or may place another person in danger of death or serious bodily injury.”⁴⁴³ Many jurisdictions follow the Model Penal Code on this score;⁴⁴⁴ many, too, criminalize types of risk imposition more specifically defined, for example drunk driving.⁴⁴⁵ Finally, inchoate crimes—attempts, conspiracies, and solicitations—might be seen as actions that risk choate crimes.⁴⁴⁶ By attempting to commit homicide, I have imposed a risk of death upon the attempted victim. Liability for inchoate crimes is a centerpiece of the Model Penal Code and of enacted criminal codes.⁴⁴⁷

John Stuart Mill famously advanced what might be termed the “Harm Principle”: Only those actions which cause harm to persons other than the actor ought to be criminally proscribed and punished.⁴⁴⁸ Joel Feinberg, one of the preeminent criminal law scholars of our time, and other contemporary theorists have adopted Mill’s view or variants of this view.⁴⁴⁹ If the

443. MODEL PENAL CODE § 211.2, *reprinted in* 2 WAYNE R. LAFAVE & AUSTIN SCOTT, JR., *SUBSTANTIVE CRIMINAL LAW* app. A (1986).

444. *See* 2 WAYNE R. LAFAVE & AUSTIN W. SCOTT, JR., *SUBSTANTIVE CRIMINAL LAW* 312-13 (1986).

445. *See* Douglas N. Husak, *Reasonable Risk Creation and Overinclusive Legislation*, 1 BUFF. CRIM. L. REV. 599, 604-05 (1998) (conceptualizing reckless driving as a “simple inchoate offense,” which targets a risk of harm and is “simple” because it is not defined with reference to an object offense).

446. *Cf. id.* at 602-03.

447. *See* 2 LAFAVE & SCOTT, *supra* note 444, at 1-179.

448. *See* John Stuart Mill, *On Liberty*, in *THE PHILOSOPHY OF JOHN STUART MILL: ETHICAL, POLITICAL AND RELIGIOUS* 187, 197 (Marshall Cohen ed., 1961) (“[T]he only purpose for which power can be rightfully exercised over any member of a civilized community, against his will, is to prevent harm to others.”); *see also* 1 JOEL FEINBERG, *THE MORAL LIMITS OF THE CRIMINAL LAW: HARM TO OTHERS* 10-16 (1984) (discussing Mill’s view).

449. *See* FEINBERG, *supra* note 369, at ix-xx. Feinberg does endorse an

Harm Principle is true, and if the kind of harm required by the principle is a welfare setback understood consequentially—that the “harmed” person be worse off than he would have been absent the harmful action—my Article has important implications regarding the criminality of risk imposition. Prohibitions on reckless endangerment, on drunk driving, on inchoate crimes, and on all other actions of risk imposition are inconsistent with the Harm Principle thus construed.⁴⁵⁰

Lawrence Crocker, attentive to the deep tension between the Harm Principle and the criminality of risk imposition,⁴⁵¹ has distinguished between “harmful” attempts, which involve an “objective risk” of harm (on Crocker’s characterization, a kind of frequentist risk⁴⁵²), and “harmless” attempts, which involve no such risk.⁴⁵³ Crocker claims that the Harm Principle permits the proscription of “harmful” attempts and other actions which may not cause physical harm but impose a frequentist risk of harm on some victims.⁴⁵⁴ My Article undercuts Crocker’s claim. Whether risk is seen in frequentist terms, or alternatively in Bayesian terms that do not refer to the actual beliefs of the victim, suffering a risk is not (without more) a welfare-reducing feature of someone’s life history. Anyone who takes the Harm Principle seriously must contemplate major changes in the structure of existing criminal law.

To be sure, Mill might have been wrong. Law and economics scholars deny that particular actions need be harmful or otherwise blameworthy to trigger criminal punishment. Desert theorists will insist on blameworthiness,

Offense Principle, but “[a]rguably . . . all the types of offensive conduct that liberals would wish to criminalize also involve harm.” A.P. Simester & A. Von Hirsch, *Some Questions about the Offense Principle 1* (unpublished paper, on file with author).

450. Making the causation of fear or other such emotional harm an element of these prohibitions—as in one variant of the crime of assault—would eliminate the inconsistency. *See generally* 2 LAFAVE & SCOTT, *supra* note 444, at 315-17 (discussing the crime of assault).

451. *See* Lawrence Crocker, *Justice in Criminal Liability: Decriminalizing Harmless Attempts*, 53 OHIO ST. L.J. 1057, 1057-63 (1992). More precisely, Crocker argues that an “imposition” principle limits the proper scope of the criminal law: “[F]ree people should not become criminally liable unless they trespass upon someone else’s moral space, that is, unless they impose in some way.” *Id.* at 1060.

452. *See id.* at 1099-1102.

453. *Id.* at 1096-1109.

454. *Id.*

but need not construe that to entail harm.⁴⁵⁵ Our old friend, harmless wrongdoing, surfaces here. Some desert theorists will argue that harmless wrongdoing is blameworthy and properly punished, at least if the actor had a culpable mental state.⁴⁵⁶ Other desert theorists, for example Larry Alexander, argue that culpability absent harm or wrongdoing is sufficient for criminal liability.⁴⁵⁷

It would be foolhardy for me to enter these debates, here, just as it would be overreaching to take a stand on the goals and preconditions of tort liability. My Article does not resolve the criminal or tort law status of risk imposition. Rather, it makes an important contribution to the debate about that status, by showing that one important argument in favor of liability—the argument that risking is itself a kind of consequential harming—must be abandoned.

CONCLUSION

This Article has rigorously analyzed the harmfulness—the welfare impact—of death and the risk of death. The analysis was lengthy and complex because welfare plausibly involves various elements—what I termed Experience, Preference, Value, and Integration—and because the nature of risk is also quite contested. My conclusion was that death is harmful, that Bayesian risk is harmful (at least the actual, first-person variant of Bayesian risk, and at least where linked with appropriate affective states), but that frequentist risk is not. In the final part of the Article, I surveyed the large implications of my analysis for risk regulation. How the CPSC, EPA, FDA, NHTSA, OSHA, and other risk-regulatory agencies should engage in cost-benefit analysis or risk-risk analysis, how they should specify health and safety “thresholds,” how they should take account of distributive considerations, and how these agencies should set priorities all depend on whether death, Bayesian risk, and frequentist risk are harmful. Tort and criminal law doctrines, too, are implicated by the claims advanced here.

455. See Heidi M. Hurd, *What in the World is Wrong?*, 5 J. CONTEMP. LEGAL ISSUES 157, 209 (1994).

456. *Id.* at 208-15.

457. See Larry Alexander, *Crime and Culpability*, 5 J. CONTEMP. LEGAL ISSUES 1, 1-2 (1994) (stating that culpable action, not harm production, is central to criminal law); *id.* at 21-22 (distinguishing between culpable action and wrongdoing in Hurd's sense).

Much more remains to be said about regulatory implications, to be sure. How, exactly, does one calculate the appropriate monetary valuation of death, " V_{death} ," one that follows from a welfarist rather than Kaldor-Hicks construal of cost-benefit analysis and from the premise that frequentist risk and hypothetical Bayesian risk are not welfare setbacks? How should the predicted number of deaths be integrated with cost considerations, if at all, in setting safety thresholds, or performing comparative risk assessment? How does one calculate a monetary value for fear states, for other epistemic-affective hybrids, and for naked belief states if beliefs alone can be harmful? Nor have I discussed the status of risk and harm with respect to moral considerations other than welfare. When is the causation of death not just harmful but *wrongful* to the victim? Can the imposition of a frequentist risk be a harmless wrong, if not a harmful one? My Article has explored part of the normative underpinnings of risk regulation—the welfarist part—but I have nowhere claimed that risk regulation is merely a matter of overall or well-distributed welfare. These vital questions, both foundational and practical, must await another day.