3-2005

Against 'Individual Risk': A Sympathetic Critique of Risk Assessment

Matthew D. Adler
University of Pennsylvania Carey Law School

University of Pennsylvania Law Review

FOUNDED 1852

Formerly
American Law Register

VOL. 153 | MARCH 2005 | NO. 4

ARTICLES

AGAINST "INDIVIDUAL RISK": A SYMPATHETIC CRITIQUE OF RISK ASSESSMENT

MATTHEW D. ADLER[†]

[†] Professor of Law, University of Pennsylvania.

Many thanks to Adam Finkel, Ken Foster, Jeff Gordon, James Hammitt, Andrew Hopkins, Jason Johnston, Ariel Porat, Amy Sinden, and participants in law school faculty workshops at the University of Chicago, Columbia University, Cornell University, University of Michigan, University of Pennsylvania, and Northwestern University for their very helpful comments and questions.

INTRODUCTION

EPA's decision to list a carcinogenic substance as a "hazardous waste," subject to stringent regulation under the Resource Conservation and Recovery Act, depends on the fatality risk that the substance would impose upon highly exposed individuals if discarded in unregulated landfills. If this risk exceeds 1 in 10,000, the substance is listed.[1] EPA's rule for cleanups under the Superfund statute is similarly risk-based: toxic waste dumps are to be remedied so that the lifetime fatality risk to the "maximally exposed individual" from carcinogens in the

[1] *See infra* text accompanying notes 116-19.

dump is within the range of 1 in 10,000 to 1 in 1 million.[2] FDA has traditionally used a 1 in 1 million threshold in determining whether carcinogenic food constituents exempt from the Delaney Clause pose a de minimis safety threat to consumers and thus should be permitted to enter or remain in the food supply.[3] OSHA, which is statutorily authorized to regulate workplace toxins that pose "significant" threats to safety, is more permissive than FDA and EPA but also focuses, in part, on individual fatality risks: the agency has generally followed the rule that carcinogens creating more than a 1 in 1000 risk for any worker are "significant," for statutory purposes, and that toxins creating a substantially smaller risk are not.[4]

In all these cases, health and safety agencies have decided to key regulatory choices to the level of "individual risk" (specifically, the "individual risk" to the maximally exposed individual or some similar construct) without any explicit statutory mandate to do so. But such mandates do exist. A salient one: when Congress in 1990 overhauled section 112 of the Clean Air Act, the section covering carcinogens and other "hazardous air pollutants," it put in place a hybrid regulatory regime that first requires polluters to use the best currently available technology for reducing emissions, and then requires EPA to consider promulgating yet more stringent emissions standards if "excess cancer risks to the individual most exposed to emissions . . . [exceed] one in one million."[5] The 1 in 1 million risk level is also invoked in another provision of the amended Clean Air Act.[6] And when the legislative regime for pesticide licensing was reworked in 1996[7]—the ban on certain carcinogenic pesticides was replaced with a "reasonable certainty

---

[2] *See infra* text accompanying notes 123-27. As explained below, EPA in the Superfund context focuses on risk given a "reasonable maximum exposure" rather than looking to the single maximally exposed individual. *See infra* Part II.A.1.c.

[3] *See infra* text accompanying notes 44; 166-76.

[4] *See infra* text accompanying notes 186-88.

[5] 42 U.S.C. § 7412(f)(2)(A) (2000). *See infra* text accompanying notes 94-102 (discussing Clean Air Act section 112).

[6] *See* 42 U.S.C § 7412(c)(9)(B)(i) (2000) (permitting EPA to remove a category of sources from the list of sources subject to pollution controls under section 112 if "no source in the category . . . emits . . . hazardous air pollutants in quantities which may cause a lifetime risk of cancer greater than one in one million to the individual in the population who is most exposed to emissions of such pollutants from the source").

[7] *See* Food Quality Protection Act of 1996, Pub. L. No. 104-170, 110 Stat. 1489 (codified as amended in scattered sections of 7 and 21 U.S.C.); *infra* text accompanying notes 136-41 (discussing the effect of the Food Quality Protection Act on pesticide regulation).

[of] no harm"[8] standard both for possible carcinogens and for pesticides that might cause other toxic effects—the official House Committee report explained that this new statutory standard ought to be construed as an "individual risk" test:

> [T]he Committee expects . . . that a [pesticide] tolerance will be considered to provide a 'reasonable certainty of no harm' if any increase in lifetime risk, based on quantitative risk assessment using conservative assumptions, will be no greater than 'negligible.' . . . [A] negligible risk [is] a one-in-a-million lifetime risk.[9]

In short: individual fatality risk plays a major role in our current system of health and safety regulation. Some examples have just been provided. Many more will be furnished below. In particular, "individual risk" is absolutely central to federal regulation of toxic chemicals. EPA employs an "individual risk"-based approach in administering all of its major statutes: the Clean Air Act (which addresses toxins present in air), the Clean Water Act and Safe Drinking Water Act (toxins in water), the Resource Conservation and Recovery Act and the Comprehensive Environmental Response, Compensation, and Liability Act (toxins that leach into the ground from waste sites), the Federal Insecticide, Fungicide, and Rodenticide Act (toxic pesticides), and the Toxic Substances Control Act (a backup statute authorizing EPA to take measures not authorized by the media-specific statutes).[10] FDA and OSHA follow a similar approach, as we have seen. But the focus on "individual risk" is not limited to toxins, or to federal agencies. For example, the Nuclear Regulatory Commission (NRC) has long taken the position that the ultimate safety goals governing its licensing and regulation of nuclear power plants partly concern the "individual risk" of immediate death, resulting from an accidental release of radiation, incurred by the average person living near a plant.[11] FDA sets acceptable levels of microbial contaminants in foods with reference to the "individual risk" of illness of a high-end consumer.[12] Although OSHA traditionally focuses on aggregate fatalities or lost days of work in regulating workplace conditions that cause injury (as opposed to illness), it has recently begun to consider the "individual risk" of injury—the rate at which workers in particular industries are injured by electric shock, falls, explosions, fires, and other such indus-

---

[8] 21 U.S.C. § 346a(b)(2)(A)(ii) (2000).

[9] H.R. REP. NO. 104-669, pt. 2, at 41 (1996).

[10] *See infra* Part II.A.

[11] *See infra* text accompanying notes 202-08.

[12] *See infra* text accompanying note 165.

trial accidents.[13] And environmental agencies in some *states* have followed EPA's lead and employ "individual risk" tests in regulating toxins.[14]

What accounts for this regulatory focus on "individual risk"? One answer is tempting, but wrong. The temptation is to say that regulatory agencies inevitably take the maximal level of "individual risk" as the test of safety, at least for substances and activities that cannot be removed from our lives without massive cost. Many, many chemicals cause cancer to animals at large enough doses, and can be predicted to cause some human deaths at actual doses in a sufficiently large group.[15] How else to determine which toxic exposures merit a regulatory response *except* by setting an "individual risk" threshold which seems very low—say, 1 in 1 million to the maximally exposed individual—and taking that as the trigger for regulatory intervention? But this response overlooks a crucial deficit in "individual risk" tests of this kind: their insensitivity to population size. Compare an isolated toxic waste dump that (under worst-case modeling) leaches contaminants to a radius of ten miles, affecting a population of 10,000; a workplace toxin employed in certain industries, to which one million workers are exposed; and a chemical in drinking water that is consumed by 100 million. For simplicity, assume that in each case every person in the exposed population incurs a 1 in 1 million risk of dying from the hazard. Then in the waste-dump case it is overwhelmingly likely that the hazard will cause *no* fatalities; in the workplace case it is reasonably likely that the hazard will cause one or more fatalities, with one incremental death the expected outcome; and in the drinking-water case it is overwhelmingly likely that the hazard will cause one or more fatalities, with 100 incremental deaths the expected outcome.[16]

---

[13] *See infra* text accompanying notes 191-97.

[14] *See infra* note 89.

[15] *See* John D. Graham, *Historical Perspective on Risk Assessment in the Federal Government*, 102 TOXICOLOGY 29, 33-35 (1995) (explaining that carcinogens traditionally have been seen to lack safety "thresholds," and describing how this no-threshold view prompted regulatory agencies to adopt "individual risk" tests for determining when exposures to carcinogens are de minimis); Dennis J. Paustenbach, *Retrospective on U.S. Health Risk Assessment: How Others Can Benefit*, 6 RISK 283, 284 (1995) ("[O]ver 300 of about 5,000 chemicals routinely used in industry have been labeled carcinogens as a result of animal studies.").

[16] *See, e.g.*, JAMES T. HAMILTON & W. KIP VISCUSI, CALCULATING RISKS?: THE SPATIAL AND POLITICAL DIMENSIONS OF HAZARDOUS WASTE POLICY 1-23 (1999) (criticizing EPA's policy for remedying Superfund sites, in part because the policy employs "individual risk" criteria that are insensitive to the size of the populations exposed to the sites).

Risk assessors typically distinguish between "individual risk"—the risk of death borne by a particular individual, either a named person or someone identified by her exposure characteristics—and "population risk."[17] "Population risk" (also sometimes called "societal risk") is the total number of fatalities resulting from a toxin, a hazardous activity, or some other threat to human life. To quote a leading textbook on risk assessment:

> [Risk assessments typically] 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 of fatalities as a function of location or population subgroup . . . .[18]

Regulatory agencies might use the level of "population risk," rather than the level of "individual risk," as their measure of health and safety.[19] This is true both for agencies operating under statutes that accord high priority to the avoidance of death, illness, and injury, as opposed to other goals, as well as for agencies operating under cost-

---

[17] *See* VINCENT T. COVELLO & MILEY W. MERKHOFER, RISK ASSESSMENT METHODS: APPROACHES FOR ASSESSING HEALTH AND ENVIRONMENTAL RISKS 230-34 (1993); NAT'L RESEARCH COUNCIL, SCIENCE AND JUDGMENT IN RISK ASSESSMENT 69-70 (1994); Pamela R.D. Williams & Dennis J. Paustenbach, *Risk Characterization*, *in* HUMAN AND ECOLOGICAL RISK ASSESSMENT: THEORY AND PRACTICE 293, 322-23 (Dennis J. Paustenbach ed., 2002).

[18] COVELLO & MERKHOFER, *supra* note 17, at 231. The term "population risk" is, admittedly, ambiguous, and is not always used to mean the total number of deaths. It might be used, instead, to mean the total number of individuals at various levels of "individual risk," *see id.*, or more generally to refer to a variety of population-size sensitive criteria for evaluating hazards. In this Article, I use "population risk" to mean aggregate deaths or other adverse cases (illnesses, injuries), and nothing else.

[19] *See* Frank B. Cross et al., *Discernible Risk—A Proposed Standard for Significant Risk in Carcinogen Regulation*, 43 ADMIN. L. REV. 61, 73-75 (1991); *cf.* National Emissions Standards for Hazardous Air Pollutants, 54 Fed. Reg. 38,044, 38,044-46 (EPA Sept. 14, 1989) (considering, but rejecting, a proposal to define a safe emission level of air pollutants as that amount causing one cancer case per year).

benefit statutes or other "balancing" statutes that permit a wider array of considerations to trump the goal of physical integrity.[20] Indeed, as we shall see, federal programs concerned with safety rather than health hazards generally seem to focus on "population risk" rather than "individual risk," and even health threats such as toxins, radiation, and pathogens are sometimes regulated with reference to "population risk."[21]

So the question just posed remains unanswered: why do EPA, OSHA, NHTSA, and many other agencies, federal and state, employ some variation on the "individual risk" construct—be it "individual risk" to the maximally exposed individual, to a highly exposed individual, to the median or average individual, or to some other person—in administering statutes that make human health and safety a (high-priority or ordinary-priority) regulatory goal? Why not use a fatality-based metric instead, for example one that looks at the effect of regulatory intervention in reducing the total number of deaths caused by fatal illnesses or injuries?

One plausible answer points to a seminal 1980 Supreme Court case, *Industrial Union Department v. American Petroleum Institute*.[22] This case, more than any other single event, triggered the rapid growth of risk assessment in the federal government.[23] And it may well have caused or at least supported the regulatory focus on "individual risk" rather than "population risk." In *Industrial Union*, a plurality of the Court invalidated an OSHA regulation lowering the maximum permissible workplace exposure to benzene, a carcinogen, from ten parts per million (ppm) to one ppm. The Occupational Safety and Health Act, as the plurality read it, authorized OSHA only to regulate "significant" risks—not to ban workplace chemicals or activities based on the mere possibility of an injury or fatality.[24] Note that this aspect of the

---

[20] *See* Matthew D. Adler, *Risk, Death and Harm: The Normative Foundations of Risk Regulation*, 87 MINN. L. REV. 1293, 1391-92 (2003) (describing the variation in risk regulation statutes, some of which give priority to health and safety, others of which permit or require cost-benefit analysis).

[21] *See infra* Part II.C.

[22] 448 U.S. 607 (1980) (plurality opinion).

[23] *See, e.g.*, Cross et al., *supra* note 19, at 68 (describing *Industrial Union* as "[t]he seminal decision in carcinogen regulation"); John D. Graham, *The Risk Not Reduced*, 3 N.Y.U. ENVTL. L.J. 382, 386 (1995) (describing *Industrial Union* as "[t]he turning point for quantitative risk assessment"); Randall S. Wentsel, *Application of Risk Assessment in Policy and Legislation in North America*, *in* HANDBOOK OF ENVIRONMENTAL RISK ASSESSMENT AND MANAGEMENT 261, 262 (Peter Calow ed., 1998) (stating that *Industrial Union* generally encouraged federal agencies to engage in risk assessment).

[24] *See* 448 U.S. at 642-52.

Industrial Union opinion does not entail a preference for regulatory attention to "individual risk." After all, in implementing the "significant risk" threshold, OSHA *could* look to aggregate premature deaths resulting from the workplace toxin or activity at issue, not "individual risk" to the maximally exposed or average worker. But, in the final portion of the opinion, Justice Stevens suggested that the statutory requirement of "significant risk" be implemented through an "individual risk" test.

> Contrary to the Government's contentions, imposing a burden on the Agency of demonstrating a significant risk of harm will not strip it of its ability to regulate carcinogens, nor will it require the Agency to wait for deaths to occur before taking any action. First, the requirement that a "significant" risk be identified is not a mathematical straitjacket. . . . Some risks are plainly acceptable and others 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.[25]

To this day, OSHA carefully follows this dictum from *Industrial Union*. The agency still uses the 1 in 1000 level of "individual risk" identified by Justice Stevens as its cutoff for regulating a workplace carcinogen.[26] More generally, although EPA and FDA do not employ that cutoff—a cutoff which Stevens characterized as a reasonable construal of the Occupational Safety and Health Act, not the only acceptable construal—the Stevens dictum may well have prodded EPA, FDA, and other agencies to focus on individual, not population, risk.[27]

A second explanation for the wide use of "individual risk" tests by regulatory agencies points to the norms of risk assessment. Risk assessment is a set of techniques for quantifying health and safety threats, paradigmatically but not exclusively threats from toxic chemicals.[28] Risk assessments are very widely employed by government

---

[25] *Id.* at 655.

[26] *See infra* text accompanying notes 186-88.

[27] *See* Natural Res. Def. Council v. EPA, 824 F.2d 1146, 1164 (D.C. Cir. 1987) (en banc) (quoting *Industrial Union* in construing the "ample margin of safety" language of the Clean Air Act); National Emission Standards for Hazardous Air Pollutants, 54 Fed. Reg. 38,044, 38,044-46 (EPA Sept. 14, 1989) (responding to *National Resources Defense Council* by adopting a test focused, at least presumptively, on "individual risk").

[28] Good general introductions to risk assessment methodology include: COVELLO & MERKHOFER, *supra* note 17; DOUGLAS J. CRAWFORD-BROWN, RISK-BASED ENVIRONMENTAL DECISIONS: CULTURE AND METHODS (1999); WILLIAM H. HALLENBECK,

agencies in setting priorities and evaluating interventions,[29] and are also used in other contexts.[30] Risk assessment techniques have become quite standardized, both as a result of governmental standardization (for example, EPA's various guidelines)[31] and because of the standardization internal to the emerging professional community of risk assessors.[32] The core of risk assessment for toxins consists of two steps: drawing a dose-response curve and predicting individual exposures. As we shall see, dose-response curves and exposure assessments *can* be integrated to generate predictions of aggregate deaths—and sometimes are—but they are also naturally deployed to generate predictions of "individual risk."[33]

In any event, whether as a result of *Industrial Union*, the professionalized techniques of risk assessment, or other factors, governmental agencies in the United States, in a host of different contexts, employ the test of "individual risk" to the maximally exposed individual or some similar test as a criterion for regulatory choice. As traditional economic regulation has become less important, particularly at the federal level, an increasing proportion of regulatory activity concerns

---

QUANTITATIVE RISK ASSESSMENT FOR ENVIRONMENTAL AND OCCUPATIONAL HEALTH (2d ed. 1993); Susan P. Felter et al., *Assessing Risks to Human Health from Chemicals in the Environment*, *in* HANDBOOK OF ENVIRONMENTAL RISK ASSESSMENT AND MANAGEMENT, *supra* note 23, at 9, 9-23; Dennis J. Paustenbach, *Hazard Identification*, *in* HUMAN AND ECOLOGICAL RISK ASSESSMENT: THEORY AND PRACTICE, *supra* note 17, at 85, 85-149; Dennis J. Paustenbach, *Exposure Assessment*, *in* HUMAN & ECOLOGICAL RISK ASSESSMENT: THEORY AND PRACTICE, *supra* note 17, at 189, 189-291; Williams & Paustenbach, *Risk Characterization*, *supra* note 17.

[29] On the role of risk assessment in regulation, see generally CARL F. CRANOR, REGULATING TOXIC SUBSTANCES: A PHILOSOPHY OF SCIENCE AND THE LAW 103-51 (1993); QUANTITATIVE RISK ASSESSMENT IN REGULATION (Lester B. Lave ed., 1982); Howard Latin, *Good Science, Bad Regulation, and Toxic Risk Assessment*, 5 YALE J. ON REG. 89, 95-126 (1988); Mark Eliot Shere, *The Myth of Meaningful Environmental Risk Assessment*, 19 HARV. ENVTL. L. REV. 409, 417-68 (1995); Symposium, *Risk Assessment in the Federal Government*, 3 N.Y.U. ENVTL. L.J. 251 (1995); *infra* text accompanying notes 40-53, Part II, and sources cited *infra* note 88.

[30] *See* Charles A. Pittinger et al., *Corporate Chemical Management: A Risk-Based Approach*, *in* HANDBOOK OF ENVIRONMENTAL RISK ASSESSMENT AND MANAGEMENT, *supra* note 23, at 379, 379-401.

[31] *See* Dennis J. Paustenbach, *Primer on Human and Environmental Risk Assessment*, *in* HUMAN AND ECOLOGICAL RISK ASSESSMENT: THEORY AND PRACTICE, *supra* note 17, at 3, 28-29 (citing EPA's main guidance documents).

[32] *See* Graham, *supra* note 23, at 386-87 (describing the emergence of an institutional framework for risk analysts, such as journals and scientific organizations); Paustenbach, *supra* note 15, at 286-87, 289-90 (describing the emergence of shared understandings among risk analysts concerning various aspects of the risk assessment process).

[33] *See infra* text accompanying note 63.

the avoidance of death and, to a lesser extent, nonfatal injury and disease.[34] In turn, "individual risk" tests have become a linchpin of government's health and safety efforts. This is misguided. In this Article, I shall launch a sustained critique of the use of "individual risk" tests by health and safety agencies. This critique does not depend on controversial normative commitments. My own commitments are welfarist and consequentialist,[35] and I have argued elsewhere in favor of cost-benefit analysis (CBA).[36] It is true that "population risk," not "individual risk," is the input to CBA as currently practiced.[37] But it emerges that normative frameworks directing agencies to accord higher priority to safety than CBA would countenance are also best specified in terms of "population risk" or cognate tests. Or so I shall argue below.

Parts I and II of the Article set the stage. Part I is a primer on risk assessment. It explains the structure of risk assessment, describes its rise to prominence as a tool for health and safety regulators, and then explores the nature of the "individual risk" numbers so central to the technique. What exactly does it mean to say that some toxin, substance, activity, or, more abstractly, some object or event imposes a 1 in *x* risk of death upon a particular individual? What concept of "risk" is being invoked here? The standard interpretation of the "individual risk" numbers generated by risk assessment invokes the *frequentist* view of risk. On the frequentist view, to say that *E* imposes a 1 in *x* risk of death upon *P* is to say this: over the long run, when people similar to *P* are exposed to events similar to *E*, a 1 in *x* fraction of those individuals will die prematurely as a result of those exposures.

Part II describes, in detail, the widespread use of "individual risk" tests by federal agencies. The practices I describe should be familiar to environmental lawyers, food and drug specialists, workplace safety scholars, and others who have specialized knowledge about EPA, FDA,

---

[34] *See* Robert L. Rabin, *Federal Regulation in Historical Perspective*, 38 STAN. L. REV. 1189, 1317-18 (1986).

[35] *Cf.* Matthew D. Adler, *Beyond Efficiency and Procedure: A Welfarist Theory of Regulation*, 28 FLA. ST. U. L. REV. 241 (2000) (arguing for the moral weight of overall well-being, but bracketing the question whether morality has a consequentialist structure or, instead, incorporates some deontological norms).

[36] *See* Matthew Adler, *Incommensurability and Cost-Benefit Analysis*, 146 U. PA. L. REV. 1371 (1998); Matthew D. Adler & Eric A. Posner, *Implementing Cost-Benefit Analysis when Preferences Are Distorted*, *in* COST-BENEFIT ANALYSIS: LEGAL, ECONOMIC, AND PHILOSOPHICAL PERSPECTIVES 269 (Matthew D. Adler & Eric A. Posner eds., 2001); Matthew D. Adler & Eric A. Posner, *Rethinking Cost-Benefit Analysis*, 109 YALE L.J. 165 (1999).

[37] *See infra* text accompanying note 300.

OSHA, or similar agencies. These important practices will not be as familiar—I hazard to guess—for public law generalists, law and economists, legal philosophers, and other scholars who may have a deep interest in the regulatory state but have not read the latest issue of *Risk Analysis* or the latest version of EPA's *Guidelines for Carcinogen Risk Assessment*.[38] And even risk-regulation specialists might be surprised to learn just how widespread the focus on "individual risk" is. In any event, Part II seeks to show that administrative decision making across a wide swath of significant governmental programs conforms to a problematic recipe. This recipe specifies the safety of a workplace, a toxic dump, a water source, a radioactive emission, a consumer product, or some other regulatory target in terms of the frequentist "individual risk" the targeted substance, activity, or product imposes on the maximally exposed individual or some other person (with 1 in 1 million most widely used as the "safe" level).

The remainder of the Article provides a rigorous, normative critique of the frequentist "individual risk" tests described in Part II. I seek to show that "individual risk" in the frequentist sense is normatively irrelevant across a range of plausible moral theories. Part III looks at *welfarist consequentialism*: the moral view undergirding welfare economics and cost-benefit analysis. It argues that the kind of risk relevant to welfarist consequentialism is Bayesian risk, not frequentist risk. Bayesian risks are measures of belief, not measures of frequency. On the Bayesian view, to say that some individual has a 1 in x probability of death means that the risk analyst, the individual herself, or someone else believes to degree 1 in x that the individual will die. Part IV explores the subtle, but crucial differences between Bayesian and frequentist risk.

Part V moves beyond welfare consequentialism and examines alternative moral views: *safety-focused* views that accord special priority to physical integrity; *deontological* views that recognize the existence of moral rights (specifically, a right not to be killed and perhaps an independent right not to be put at risk of death); *contractualist* views that evaluate governmental choices by asking whether citizens in a suitable, hypothetical contracting scenario would approve the choices; and *democratic* views that see democratic responsiveness (including responsiveness to popular judgments about risk) as morally important. Part

[38] *See* Notice of Availability and Opportunity to Provide Comment on the Draft Final Guidelines for Carcinogen Risk Assessment, 68 Fed. Reg. 10,012 (EPA Mar. 3, 2003); Guidelines for Carcinogen Risk Assessment, 51 Fed. Reg. 33,992 (EPA Sept. 24, 1986).

V argues that none of these moral views warrants regulatory attention to "individual risk" in the frequentist sense.

Part VI shifts critical focus to a different feature of the current regulatory practices described in Part II. Those practices are, in effect, doubly misguided. First, they make "individual risk," in the frequentist rather than Bayesian sense, a determinant of regulatory choice. Second, they are insensitive to population size. Whether regulators should intervene to abate some hazard depends, morally, on the number of persons at risk from the hazard. A specialty food consumed by a very small group, an industrial toxin with which many more workers come into contact, and an airborne pollutant that we all breathe might impose the very same "individual risk" on the maximally exposed, high-end, median, and average exposed individual. But the morally warranted regulatory responses in these cases will, or at least may, be very different. Part VI argues that both welfare consequentialism and alternative moral views (safety-focused, deontological, contractualist, and democratic views) demand risk assessment procedures that are sensitive to population size.

What this Article offers, in short, is a sympathetic critique of risk regulation and risk assessment. Much of the legal scholarship in this area is more radically critical. Regulation guided by risk assessment is allegedly flawed to the core—for example, because it is undemocratic, or because it is beset with uncertainties about the mechanisms of cancer, dose-response relationships, and exposure pathways. The very enterprise of quantifying safety is seen as misguided.[39] I do not believe that the very enterprise of quantifying safety is misguided. Risk assessment represents a giant leap forward for public rationality, in my view. Dose-response curves and exposure assessments are, properly, central for the regulation of toxins; parallel techniques are central for agencies that focus on other threats to life and limb. But these impressive techniques should be used to illuminate what is truly at stake in risk regulation, not to distract us with morally unimportant information. Risk regulation needs to be changed in two ways. First, it should adopt a new understanding of risk, Bayesian rather than fre-

[39] Scholarship in this vein includes: Adam Babich, *Too Much Science in Environmental Law*, 28 COLUM. J. ENVTL. L. 119 (2003); Jeremy D. Fraiberg & Michael J. Trebilcock, *Risk Regulation: Technocratic and Democratic Tools for Regulatory Reform*, 43 MCGILL L.J. 835 (1998); Eileen Gay Jones, *Risky Assessments: Uncertainties in Science and the Human Dimensions of Environmental Decisionmaking*, 22 WM. & MARY ENVTL. L. & POL'Y REV. 1 (1997); Latin, *supra* note 29; Thomas O. McGarity, *A Cost-Benefit State*, 50 ADMIN. L. REV. 7 (1998); Shere, *supra* note 29; Wendy E. Wagner, *The Science Charade in Toxic Risk Regulation*, 95 COLUM. L. REV. 1613 (1995).

quentist. Second, it should adopt choice criteria that are sensitive to population size—paradigmatically, "population risk" criteria. Risk assessment, in turn, must be reworked so that it can inform regulatory choice thus revised.

I. RISK ASSESSMENT: A PRIMER

A. *History and Structure*

Risk assessment, generically, is a set of techniques for quantifying the fatalities or fatality risks resulting from various hazards. These techniques can also be used to quantify nonfatal illness or injury, or the risk of nonfatal illness or injury. However, because death is the central and paradigmatic harm addressed by health and safety regulators, my presentation will focus there. The best-developed variant of risk assessment, in current regulatory practice, is *toxic* risk assessment: a quantitative description of the fatalities and fatality risks caused by toxic chemicals. But risk assessment with respect to a much wider array of death's causes is also possible and, to some extent, practiced.

Let's start with the toxins. Toxic risk assessment is, in effect, quantitative toxicology and dates from the nineteenth century.[40] Toxic risk assessment by U.S. governmental entities is more recent than that, but still has a substantial history.[41] Toxic risk assessment at the federal level was pioneered by FDA. This agency is charged with implementing a statute that generally precludes foods containing "poisonous or deleterious" substances[42] and that imposes even stricter standards on "food additives": such additives must be "safe,"[43] and carcinogenic food additives are flatly prohibited by the (in)famous "Delaney Clause."[44] FDA began systematically to engage in the risk assessment of the noncancer toxicity of food constituents in the 1950s, developing the so-called NOAEL/safety factor method which is the

---

[40] *See* Graham, *supra* note 15, at 31.

[41] *See id.* at 33-40.

[42] 21 U.S.C. § 342(a) (2000).

[43] *Id.* § 348(c)(3)(A).

[44] *See* Food Additives (Delaney) Amendment of 1958, Pub. L. No. 85-929, sec. 2, § 409(c)(3)(A), 72 Stat. 1784, 1786 (codified at 21 U.S.C. § 348(c)(3)(A) (2000)) ("[N]o additive shall be deemed to be safe if it is found to induce cancer when ingested by man or animal, or if it is found, after tests which are appropriate for the evaluation of the safety of food additives, to induce cancer in man or animal . . . .").

standard method for noncarcinogens today.[45] As for the threat of cancer: there are various escape routes around the absolutism of the Delaney Clause—for example, FDA takes the position that the Clause does not apply to the nonfunctional carcinogenic constituents of additives[46]—and in the 1970s FDA commenced a practice of quantifying the potency of certain food carcinogens. The 1 in 1 million cutoff for individual cancer risk derives from FDA practice during this period.[47]

The widespread use of toxic risk assessment at other federal agencies began in the 1980s. This development had multiple triggers, including three apparent ones: (1) the *Industrial Union* case, which forced OSHA to follow FDA's lead and, more generally, made clear that the courts would not permit federal agencies engaged in health and safety regulation to impose large costs on the regulated entities without some effort to justify such costs through a quantification of benefits, even in cases (as with OSHA) where the underlying statute was quite pro-safety;[48] (2) the 1983 appointment, as EPA administrator, of William Ruckleshaus, who made risk assessment his top priority and actually succeeded in infusing such techniques into administrative routines throughout the large EPA bureaucracy;[49] and (3) the publication, also in 1983, of a seminal study by the National Research Council, *Risk Assessment in the Federal Government: Managing the Process*[50] (the so-called *Red Book*), which further popularized the practice of risk assessment and, perhaps more importantly, did much to standardize it. By the 1990s, risk assessment had become such a familiar feature of the regulatory landscape that OMB, in its guidance to federal agencies regarding Executive Order 12,866,[51] instructed that the regulatory impact analysis required by this Executive Order prior to the issuance of major rules should include a "risk assessment":

---

[45] *See* Graham, *supra* note 15, at 32. See *infra* Part II.B.1 for a detailed description of current FDA risk assessment practices.

[46] *See infra* text accompanying notes 167-68.

[47] *See* Graham, *supra* note 15, at 34-35; Joseph V. Rodricks et al., *Evaluating the Safety of Carcinogens in Food—Current Practices and Emerging Developments*, 46 FOOD DRUG COSM. L.J. 513, 533-35 (1991).

[48] *See supra* text accompanying notes 22-27.

[49] *See* Graham, *supra* note 15, at 39.

[50] COMM. ON THE INSTITUTIONAL MEANS FOR ASSESSMENT OF RISKS TO PUB. HEALTH, NAT'L RESEARCH COUNCIL, RISK ASSESSMENT IN THE FEDERAL GOVERNMENT: MANAGING THE PROCESS (1983) [hereinafter RED BOOK].

[51] Exec. Order No. 12,866 § 6(a)(3), 3 C.F.R. 638, 645-46 (1994), *reprinted in* 5 U.S.C. § 601 (2000).

> Estimating the benefits and costs of risk-reducing regulations [requires, inter alia] . . . a risk assessment that . . . characterizes the probabilities of occurrence of outcomes of interest . . . .
>
> . . . The risk assessment should generate a credible, objective, realistic, and scientifically balanced analysis; present information on hazard, dose-response, and exposure (or analogous material for non-health assessments); and explain the confidence in each assessment . . . .[52]

So risk assessment is now standard practice for federal agencies that regulate toxins, as well as other health and safety agencies, at the major rulemaking stage. But the practice of toxic risk assessment is really much broader than that. For example, the overwhelming majority of EPA risk assessments do not involve major rules, but other categories of administrative decision, such as clean-up decisions with respect to individual Superfund sites.[53]

The *Red Book* framework for toxic risk assessment has been the canonical framework[54] since its publication and runs as follows. There are four parts to toxic risk assessment: hazard identification, dose-response

---

[52] OFFICE OF MGMT. & BUDGET, ECONOMIC ANALYSIS OF FEDERAL REGULATIONS UNDER EXECUTIVE ORDER 12866, at § III.A.4 (1996), *available at* http://www.whitehouse.gov/omb/inforeg/riaguide.html.

[53] Robert Hahn and coauthors collected all the regulatory impact analyses that were prepared for nontransfer major rules finalized between April 1996 and July 1999. Since impact analyses are typically not undertaken for transfer rules, this collection should include most of the major rule impact analyses completed during this three-year period. *See* 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, 862-65 (2000). Of the forty-eight regulatory impact analyses, twenty-three were issued by EPA, or less than eight per year. *See id.* app. 2, at 881-85 (listing impact analyses by issuing agency).

By contrast, Hamilton and Viscusi identified at least 268 Superfund sites where "Records of Decisions" were signed during 1991 or 1992. *See* HAMILTON & VISCUSI, *supra* note 16, app. A, at 245. At most of these sites, presumably, site-specific risk assessments were undertaken. *See id.* app. A, at 247. Hamilton and Viscusi ultimately analyzed a subsample of 150 of the 268 sites; it appears that every site in their subsample included a site-specific risk assessment. *See id.* For another example of an EPA program in which many risk assessments are undertaken outside the major rule context, see Lorenz R. Rhomberg, *A Survey of Methods for Chemical Health Risk Assessment Among Federal Regulatory Agencies*, 3 HUM. & ECOLOGICAL RISK ASSESSMENT 1029, 1126 (1997), noting that EPA undertakes more than 2000 risk assessments annually, albeit rudimentary ones, in connection with the premanufacture notice requirement of the Toxic Substances Control Act.

[54] *See, e.g.*, Paustenbach, *supra* note 15, at 289 ("Risk assessment has (by convention) been separated into four subdisciplines: hazard identification, dose-response assessment, exposure assessment and risk characterization."); Rhomberg, *supra* note 53, at 1085 (stating that EPA and other agencies see the *Red Book* as providing "overarching guiding principles").

assessment, exposure assessment, and risk characterization.[55] Hazard identification is a preliminary step: the risk analyst verifies that the allegedly toxic substance is indeed a toxin, that there is sufficient evidence of a causal link to disease and death.[56] If so, the analysis moves on to the two central parts of the risk-assessment inquiry, namely dose-response assessment and exposure assessment. Dose-response assessment means quantifying the link between different doses of the toxin and premature death. This inquiry is, in effect, *physiological*: it seeks to determine how frequently the ingestion, inhalation, or dermal uptake of the toxin, into humans' bodies, leads to cancer or other fatal illnesses. This physiological inquiry is almost always grounded in two types of data—rodent bioassays, in which the differing rates of fatal illness in groups of rodents fed different doses of the toxin are measured, and human epidemiological data—and eventuates in a dose-response curve.[57] The X-axis of the curve represents human doses of the toxin; the Y-axis, the risk to a person exposed to that dose of dying prematurely as a result.

Figure 1: Dose-Response Curve for Carcinogen C

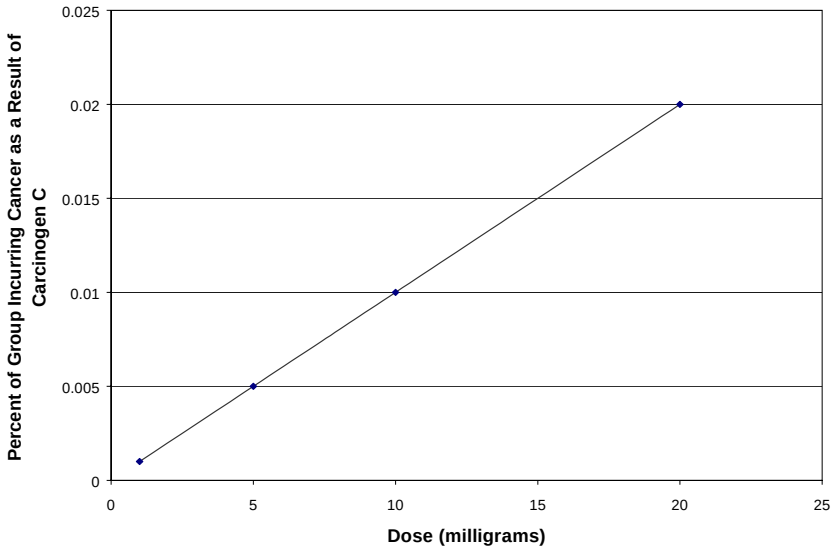

[55] *See* RED BOOK, *supra* note 50, at 19-20.

[56] *See* Paustenbach, *Hazard Identification*, *supra* note 28, at 85.

[57] Good surveys of dose-response assessment methods include: COVELLO & MERKHOFER, *supra* note 17, at 137-79; CRAWFORD-BROWN, *supra* note 28, at 95-121; HALLENBECK, *supra* note 28, at 33-103; Suresh H. Moolgavkar & E. Georg Leubeck, *Dose-Response Modeling for Cancer Risk Assessment*, *in* HUMAN AND ECOLOGICAL RISK ASSESSMENT: THEORY AND PRACTICE, *supra* note 17, at 151, 151-88.

Exactly how to draw this curve, based on the animal or epidemiological evidence, is a technical (but important!) issue in risk assessment.

While dose-response assessment is physiological, exposure assessment is *topographical* and *demographic*. The aim here is to characterize the pattern of exposures to the toxin that will occur under various contingencies (for regulatory purposes, as a result of different regulatory options, including the status quo option of inaction).[58] This assessment depends on where the toxin is located; on how much toxin currently exists, and would be produced under various contingencies, at the source; on the toxin's so-called "fate and transport," i.e., how the toxin migrates through the air, the water, and other environmental pathways; and on how the human population is distributed at varying distances from the source of the toxin. What "source" means depends, of course, on the toxin and the regulatory program. It might be a single waste dump, a group of smokestacks (for example all smokestacks in factories in a given industrial category), a food type (in which case the relevant pathway is direct ingestion by food consumers), containers of hazardous workplace chemicals (in which case the relevant pathway is air transport to workers or direct worker contact with the containers), and so on. A relatively complete exposure assessment will predict the dose of the toxin that each member of the population will receive as a result of the evaluated source. Typically, members of the population will not be identified by name, but rather by their doses. That is: a relatively complete exposure assessment will produce a predicted distribution, by numbers and percentiles, of lifetime doses resulting from the analyzed source, for the status quo option of regulatory inaction and, ideally, for each regulatory contingency being assessed.[59]

---

[58] *See* COVELLO & MERKHOFER, *supra* note 17, at 91-125; CRAWFORD-BROWN, *supra* note 28, at 69-93; ALISON C. CULLEN & H. CHRISTOPHER FREY, PROBABILISTIC TECHNIQUES IN EXPOSURE ASSESSMENT: A HANDBOOK FOR DEALING WITH VARIABILITY AND UNCERTAINTY IN MODELS AND INPUTS (1999); Paustenbach, *Exposure Assessment*, *supra* note 28.

[59] An even more complete exposure assessment will present, and ascribe probabilities to, different possible exposure distributions, rather than merely presenting the most likely distribution. *See, e.g.*, CULLEN & FREY, *supra* note 58, at 217-37, 304-06.

Figure 2: Exposure Assessment for Carcinogen C in Source S

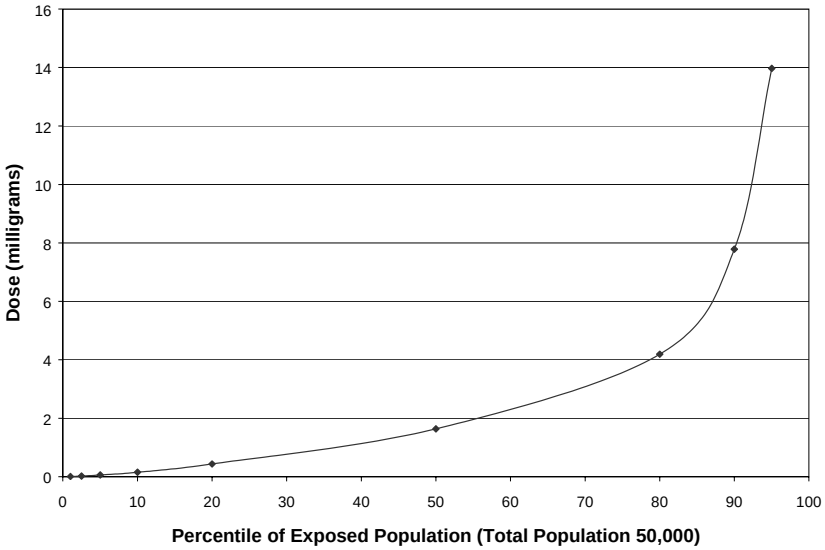

To be sure, the toxic exposure assessments produced by regulatory agencies or their contractors are often not this detailed. If the agency's program focuses on risk to the maximally exposed individual,[60] then the full pattern of dosages that will occur in the status quo, or as a result of various regulatory interventions, is irrelevant. The analyst might estimate the dosage received by the maximally exposed individual by generating a full distribution of doses across the population, then using a very high percentile as maximum exposure; or she might do so more directly by maximizing the parameters underlying her exposure model and determining what dosage results. Concretely, this might mean estimating the maximum exposure to a toxic air pollutant emitted from a factory by looking at the exposure incurred by the person living closest to the factory, on the conservative assumptions that he lives there for his entire lifetime and that his inhalation rate is at the high end of the population distribution of such rates.[61]

---

[60] *See infra* text accompanying notes 94-102 (describing Clean Air Act section 112, which focuses on risk to the maximally exposed individual).

[61] To be sure, most federal agencies now focus on risk to average or highly exposed individuals, not maximally exposed individuals, *see infra* Parts II.A-II.B, but truncated exposure assessments might also be employed in these contexts by using point estimates (average or high-end) of the inputs to the exposure model to generate a

The final stage of risk assessment, risk characterization, is the prime focus of this Article. "Risk characterization" means combining the dose-response assessment (which correlates doses and fatality risks) and the exposure assessment (which predicts doses, across the population or at least for some segment) so as to generate a prediction of the fatalities and fatality risks resulting from the toxin under various contingencies.[62] The risk assessment jargon for total fatalities, as I have already noted, is "population risk"; the jargon for the risk of death incurred by one or another individual is "individual risk." Using the dose-response assessment and exposure assessment to predict "population risk" is somewhat laborious. In general, to do that, the analyst needs a full population distribution of doses, and even then the analyst cannot simply "read" an estimate of total deaths off the dose-response curve, but instead must use probability theory to generate a probability distribution of total deaths and then a point estimate of "population risk" equaling the mean number of total deaths.[63]

Generating a prediction of "individual risk" is more straightforward. For example, if the analyst possesses a full or truncated exposure assessment showing the dosage of the toxin to the "maximally exposed" individual, then the "individual risk" to that person is simply the risk corresponding to that dosage given by the dose-response curve. And if the analyst possesses a full or truncated exposure assessment showing the exposure to the "representative" individual—the person at the median or mean of the dosage distribution—then the "individual risk" incurred by this "ordinary Joe" is the risk for *his* dosage predicted by our physiological graph, the dose-response curve.

I have focused, to this point, on toxic risk assessment, since the standard practices in this area readily lend themselves to predictions

---

point estimate of average or high-end dose. *See* CULLEN & FREY, *supra* note 58, at 2-8 (distinguishing between "deterministic" exposure assessment, where point estimates of inputs to exposure models are employed, and "probabilistic" exposure assessment, where variability in, as well as uncertainty about, exposure is characterized).

[62] *See* COVELLO & MERKHOFER, *supra* note 17, at 203-37; Williams & Paustenbach, *supra* note 17, at 322-24.

[63] *Cf.* COMM. ON TECHNICAL BASES FOR YUCCA MOUNTAIN STANDARDS, NAT'L RESEARCH COUNCIL, TECHNICAL BASES FOR YUCCA MOUNTAIN STANDARDS 47-59, 59 (1995) (arguing that the safety of Yucca Mountain repository for high-level waste should be assessed in terms of "individual risk" to maximally exposed individuals, rather than "population risk," in part given the "great uncertainty about the number of health effects that would be imposed on the global population"). In the special case where the analyst has estimated the average dose, and the dose-response curve is linear, generating a prediction of total deaths is easier. *See* Guidelines for Exposure Assessment, 57 Fed. Reg. 22,888, 22,901 (EPA May 29, 1992).

of "individual risk" (as I have tried to show) and since (as we shall see) many of the important cases of "individual risk"-based decision making by agencies involve toxins. But nontoxic risk assessment—quantitative assessment of the wide variety of threats to human health and safety posed by substances or activities other than toxic chemicals—is also quite important in governmental practice. A salient example here is radiation risk assessment, as pioneered at the federal level by the Nuclear Regulatory Commission. The famous *Reactor Safety Study* (WASH-1400),[64] commissioned by NRC's predecessor agency and published in 1975, was the first full-blown, probabilistic evaluation of the core damage accidents at nuclear reactors that could lead to dangerous releases of radiation.[65] This study, together with the Three Mile Island accident four years later, prompted NRC to make risk assessment integral to the licensing and regulation of nuclear plants.[66] Reactor risk assessment divides, very roughly, into two parts: (1) evaluating the probability of different types of releases (releases of various amounts of various radioactive isotopes); and (2) evaluating the safety threat for any given release.[67] This second component closely tracks the standard methodology for toxic risk assessment. For any given release, an exposure assessment can be performed evaluating the possible fate and transport of the released isotopes and the exposure to those substances of various members of the population; this information, when combined with a dose-response curve correlating radiation doses with an individual's risk of dying as a result of the exposure, can be used to predict the "individual risk" of death imposed on different individuals by any given release and therewith (if desired) a prediction of "population risk."[68]

With the exception of radiation, the norms of nontoxic risk assessment are less well established than for toxics. NHTSA, for example, certainly engages in risk assessment of a sort—the quantitative

---

[64] U.S. NUCLEAR REGULATORY COMM'N, REACTOR SAFETY STUDY: AN ASSESSMENT OF ACCIDENT RISKS IN U.S. COMMERCIAL NUCLEAR POWER PLANTS (1975).

[65] *See* Robert M. Bernero, *Probabilistic Risk Analyses: NRC Programs and Perspectives*, 4 RISK ANALYSIS 287, 287-88 (1984).

[66] *See id.* at 288-97; Ian B. Wall et al., *Recent Applications of PSA for Managing Nuclear Power Plant Safety*, 39 PROGRESS NUCLEAR ENERGY 367, 367-72 (2001).

[67] *See* Wall et al., *supra* note 66, at 369 n.d.

[68] *See* Thomas E. Widner & Susan M. Flack, *Dose Reconstructions for Radionuclides and Chemicals: Case Study Involving Federal Facilities at Oak Ridge, Tennessee*, *in* HUMAN AND ECOLOGICAL RISK ASSESSMENT: THEORY AND PRACTICE, *supra* note 17, at 735, 752-69.

evaluation of motor vehicle safety[69]—but lacks a clear template analogous to the exposure/dose-response framework made canonical, for toxics regulators, by the 1983 *Red Book*. Extending the "exposure" construct from toxins to radiation or certain other health hazards (such as pathogens) is straightforward; extending it further, to car crashes, industrial accidents, dangerous consumer products, or other sources of bodily injury targeted by federal regulators, is not quite so easy. But some such extrapolation is often possible. As one researcher in the area of occupational injury notes:

> Injuries are acute events associated with the transfer of hazardous levels of energy. A fatality only occurs when the energy source contacts the worker in a specific way (e.g., a tree falling on a logger's leg may cause a severe fracture, but probably not death, while the same tree striking the logger's head will usually cause death). Since the worker is only exposed to a potential fatal injury hazard for a portion of the workday, the estimation of exposure for traumatic injuries is complex . . . .[70]

Implicitly, here, "exposure" is understood as physical proximity to some machine or other workplace object that might cause death or (yet more abstractly) as the occupying of a certain spatiotemporal location in the workplace that makes some type of injury possible (falling from a high place). Some such conception of "exposure" can, in principle, provide a foundation for "exposure" assessments and "exposure"-response curves for occupational injury as well as car crashes, dangerous products, and other safety hazards.[71] Safety agencies, particularly OSHA, are just now beginning to develop risk assessment techniques along these lines.[72]

---

[69] *See infra* text accompanying notes 245-46 (finding that NHTSA regularly quantifies total deaths and injuries averted by motor vehicle safety standards).

[70] David E. Fosbroke et al., *Working Lifetime Risk of Occupational Fatal Injury*, 31 AM. J. INDUS. MED. 459, 465 (1997).

[71] *See* COVELLO & MERKHOFER, *supra* note 17, at 92 ("Physical risk agents such as mechanical force or heat can also be the subject of an exposure assessment."). Indeed, determining the "individual risk" of traumatic injury associated with "exposure" to different workplace hazards is a cutting-edge topic in occupational risk assessment. For examples or discussion of work in this area, see Fosbroke et al., *supra* note 70; Special Issue, *Occupational Injury Risk Assessment*, 4 HUM. & ECOLOGICAL RISK ASSESSMENT 1255-1441 (1998); Lauren Zeise et al., *Improving Risk Assessment: Research Opportunities in Dose Response Modeling to Improve Risk Assessment*, 8 HUM. & ECOLOGICAL RISK ASSESSMENT 1421, 1430-31 (2002).

[72] *See infra* text accompanying notes 192-99.

B. *What Is Risk? The Frequentist Answer*

What conception of "risk" is involved in risk assessment? When dose-response curves correlate an exposure amount with an "individual risk," what precisely does that risk number *mean*?

The answer: risk assessment trades upon a frequentist, rather than a Bayesian, conception of risk. Frequentism and Bayesianism are the two great traditions in the intellectual history of risk and probability.[73] Bayesianism has been hugely influential within economics and social science;[74] but the frequentist view of risk is the mainstream view within experimental science and, derivatively, within risk assessment, which has been dominated by toxicologists and other applied scientists.[75]

Scientific models are often probabilistic, and bedrock physical laws may be irreducibly probabilistic, as is now thought to be true of quantum mechanics.[76] Generically, scientists need to be able to attach a probability to a proposition asserting that some *event* will have some *attribute*. The Bayesian suggests that the probability of a proposition concerning an event, more generally the probability of any proposition, is simply someone's degree of belief in the proposition: the actual scientist's degree of belief, a "reasonable scientist's" degree of belief, an idealized observer's degree of belief, etc.[77] But scientists and, traditionally, statisticians have eschewed the suggestion, because it makes essential reference to *minds*—to beliefs—and thus has seemed too subjective for scientific purposes.[78]

Instead, following the lead of the great Austrian probabilist Richard von Mises,[79] scientists typically see probabilities as frequencies

---

[73] *See* Adler, *supra* note 20, at 1312-15.

[74] *See, e.g.*, JACK HIRSHLEIFER & JOHN G. RILEY, THE ANALYTICS OF UNCERTAINTY AND INFORMATION 9-10 (1992) (suggesting that the relevant probabilities for economic modeling are Bayesian, not frequentist).

[75] *See* COLIN HOWSON & PETER URBACH, SCIENTIFIC REASONING: THE BAYESIAN APPROACH 224 (2d ed. 1993) (noting that "classical," i.e., frequentist, statistics "now exert[s] an enormous influence on statistical workers and on scientists concerned with statistical hypotheses"). Although there is an emerging Bayesian school within toxic risk assessment, this is still quite new, and even here Bayesian probability typically has a second-order function—to quantify the analyst's beliefs about the first-order, frequentist probabilities. Frequentist probability therefore remains central. *See infra* note 328-31 and accompanying text.

[76] *See* HOWSON & URBACH, *supra* note 75, at 7-8; Adler, *supra* note 20, at 1361 & n.190.

[77] *See* Adler, *supra* note 20, at 1312 n.73 (citing sources).

[78] *See* HOWSON & URBACH, *supra* note 75, at 11-12.

[79] *See* RICHARD VON MISES, PROBABILITY, STATISTICS AND TRUTH 1-65 (2d rev. English ed., 1957) (formulating the frequentist account of probability).

within reference classes. Consider a very simple reference class: for simplicity, a class of physical objects (the class of all rocks now existing, say) rather than a class of events, which constitute a more esoteric kind of object. What is the probability that "rocks weigh 100 pounds"? Probability numbers range from zero to one, and have some other basic characteristics: if the probability that "rocks weigh 100 pounds" is p, then the probability that "rocks do not weigh 100 pounds" is $1 - p$; if the probability that "rocks weigh 50 pounds" is q, and rocks can't be both 50 pounds and 100 pounds, then the probability that "rocks weigh 50 pounds or 100 pounds" is $p + q$.[80] Notice that frequency numbers within the class of rocks satisfy just these mathematical characteristics. The frequency or proportion of 100-pound rocks can't be less than zero or greater than one; because rocks must be either 100 pounds or not 100 pounds, the proportion of 100-pound rocks plus the proportion of not-100-pound rocks is unity; and so on. Finally, note that the proportions here, hence the probabilities, are no more mind-dependent than the underlying objects and attributes: whether a rock exists, and has a given weight, does not depend on anyone's beliefs, and neither, then, does the proportion of 100-pound rocks.

Mises formalized the intuitive connection between frequencies and probabilities, by defining probability as the limit of the frequency of some attribute within an infinite, sequentially ordered reference class of events.[81] Imagine the infinite class of $\{E_1, E_2, \ldots\}$. This class has a series of segments, each of which includes the one before: $\{E_1\}$, $\{E_1, E_2\}$, $\{E_1, E_2, E_3\}$, and so on. For each such segment, one can calculate the proportion of the events with the attribute of interest A. If, as the segments become longer, the proportion of A-type events approaches p as a limit, then (on the Misean construct) p is the probability of events within this infinite class having attribute A. Mises's use of infinite reference classes accommodates the intuition that changes in observed frequencies in a finite series of experiments do not entail changes in the underlying probability, and has been somewhat controversial.[82] This controversy need not occupy us. The important point to understand is that Mises's construct, or some frequentist variant, is what underlies probability ascriptions within contemporary science and therewith quantitative risk assessment.

---

[80] *See* Adler, *supra* note 20, at 1312.

[81] *See* MISES, *supra* note 79, at 28-29; *see also* Adler, *supra* note 20, at 1314 n.79 (citing sources explicating the frequentist account of probability).

[82] *See* Adler, *supra* note 20, at 1314.

In the context of toxic risk assessment, it is pretty easy to see how this goes. The events of interest, here, are toxic exposures: the passage into a human's body of a particular dose of some toxin. For any given dose, one can imagine exposing humans to that dose, over and over again, indefinitely. This "dosing class" is an infinite class of events, specifically the infinite class of hypothetical human exposures to the particular dose. The relevant event-attribute is "causing death." Each exposure either causally contributes to, or fails to causally contribute to, the death of the person receiving that dose. For the first 1000 exposures, say, the number of subjects who die as a result of those exposures is five. For the first 10,000, say, the number is forty-eight. For the first 100,000, the number is 491. If the fractions converge, in the limit, to (say) 0.0049, then that is the frequentist, Misean probability of an event within this dosing class causing death.

To be sure, the Misean probability of death-causation within a dosing class cannot be directly observed; it is not within human capabilities to perform an infinite series of experiments. But by performing or observing a finite series of dosings, we can use statistical techniques to generate more or less precise estimates of the true frequentist probability.[83] Whatever epistemological difficulties might attend the estimation of relative frequencies, those difficulties will not be the focus of my normative critique in Parts III through V below. Rather, my focus will be conceptual. Conceptually, for the frequentist, risk is relative to reference classes.[84] Thus the ascription of risk to a particular event is arbitrary—as arbitrary as choosing a reference class, among the multiplicity of possible ones, within which to subsume a particular event. This so-called "problem of the reference class" has been much mooted by philosophers of science and probability.[85] Astonishingly, the problem is almost never mentioned within the risk assessment literature, even within scholarship that is otherwise extremely sophisticated.

Consider any particular exposure event. *P* ingests a 100-gram dose of the toxin. This particular event can be characterized in a multiplicity of ways. First, it can be characterized without a precise specification of *P*'s dose. *P* has received a dose between fifty and 150 grams. He has, at the same time, received a dose between twenty and 500 grams. He has,

---

[83] *See* HALLENBECK, *supra* note 28, at 33-62.

[84] *See* Adler, *supra* note 20, at 1345-48; Stephen R. Perry, *Risk, Harm, and Responsibility*, *in* PHILOSOPHICAL FOUNDATIONS OF TORT LAW 321, 333-35 (David G. Owen ed., 1995).

[85] *See* Adler, *supra* note 20, at 1345 n.158.

at the same time, received a dose between ninety and 200 grams. Second, the event can be characterized with a precise specification of *P*'s dose, and with further description of him. *P* has received a 100-gram dose, and *P* is forty years old. *P* has received a 100-gram dose, and *P* has a family history of cancer. *P* has received a 100-gram dose, and *P* has a family history of cancer and is a smoker. Finally, the event can be characterized with a precise specification of *P*'s dose, and with no further description of him. *P* has received a 100-gram dose. Each of these possible characterizations of the exposure event generates a different reference class: the class of all dosings between fifty and 100 grams; the class of all dosings between twenty and 500 grams; the class of all dosings between ninety and 200 grams; the class of all 100-gram dosings to forty-year-olds; the class of all 100-gram dosings to those with a family history of cancer; the class of all 100-gram dosings to smokers with a family history of cancer; and, finally, the class of all 100-gram dosings. *But the frequency with which toxic exposure causes death, within these various reference classes, may well be different.*[86] Those who receive a dose between fifty and 100 grams die as a result less frequently, one imagines, than those who receive a dose between ninety and 200 grams. Smokers who are exposed die more frequently than nonsmokers. And so on.

The reference class standardly used to calculate "individual risk" is the third type of class[87]—what might be called the *canonical dosing class*. Canonical dosing classes are composed of all exposures to hu-

---

[86] It is universally accepted by risk assessors that the frequency of death within groups of persons receiving different doses of some toxin may well be different. Otherwise, dose-response curves would be flat. And it is increasingly recognized that individual characteristics other than dose, such as age, health history, or genetic makeup, may well make a difference to toxicity. On this latter point, see, for example, NAT'L RESEARCH COUNCIL, *supra* note 17, at 200-03; Adam M. Finkel, *A Quantitative Estimate of the Variations in Human Susceptibility to Cancer and its Implications for Risk Management*, *in* LOW-DOSE EXTRAPOLATION OF CANCER RISKS: ISSUES AND PERSPECTIVES 297, 299-305 (Stephen S. Olin ed., 1995); D. Hattis & K. Barlow, *Human Interindividual Variability in Cancer Risks—Technical and Management Challenges*, 2 HUMAN & ECOLOGICAL RISK ASSESSMENT 194 (1996). An excellent recent anthology on the issue of variability in susceptibility to toxins is HUMAN VARIABILITY IN RESPONSE TO CHEMICAL EXPOSURES: MEASURES, MODELING, AND RISK ASSESSMENT (David A. Neumann & Carole A. Kimmel eds., 1998).

[87] *See* NAT'L RESEARCH COUNCIL, *supra* note 17, at 206; Finkel, *supra* note 86, at 299; Dale Hattis, *Human Interindividual Variability in Susceptibility to Toxic Effects: From Annoying Detail to a Central Determinant of Risk*, 111 TOXICOLOGY 5, 6-7 (1996). The main exception to the general practice of calculating risk relative to canonical dosing classes appears to be some areas of noncarcinogen risk assessment, where regulators have taken account of the fact that certain "sensitive subpopulations" incur a higher frequency of death or adverse effect than the general population. *See* Finkel, *supra* note 86, at 299.

mans precisely specified with respect to dose, and otherwise unspecified. Consider the exposure of *P*, a forty-year-old Caucasian smoker with a family history of cancer, to a 100 milligram (mg) dose of benzene. The canonical dosing class subsuming this exposure is the class of all events whereby human persons, of any age and with any other behavioral or genetic characteristics, are exposed to a 100 mg dose of benzene. Remember the form of the classic dose-response curve. This correlates exposures (defined as precisely as possible, in terms of a real number) with "risks," i.e., frequencies. And the curve is valid for all humans, not for a more precisely characterized subset: risk analysts typically use a single dose-response curve per toxin, rather than (say) one dose-response curve for forty-year-old smokers with a family history of cancer, another for forty-one-year-old smokers with a family history of cancer, another for forty-year-old nonsmokers, etc.

Thus, when EPA, OSHA, FDA, or a state agency undertakes an exposure assessment; determines that the "maximally exposed individual" or "highly exposed individual" or "representative individual" receives a particular dose; uses the generic dose-response curve to attach a 1 in 100,000 risk to this exposure; and concludes that the "maximally exposed" or "highly exposed" or "representative individual" is subjected to a 1 in 100,000 risk, what this means is the following: by characterizing this individual and exposure in the canonical way, abstracting from everything about the exposure except the dose received, we generate a class of exposures 1 in 100,000 of which result in death. But viewed another way (subsumed in a different, noncanonical dosing class), the "maximally exposed" or "highly exposed" or "representative individual" is subjected to a much lower risk. And viewed yet another way, she is subjected to a much higher one. Consider this analogy: whether some particular person's hair color is "unusual" depends on how we characterize the color ("bright red," "red," "within the red-brown range") and who we compare the person to. The very same adult male might come out as "unusual" if viewed as a bright red head and compared to all males, but not "unusual" if viewed as a person with hair in the red-brown range and compared to all persons in a particular ethnic group. The canonical "individual risk" number ascribed to the individual maximally exposed, or highly exposed, or receiving an average exposure from some toxin, is no more unique than a description of his hair color as "unusual" or "not unusual" generated in some standard way (by using the scheme for describing colors, and generating comparison classes, employed by the National Association of Barbers, for example). And once we un-

derstand *this* about these canonical risk ascriptions, it becomes seriously questionable why we should care about them.

II. RISK REGULATION AND "INDIVIDUAL RISK": A SURVEY OF GOVERNMENTAL PRACTICE

The preceding Part was a primer on risk assessment. In particular, it showed how the concept of "individual risk" is central to dose-response curves, a central component of risk assessment as currently structured, and it explicated the "frequentist" understanding of "individual risk." Regulatory programs relying on risk assessment as an input need not be focused on "individual risk"—dose-response curves and exposure assessments might instead be used to generate predictions of "population risk"—but the current structure of risk assessment certainly facilitates a focus on "individual risk."

This Part shows, in detail, that health and safety agencies do indeed place substantial emphasis on frequentist "individual risk." This is especially true for carcinogens, other toxic chemicals, pathogens, and radiation, but also encompasses the regulation of certain other hazards. The relevant "individual risk" is sometimes the risk to a maximally or highly exposed individual, sometimes the risk to an average individual. And these "individual risk" numbers play a range of regulatory roles. Sometimes they serve as decisional triggers: a toxin or other hazard is placed on the regulatory agenda, as it were, if the "individual risk" number is sufficiently high. Sometimes (a related idea) they serve as regulatory triggers: preexisting standards come into play if the hazard is sufficiently harmful, as measured by "individual risk" to the maximally exposed individual, the highly exposed individual, or the average individual. Finally, "individual risk" levels sometimes serve as criteria for shaping regulatory measures: rules, cleanups, or other measures should be sufficiently stringent to bring the "individual risk" (consequent upon a maximal, high-end, or average exposure) below some level. In this last role, "individual risk" might serve as a master criterion, or might instead be balanced against other criteria (for example, criteria measuring cost, technological feasibility, or "population risk").

Section A focuses on EPA, the most important health and safety agency in the United States, and the agency where attention to "individual risk" is most pervasive. Section B describes the "individual risk" based practices of other federal health and safety agencies, specifically

FDA, OSHA, NRC, and the Consumer Product Safety Commission (CPSC).[88] Section C briefly discusses the role of "population risk" in regulatory choice: even at EPA, FDA, OSHA, NRC, and CPSC, "population risk" considerations do play a role, and they certainly do at other agencies, such as NHTSA.

My ambition in this Part is *not* to provide a comprehensive overview of governmental risk assessment practices. It is rather to present the major examples of current federal[89] health and safety programs where "individual risk" has a function in regulatory choice. Unless otherwise noted, "individual risk" means the risk of death.[90] The practices here described are the target for the revisionary, normative analysis provided in Parts III to VI below.

Two final preliminary points: first, as I have already discussed, the dominant understanding of "individual risk" within the risk assessment community is frequentist, not Bayesian. This understanding pervades the regulatory practices described in this Part. "Individual risk" is, at least implicitly, understood by EPA, FDA, OSHA, NRC, and CPSC as the frequency of death relative to a canonical dosing class or some other reference class.[91] The methodologies employed by these agencies to generate the "individual risk" numbers to which they then

---

[88] In these sections, I rely heavily on three extremely helpful surveys of governmental risk assessment: Rhomberg, *supra* note 53; Alon Rosenthal et al., *Legislating Acceptable Cancer Risk from Exposure to Toxic Chemicals*, 19 ECOLOGY L.Q. 269 (1992); March Sadowitz & John D. Graham, *A Survey of Residual Cancer Risks Permitted by Health, Safety and Environmental Policy*, 6 RISK 17 (1995).

[89] State agencies also use "individual risk" tests. *See* Rosenthal et al., *supra* note 88, at 330-32; Sadowitz & Graham, *supra* note 88, at 31-33. However, it would be a very large task to identify the major examples of state regulatory reliance on "individual risk," and a task not necessary for my purposes—namely to show that "individual risk" tests are an established practice, worthy of the sustained critical attention provided in Parts III to VI below. I therefore focus on federal agencies in this Part.

[90] And unless otherwise noted, the quantitative risk levels are lifetime fatality risks—the exposed individual's risk of dying as a result of the hazard at some point during her lifetime—rather than annual or other periodic risks.

[91] For a description of EPA's approach to risk regulation, including its risk assessment methodologies, see Rhomberg, *supra* note 53, at 1080-86; OFFICE OF SCI. ADVISOR, EPA, AN EXAMINATION OF EPA RISK ASSESSMENT PRINCIPLES AND PRACTICES (2004), *available at* http://www.epa.gov/osa/ratf-final.pdf. FDA's risk assessment methodology is surveyed by D.W. Gaylor et al., *Health Risk Assessment Practices in the U.S. Food and Drug Administration*, 26 REG. TOXICOLOGY & PHARMACOLOGY 307 (1997). On OSHA and CPSC methodologies, see Rhomberg, *supra* note 53, at 1060-80. NRC's approach is generally discussed by the sources cited *infra* note 200. *See also* U.S. GEN. ACCOUNTING OFFICE, CHEMICAL RISK ASSESSMENT: SELECTED FEDERAL AGENCIES' PROCEDURES, ASSUMPTIONS, AND POLICIES (2001) (surveying EPA, FDA, OSHA, and DOT methodologies for risk assessment of toxic chemicals), *available at* http://www.gao.gov/new.items/d01810.pdf.

attach (some kind of) weight are frequentist methodologies. This point is crucial, but just because it is so general it is stated here, once and for all, and will not be belabored below.

Second, it is possible, in principle, for decisional criteria to make reference to "individual risk" yet be sensitive to population size. For example, an agency might calculate the total number of individuals incurring different levels of "individual risk" and attempt to minimize the totals within each category.[92] But, as shall emerge in the following survey of agency practice, federal agency consideration of "individual risk" typically does *not* take this form. Instead, the fact that the "individual risk" borne by some person in the exposure distribution lies above or below some stipulated level (be it 1 in 1000, 1 in 10,000, 1 in 100,000, or 1 in 1 million) functions as a deliberational or regulatory trigger, or an index of regulatory success or failure, regardless of the number of persons in the exposure distribution—regardless of the size of the population exposed to the hazard. This feature of current administrative practices—the use of "individual risk"-based decisional criteria which are insensitive to population size—is one way in which those practices are normatively misguided, as I shall argue in Part VI.

A. *"Individual Risk" and Agency Practice: The Environmental Protection Agency*

EPA is the largest health and safety agency in the federal government, and "individual risk" has a central role in this agency's decision making. Most risk assessment at EPA concerns environmental toxins, EPA's predominant regulatory target. Although EPA does also have a role in regulating pathogens and radiation, and "individual risk" considerations do come into play here,[93] this survey of EPA practice will

---

[92] *See* National Emission Standards for Hazardous Air Pollutants, 54 Fed. Reg. 38,044, 38,044-45 (EPA Sept. 14, 1989) ("In protecting public health with an ample margin of safety under section 112 [of the Clean Air Act], EPA strives to provide maximum feasible protection against risks to health from hazardous air pollutants by [inter alia] protecting the greatest number of persons possible to an individual lifetime risk level no higher than approximately 1 in 1 million . . . . "); *infra* notes 101-02 and accompanying text (discussing the role of this test under the current version of section 112).

[93] *See* National Primary Drinking Water Regulations: Long Term 2 Enhanced Surface Water Treatment Rule, 68 Fed. Reg. 47,640, 47,669-71 (proposed Aug. 11, 2003) (relying on "individual risk" in proposing treatment requirements for pathogen *Cryptosporidium* pursuant to Safe Drinking Water Act); *infra* text accompanying notes 211-20 (describing EPA's use of an "individual risk" approach in setting standards for the storage of high-level radioactive waste at Yucca Mountain).

focus on toxins. Risk assessment for toxins falls into two subcategories: cancer risk assessment and noncancer risk assessment. EPA practices in these two areas will be described separately since there are important technical differences between cancer and noncancer dose-response curves, to be explained anon.

1. Cancer Risk Assessment and "Individual Risk"

What follows are the central examples of EPA programs where the "individual risk" of cancer is determinative, wholly or partly, of the agency's regulatory choices.

a. *Air pollution (Clean Air Act section 112)*

The chief provision of the Clean Air Act governing carcinogens is section 112.[94] In its original form, as enacted in 1970, section 112 required EPA to identify "hazardous air pollutants"; once a chemical was thus listed, EPA was required to regulate emissions of the pollutant so as to "protect the public health" with an "ample margin of safety."[95] The agency was slow to implement this provision, in part because of a dilemma created by the toxicology of carcinogens. Carcinogens do not have a physiological threshold below which they lack toxicity; in other words, at any nonzero dose, the incremental fatality risk (incremental frequency of cancer within that dosing class) is greater than zero. Thus EPA faced the dilemma of either banning all emissions of carcinogenic pollutants (with huge economic costs) or allowing some emissions and therewith a nonzero probability of some deaths, in apparent violation of the "ample margin of safety" language of section 112.[96]

EPA's eventual response to this dilemma was to link section 112 standard-setting to technological feasibility: for nonthreshold toxins such as carcinogens, polluters would be required to employ the "best available technology" for limiting emissions. The D.C. Circuit, in a

---

[94] 42 U.S.C. § 7412 (2000).

[95] 42 U.S.C. § 7412(b)(1)(B) (1988), *amended by* Pub. L. No. 101-549, § 301, 104 Stat. 2399, 2531 (1990).

[96] For discussions of Clean Air Act section 112, including its history, see Bradford C. Mank, *What Comes After Technology: Using an "Exceptions Process" to Improve Residual Risk Regulation of Hazardous Air Pollutants*, 13 STAN. ENVTL. L.J. 263 (1994); Patricia Ross McCubbin, *Amending the Clean Air Act to Establish Democratic Legitimacy for the Residual Risk Program*, 22 VA. ENVTL. L.J. 1 (2003); Arnold W. Reitze, Jr. & Randy Lowell, *Control of Hazardous Air Pollution*, 28 B.C. ENVTL. AFF. L. REV. 229 (2001); Rhomberg, *supra* note 53, at 1132-34; Rosenthal et al., *supra* note 88, at 300-04, 323-27.

1987 decision involving the regulation of vinyl chloride, struck down EPA's interpretation of section 112,[97] and EPA thereupon settled on a new interpretation—one that looked to "individual risk." EPA's new test for determining a permissible emission level for carcinogenic air pollutants ran as follows: (1) the risk to the maximally exposed individual could not exceed 1 in 10,000; and (2) EPA would consider further reductions in the permissible level so as to minimize the number of individuals whose risk exceeded 1 in 1 million, but would also consider cost and feasibility considerations at this stage of the analysis.[98]

Congress quickly stepped into the fray and overhauled section 112 in 1990. The statute still applies to "hazardous air pollutants," but now provides a specific list of 188 pollutants, which EPA can revise under certain conditions. The old "ample margin of safety" test has been replaced with a more complicated structure. Specifically, EPA must set emissions limits for carcinogenic pollutants as follows: (1) the limit must be no higher than attainable using the best available technology;[99] and (2) EPA must eventually consider even lower limits if the technology-based limits "do not reduce lifetime excess cancer risks to the individual most exposed to emissions . . . to less than one in one million."[100] This explicit statutory "individual risk" provision seemingly functions as a decisional *trigger* requiring EPA to consider lower limits, not a substantive requirement that the lower limits must meet,[101] and the statute appears to contemplate that EPA will employ its hybrid, pre-1990 test (no "individual risk" above 1 in 10,000; minimize the number of individuals incurring a risk above 1 in 1 million,

---

[97] *See* Natural Res. Def. Council v. EPA, 824 F.2d 1146, 1163-66 (D.C. Cir. 1987) (en banc).

[98] National Emission Standards for Hazardous Air Pollutants, 54 Fed. Reg. 38,044, 38,044-45 (Sept. 14, 1989).

[99] More precisely, limits for "major" sources are to be based on "maximum achievable control technology," which is in turn defined differently for existing and new major sources, while limits for "area" sources are to be based on "generally available control technology." *See* Reitze & Lowell, *supra* note 96, at 247, 256-60.

[100] 42 U.S.C. § 7412(f)(2)(A) (2000). The 1 in 1 million level is also explicitly employed in a second subsection of section 112, which allows EPA to exempt a category of sources from regulation if "no source in the category . . . emits [carcinogenic] hazardous air pollutants in quantities which may cause a lifetime risk of cancer greater than one in one million to the individual in the population who is most exposed to emissions of such pollutants from the source." 42 U.S.C. § 7412(c)(9)(B)(i) (2000).

[101] *See* Rhomberg, *supra* note 53, at 1134.

taking into consideration cost and feasibility) as the substantive criterion here.[102]

b. *Water pollution (Clean Water Act and Safe Drinking Water Act)*

The Clean Water Act[103] requires that point sources of water pollution be licensed, either by EPA or by a state acting as EPA's delegate. License conditions must ensure both that the point source meets EPA's technology-based limitations and that discharges not result in violations of state water quality standards. These standards (which are also independently enforced by the states) are reviewed and approved by EPA.[104] Historically, EPA encouraged states to set standards for carcinogens so that an individual consuming a quantity of water (two liters a day) that was seen as somewhat above average but below the "high end" of the distribution, and eating a similar amount of fish (6.5 grams a day), would incur a fatality risk in the range of 1 in 10 million to 1 in 100,000. In effect, then, EPA's review of state water quality standards focused on "individual risk" to an above-average but nonmaximal individual.[105] Several years ago, EPA revised this policy to focus on the 90th percentile consumer, rather than someone whose fish and water intake is closer to average, and to adopt the 1 in 1 million "individual risk" level as its water quality goal.[106] It should be stressed that EPA's articulated "individual risk" goals under the Clean Water Act (historically a range of 1 in 100,000 to 1 in 10 million, now 1 in 1 million) are goals rather than rigid requirements, and EPA has been

---

[102] EPA has announced its intention to use the pre-1990 test. *See* OFFICE OF AIR & RADIATION, EPA, RESIDUAL RISK: REPORT TO CONGRESS, at ES-11, 124 (1999), *available at* http://www.epa.gov/ttn/oarpg/t3/reports/risk_rep.pdf; National Emissions Standards for Coke Oven Batteries, 69 Fed. Reg. 48,338, 48,339-40 (proposed Aug. 9, 2004).

[103] 33 U.S.C. §§ 1251-1387 (2000).

[104] *See* 3 LAW OF ENVIRONMENTAL PROTECTION §§ 13.31-.92 (Sheldon M. Novick et al. eds., 2004) (describing the Clean Water Act).

[105] *See* Revisions to the Methodology for Deriving Ambient Water Quality Criteria for the Protection of Human Health, 65 Fed. Reg. 66,444, 66,445-47 (Nov. 3, 2000); Rhomberg, *supra* note 53, at 1162-70; Rosenthal et al., *supra* note 88, at 309-13; Sadowitz & Graham, *supra* note 88, at 26-27.

[106] *See* OFFICE OF SCI. & TECH., EPA, METHODOLOGY FOR DERIVING AMBIENT WATER QUALITY CRITERIA FOR THE PROTECTION OF HUMAN HEALTH (2000), at 1-8, 2-6, 4-23, 4-25 (2000), *available at* http://www.epa.gov/waterscience/humanhealth/method/complete.pdf; Revisions to the Methodology for Deriving Ambient Water Quality Criteria for the Protection of Human Health, 65 Fed. Reg. at 66,445-47.

willing to approve state water quality standards that are somewhat less protective.[107]

The other key water pollution statute enforced by EPA is the Safe Drinking Water Act.[108] Under this statute, EPA directly promulgates federal water standards, which apply to virtually all public water supplies.[109] These standards are set in a two-step process: first, for each drinking water contaminant EPA sets a maximum contaminant level goal (MCLG), namely that level of the contaminant "at which no known or anticipated adverse effects on the health of persons occur and which allows an adequate margin of safety."[110] Then, EPA establishes the legally permissible level—a higher level—by taking into consideration nonhealth factors. Historically, the statutory language governing this second step focused on technological feasibility but also adverted to cost considerations: EPA was enjoined to set the legal standards as close to the MCLGs as "feasible," defined to mean "feasible with the use of the best technology, treatment techniques and other means which . . . are available (taking cost into consideration)."[111] In 1996, the statute was amended to permit a cost-benefit analysis in lieu of the feasibility analysis at the second step.[112]

EPA relies, in part, on "individual risk" in setting MCLGs. Where EPA has strong evidence that a substance is a carcinogen, it sets the MCLG at zero; where it has weaker evidence (say, some evidence of carcinogenicity in animal tests and no direct epidemiological evidence), EPA sets the MCLG so that the "individual risk" to the individual drinking an above-average amount of water is within the range of 1 in 100,000 to 1 in 1 million.[113]

---

[107] *See* OFFICE OF SCI. & TECH., *supra* note 106, at 2-6 to 2-7; Revisions to the Methodology for Deriving Ambient Water Quality Criteria for the Protection of Human Health, 65 Fed. Reg. at 66,449; Sadowitz & Graham, *supra* note 88, at 26-27.

[108] 42 U.S.C. §§ 300f to 300j-26 (2000).

[109] *See* 3 LAW OF ENVIRONMENTAL PROTECTION, *supra* note 104, at §§ 17.1-.24 (describing the Safe Drinking Water Act).

[110] 42 U.S.C. § 300g-1(b)(4)(A) (2000).

[111] 42 U.S.C. §§ 300g-1(b)(4), (5) (1994).

[112] *See* 42 U.S.C. § 300g-1(b)(6) (2000); Safe Drinking Water Act Amendments of 1996, Pub. L. No. 104-182, § 104(a)(6), 110 Stat. 1613, 1623-25 .

[113] *See* National Primary Drinking Water Regulations, 67 Fed. Reg. 19,030, 19,051 n.7 (Apr. 17, 2002); Revisions to the Methodology for Deriving Ambient Water Quality Criteria for the Protection of Human Health, 65 Fed. Reg. at 66,445-47; Rhomberg, *supra* note 53, at 1162-70; Rosenthal et al., *supra* note 88, at 309-13; Sadowitz & Graham, *supra* note 88, at 27. *But see* National Primary Drinking Water Regulations, 68 Fed. Reg. 49,548, 49,576-79 (proposed Aug. 18, 2003) (proposing non-zero MCLG for chloroform after D.C. Circuit vacated zero MCLG).

And, historically, EPA has also relied on "individual risk" in deriving enforceable federal standards from the MCLGs. Notwithstanding the statutory "feasibility" language, EPA typically set enforceable standards so that "individual risks" to above-average individuals lay in the range of 1 in 1 million to 1 in 10,000.[114] This practice has been changed, although not radically so, by the 1996 amendments. In recent Safe Drinking Water Act rulemakings, EPA has relied both on the traditional "individual risk" test just described and on cost-benefit analyses incorporating information about aggregate deaths in setting enforceable federal drinking water standards.[115]

c. *Solid waste (RCRA and CERCLA)*

The Resource Conservation and Recovery Act (RCRA)[116] and EPA's implementing regulations subject solid, hazardous waste to numerous stringent requirements, covering the generation, transportation, and (most stringently) the disposal of such wastes. These requirements are generally framed in fairly specific terms—for example, requiring disposers to employ certain technologies—and not in risk terms.[117] Rather, risk assessment comes into play under RCRA in EPA's initial determination whether to "list" particular substances as "hazardous wastes," thereby subjecting those substances to the onerous requirements just mentioned. Risk assessment also has a role to play in specifying generic toxicity characteristics such that substances with these characteristics are automatically "hazardous wastes" even without being "listed" by EPA.[118]

With respect to "listing" decisions, EPA policy has been to categorize a carcinogen-containing substance as a "hazardous waste" if unregulated disposal of the substance—specifically, disposal in ordinary municipal landfills and subsequent dissemination through groundwater—would produce an "individual risk" to a highly exposed individual (an individual in the 85th or 90th percentile of exposure) greater than 1 in 10,000. The agency exercises discretion in the range between 1 in 10,000 and 1 in 1 million. Previously "listed" substances

---

[114] *See* sources cited *supra* note 113.

[115] *See* National Primary Drinking Water Regulations, 66 Fed. Reg. 6976, 7020-23 (Jan. 22, 2001) (arsenic); National Primary Drinking Water Regulations, 65 Fed. Reg. 76,708, 76,712-15 (Dec. 7, 2000) (uranium).

[116] 42 U.S.C. §§ 6901-6992k (2000).

[117] For an overview of RCRA, see JOHN W. TEETS ET AL., RCRA: RESOURCE CONSERVATION AND RECOVERY ACT (2003).

[118] *See id.* at 21-43 (discussing the definition of hazardous waste).

can be delisted if their "individual risk" level is below 1 in 1 million. EPA considers substances to be generically toxic if they leach certain chemicals above specified concentrations, and in setting these concentrations the agency has employed an "individual risk" level of 1 in 100,000.[119]

RCRA also empowers EPA to order remedial action at active waste-disposal sites. The general agency practice, here, is apparently to perform a detailed study of remedial options for particular sites producing more than a 1 in 1 million risk to the highly exposed individual, and then to consider cost and feasibility in choosing among these options but reject remedial options resulting in more than a 1 in 10,000 risk (again to the highly exposed individual).[120] Site-specific risk assessments may be performed, too, for hazardous waste incinerators, which under RCRA require licenses. EPA apparently once took the position that a site-specific risk assessment demonstrating an "individual risk" (to a highly exposed individual) below 1 in 100,000 was a precondition for licensure.[121] More recently, EPA in a major rulemaking determined that technology-based standards for incinerators required by the Clean Air Act would generally reduce risks from dioxin emissions to individuals in the 90th percentile of exposure below 1 in

---

[119] *See* Hazardous Waste Management System, 66 Fed. Reg. 10,060, 10,088 (proposed Feb 13, 2001); Hazardous Waste Management System, 55 Fed. Reg. 11,798, 11,814-15 (Mar. 29, 1990); Rhomberg, *supra* note 53, at 1170-77; Rosenthal et al., *supra* note 88, at 314-17; Sadowitz & Graham, *supra* note 88, at 29-30; *see also* Hazardous Waste Identification Rule, 64 Fed. Reg. 63,382 , 63,440-41 (proposed Nov. 19, 1999) (proposing to use chemical-specific exemptions, determined on the basis of "individual risk," to exempt listed hazardous waste from RCRA requirements). EPA once employed an "individual risk" level of 1 in 100,000 in its listing decisions, but current policy appears to be that the agency will exercise discretion within the range of 1 in 10,000 to 1 in 1 million. *See* Rhomberg, *supra* note 53, at 1177.

[120] *See* Rhomberg, *supra* note 53, at 1170-77; Rosenthal et al., *supra* note 88, at 314-17; Sadowitz & Graham, *supra* note 88, at 29-30. EPA proposed adopting this policy as part of a comprehensive corrective action rule for RCRA, but eventually decided not to promulgate such a rule. *See* Corrective Action for Solid Waste Management Units at Hazardous Waste Management Facilities, 55 Fed. Reg. 30,798, 30,804 (proposed July 27, 1990); Corrective Action for Solid Waste Management Units at Hazardous Waste Management Facilities, 64 Fed. Reg. 54,604, 54,606 (Oct. 7, 1999). However, in deciding not to promulgate the rule, EPA affirmed that a 1996 advance notice of proposed rulemaking "should be considered the primary corrective action implementation guidance." *See id.* at 54,607. That 1996 document, in turn, affirms the "individual risk"-based policy for RCRA cleanups. Corrective Action for Releases from Solid Waste Management Units at Hazardous Waste Management Facilities, 61 Fed. Reg. 19,432, 19,446, 19,449-50 (May 1, 1996).

[121] *See* Rhomberg, *supra* note 53, at 1170-77.

10,000, and therefore that site-specific risk assessments were not presumptively required.[122]

While RCRA is forward-looking, the Superfund statute, the Comprehensive Environmental Response, Compensation, and Liability Act (CERCLA),[123] is backward-looking. CERCLA authorizes EPA to order clean-ups of inactive waste sites, with the expense of the clean-up borne by a range of private actors associated with the site.[124] Site-specific risk assessments are absolutely central to EPA's remedial decisions under CERCLA. And the goal of these risk assessments, like those under the Clean Air Act, the Clean Water Act, RCRA, and (with some recent exceptions) the Safe Drinking Water Act, is not to generate information about "population risk." Instead, CERCLA site-specific assessments focus on "individual risk" given a "reasonable maximum exposure." "Reasonable maximum exposure," as that construct is specified by EPA, is the "maximum exposure that is reasonably expected to occur at a site."[125] More specifically, "reasonable maximum exposure" is calculated by using:

> [some] values for exposure factors . . . that are mean estimates (body weight) and some parameter values that are upper bounds (e.g., exposure duration). For the concentration of the chemical at the site, the EPA guidance directs that the 95th upper confidence limit on the estimate of the mean concentration at the site or the maximum detected concentration be used, whichever is lower.[126]

---

[122] *See* National Emission Standards for Hazardous Air Pollutants: Proposed Standards for Hazardous Air Pollutants for Hazardous Waste Combustors, 69 Fed. Reg. 21,198, 21,325-31 (proposed Apr. 20, 2004); NESHAPS: Final Standards for Hazardous Air Pollutants for Hazardous Waste Combustors, 64 Fed. Reg. 52,828, 52,839-43 (Sept. 30, 1999), *vacated by* Cement Kiln Recycling Coalition v. EPA, 255 F.3d 855, 871-72 (D.C. Cir. 2001) (per curiam); Revised Standards for Hazardous Waste Combustors, 61 Fed. Reg. 17,358, 17,370-72 (proposed Apr. 19, 1996). The most recent EPA draft guidance document for site-specific incinerator risk assessments does not specify "individual risk" levels that licensed incinerators should meet. *See* EPA, 1 HUMAN HEALTH RISK ASSESSMENT PROTOCOL FOR HAZARDOUS WASTE COMBUSTION FACILITIES, Peer Review Draft at 7-9 (1998), *available at* http://www.epa.gov/epaoswer/hazwaste/combust/risk.htm.

[123] 42 U.S.C. §§ 9601-9675 (2000).

[124] For an overview of the CERCLA program, see 3 LAW OF ENVIRONMENTAL PROTECTION, *supra* note 104, at §§ 14:100-14:127.

[125] OFFICE OF EMERGENCY & REMEDIAL RESPONSE, EPA, RISK ASSESSMENT GUIDANCE FOR SUPERFUND, VOLUME 1: HUMAN HEALTH EVALUATION MANUAL (PART A), Interim Final at 6-19 (1989), *available at* http://www.epa.gov/oswer/riskassessment/ragsa.

[126] HAMILTON & VISCUSI, *supra* note 16, at 13.

EPA policy (more formalized than in the RCRA context) is to order clean-up at a CERCLA site where the individual cancer risk consequent upon "reasonable maximum exposure" exceeds 1 in 10,000; to refrain from clean-up where this risk is less than 1 in 1 million; and to exercise discretion in the range between the two levels. A clean-up, if ordered, must bring "individual risk" to the reasonably maximally exposed individual to within this same range of "individual risk" levels (1 in 10,000 to 1 in 1 million); once more, the agency has discretion, within this range, to consider factors other than "individual risk."[127]

In exercising its discretionary authority to order a clean-up and to set remedial goals under CERCLA, EPA considers factors such as cost and feasibility, but not "population risk"—at least not in any formal way, since the aggregate deaths caused by existing contamination and avoided by different remedial interventions are not quantified.[128] Some commentators have suggested that "EPA is most likely to [order a remedy within the 1 in 1 million to 1 in 10,000 range] . . . when population density suggests potentially high incidence of disease"[129] and that "EPA may be more inclined to select a more stringent remedy if a large number of people may be exposed to risks from the site."[130] But it is clear that "population risk" has at most a subordinate role in EPA's administration of the CERCLA program. Hamilton and Viscusi studied a sample of 150 sites where EPA had ordered remediation and found that, although 731 cancer cases would be averted, most of these were clustered at a few sites.[131] "[A]t the majority of sites

---

[127] On the role of "individual risk" considerations in EPA's clean-up policy under CERCLA, see HAMILTON & VISCUSI, *supra* note 16, at 25-29; Rhomberg, *supra* note 53, at 1178-86; Rosenthal et al., *supra* note 88, at 317-20; Sadowitz & Graham, *supra* note 88, at 28-29. The central features of the policy, including the "individual risk" targets, have been codified in a formal regulation. *See* 40 C.F.R. § 300.430(e)(2)(i)(A)(2) (2004).

[128] *See, e.g.*, HAMILTON & VISCUSI, *supra* note 16, at 14-15 ("[EPA does not] assess expected cancer cases averted through site cleanups. The size of the exposed population does not explicitly appear in EPA's analysis of site remediation alternatives."); Rhomberg, *supra* note 53, at 1182 ("Under EPA Superfund policy, population risks are not formally considered, so quantitative estimates of population risk rarely appear in risk assessments.").

[129] Rosenthal et al., *supra* note 88, at 320.

[130] Philip E. Karmel, *Achieving Radical Reductions in Cleanup Costs*, *in* NEW SOLUTIONS TO ENVIRONMENTAL PROBLEMS IN BUSINESS & REAL ESTATE DEALS 2003, at 371, 398 (PLI Real Estate L. & Prac. Course, Handbook Series No. 499, 2003). *See also* Rhomberg, *supra* note 53, at 1182 ("[S]ome EPA Regional Project Managers (RPMs) have unofficially acknowledged that the magnitude of the potentially exposed population sometimes informally affects remedial decisions.").

[131] *See* HAMILTON & VISCUSI, *supra* note 16, at 15, 104-07.

the expected number of cancer cases averted is less than 0.1 cases per site based on conservative risk parameter estimates."[132]

d. *Pesticides (FIFRA and FQPA)*

EPA administers two significant statutes covering pesticides. The Federal Insecticide, Fungicide, and Rodenticide Act (FIFRA)[133] is the more general statute and requires all pesticides to be "registered" (licensed) by EPA. FIFRA gives lower priority to safety than many other EPA statutes. Its central criterion is, explicitly, a balancing criterion; pesticides and pesticide uses are to be permitted absent "unreasonable risks" to human life or the environment.[134] Nonetheless, EPA in administering FIFRA has given substantial weight to "individual risk," and has employed the same range of "individual risk" levels (1 in 1 million to 1 in 10,000) as are operative in its other regulatory programs. Specifically, in evaluating pesticide licenses and license conditions, EPA considers the risks posed by pesticides to food consumers, workers, and the general public (nonworkers exposed to pesticides through pathways other than food consumption, e.g., water contamination or contact with pesticides used in the home). And it typically seeks to reduce exposures to carcinogenic pesticides, in all three categories, below an "individual risk" level of 1 in 1 million, with some tolerance for higher levels. Exposure modeling assumptions blend mid-range and high-range parameter values, so the "individual risk" here is that of an individual receiving an above-average but nonmaximal exposure.[135]

---

[132] *Id.* at 15. It appears that at only a relatively small number of the 150 sites was "no action" selected as the remedy. *See id.* at 248; *see generally id.* at app. A (describing the study methodology).

[133] 7 U.S.C. §§ 136-136y (2000).

[134] For an overview of FIFRA, see 3 LAW OF ENVIRONMENTAL PROTECTION, *supra* note 104, at §§ 18:1 to :68.

[135] *See* Rhomberg, *supra* note 53, at 1111-23; Rosenthal et al., *supra* note 88, at 304-06. This paragraph summarizes FIFRA practices as of the mid-1990s, when the Rhomberg report was written. Since then, it appears, EPA has changed its risk assessment practices under FIFRA, as a result of the Food Quality Protection Act, by focusing on aggregate risks to consumers from the combination of food, water, and residential exposures, and on cumulative risks resulting from multiple pesticides with a common mechanism of toxicity. *See* OFFICE OF PESTICIDE PROGRAMS, EPA, GENERAL PRINCIPLES FOR PERFORMING AGGREGATE EXPOSURE AND RISK ASSESSMENTS (2001), *available at* http://www.epa.gov/pesticides/trac/science/aggregate.pdf; OFFICE OF PESTICIDE PROGRAMS, EPA, GUIDANCE ON CUMULATIVE RISK ASSESSMENT OF PESTICIDE CHEMICALS THAT HAVE A COMMON MECHANISM OF TOXICITY (2002), *available at* http://www.epa.gov/pesticides/trac/science/cumulative_guidance.pdf.

The Food, Drug, and Cosmetic Act has a special section covering pesticides,[136] and these provisions are administered by EPA rather than FDA. EPA sets "tolerances," that is, maximum permissible concentrations, for pesticide residues on raw or processed food. Foods exceeding the tolerances are, legally, "adulterated" and subject to seizure.[137] Prior to 1996, EPA was required to set a zero tolerance for carcinogenic pesticide residues in processed foods, under certain conditions (by virtue of the Delaney Clause); otherwise, the statute instructed EPA to employ a balancing test in setting tolerances. In practice, for carcinogenic pesticides not covered by the zero-risk standard, EPA followed its general approach under FIFRA, namely to aim at reducing the cancer risk to an above-average food consumer below 1 in 1 million.[138]

Congress overhauled this statutory regime for pesticide tolerances with the Food Quality Protection Act of 1996,[139] which repealed the applicability of the Delaney Clause to pesticides. Tolerances for carcinogenic pesticide residues on both raw and processed foods are now to be set so that there is a "reasonable certainty that no harm will result from aggregate exposure to the pesticide."[140] Despite the switch from balancing language to language giving greater weight to safety, the House Committee report states explicitly that EPA should implement the new statutory provisions through the 1 in 1 million "individual risk" test that the agency had used (outside the Delaney context) prior to 1996.[141]

---

[136] 21 U.S.C. § 346a (2000).

[137] For a summary of both the current and the pre-1996 systems of pesticide "tolerances," see Dominic P. Madigan, Note, *Setting an Anti-Cancer Policy: Risk, Politics, and the Food Quality Protection Act of 1996*, 17 VA. ENVTL. L.J. 187, 188-205 (1998).

[138] *See* Rhomberg, *supra* note 53, at 1111-23; Rosenthal et al., *supra* note 88, at 304-06.

[139] Food Quality Protection Act of 1996, Pub. L. No. 104-170, 110 Stat. 1489 (1996) (codified as amended in scattered sections of 7 and 21 U.S.C.); *see also* Madigan, *supra* note 137, at 198-205.

[140] 21 U.S.C. § 346a(b)(2)(A)(ii) (2000).

[141] H.R. REP. NO. 104-669, pt. 2, at 41 (1996). "Individual risk" also plays a role in EPA's "Threshold of Regulation" policy for pesticide tolerances, which provides that tolerances will not be needed for pesticide uses where there are no detected residues and the incremental "individual risk" from any "theoretically possible" residues is below 1 in 1 billion. *See* THRESHOLD OF REGULATION POLICY: DECIDING WHETHER A PESTICIDE WITH A FOOD USE PATTERN NEEDS A TOLERANCE 7 (1999) (referenced in Pesticides, 64 Fed. Reg. 57,881, 57,881 (Oct. 27, 1999)), *available at* http://www.epa.gov/fedrgstr/EPA-PEST/1999/October/Day-27/6042.pdf.

e. *Toxic Substances Control Act*

The Toxic Substances Control Act (TSCA)[142] authorizes EPA to control, through rulemaking, any chemical that "present[s] an unreasonable risk of injury to health or the environment."[143] This sweeping authority is hedged by a provision that makes TSCA a back-up statute—toxics are to be regulated under TSCA only if other statutes are insufficient to meet the risk[144]—and in practice has been rarely used.[145] The office within EPA responsible for administering TSCA takes the position that carcinogens expected to cause fewer than one death per year, or less than a 1 in 1 million risk to a highly exposed individual, do not warrant regulatory intervention. Conversely, in its stated justifications for the rules that have been promulgated under TSCA, EPA has invoked both "individual risk" and "population risk."[146]

f. *Title VI*

A very different, but significant, context where "individual risk" will likely play a role in EPA practice concerns the racial impact of state and local environmental decisions. Title VI of the Civil Rights Act of 1964 authorizes federal agencies to issue regulations prohibiting recipients of federal funds, including state actors, from activities that have a disparate racial impact.[147] EPA has promulgated Title VI regulations, and recently published an important guidance document explaining how Title VI challenges to licensing decisions by state or local environmental agencies (e.g., a decision to allow a waste dump in a minority neighborhood) will be evaluated.[148] The guidance makes clear that a disparity between the minority population affected by the licensed facility and a comparison population, not merely in overall death rates, but in "individual risks" to representative or highly

---

[142] 15 U.S.C. §§ 2601-2692 (2000).

[143] *Id.* § 2605(a).

[144] *Id.* § 2608.

[145] *See* 3 LAW OF ENVIRONMENTAL PROTECTION, *supra* note 104, at §§ 16.3-.4.

[146] *See* Rhomberg, *supra* note 53, at 1123-30; Rosenthal et al., *supra* note 88, at 306-09.

[147] 42 U.S.C. § 2000d-1(2000).

[148] 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); *see also* Adler, *supra* note 20, at 1426-28 (discussing Title VI and the "individual risk" approach of the EPA guidelines).

exposed individuals, could constitute an illegal "disparate impact" for Title VI purposes.[149]

2. Risk Assessment of Noncarcinogens

Traditionally, noncancer and cancer risk assessment have been performed somewhat differently. Toxic effects other than cancer have been seen to have a physiological threshold. The difference, crudely, stems from the special causal mechanism for cancer—DNA damage to some cells, followed by proliferation of those cells—such that a dose so small as to be genotoxic to but a single cell might, in unfortunate circumstances, lead to fatal cancer for the organism.[150]

Dose-response evaluation is therefore performed differently for noncancer toxicity than for cancers.[151] Experiments and epidemiological studies are still used to produce dose-response data points pairing doses of the toxin with incremental risks (frequencies) of death or some other adverse effect, relative to background. But instead of fitting a linear function to these data points, or some other function without a threshold, the analyst instead identifies the so-called NOAEL ("no observed adverse effect level"): the "highest tested dose at which no statistically significant elevation over background in the incidence of the adverse effect was observed."[152] So-called safety factors are then applied to the NOAEL dose to produce a conservative estimate of the physiologically safe level. Typically, this means dividing

---

[149] Yet another area in which EPA has supported the use of "individual risk" tests is so-called comparative risk assessment—regulatory priority-setting at a high level. EPA's guidebook, here, suggests that "individual risk" should play a role. *See* Adler, *supra* note 20, at 1431-36.

[150] *See* HALLENBECK, *supra* note 28, at 21-25; JOSEPH V. RODRICKS, CALCULATED RISKS: UNDERSTANDING THE TOXICITY AND HUMAN HEALTH RISKS OF CHEMICALS IN OUR ENVIRONMENT 166-70 (1992); Graham, *supra* note 15, at 31-34.

[151] Good summaries of noncancer risk assessment include: CRAWFORD-BROWN, *supra* note 28, at 123-56; HALLENBECK, *supra* note 28, at 29-42, 118-26; Felter et al., *supra* note 28, at 9-16; Mark C. Gibson et al., *Comparison of Noncancer Risk Assessment Approaches for Use in Deriving Drinking Water Criteria*, 26 REG. TOXICOLOGY & PHARMACOLOGY 243, 245-54 (1997); Rhomberg, *supra* note 53, at 1105-08; Williams & Paustenbach, *supra* note 17, at 315-18, 331-35.

[152] Rhomberg, *supra* note 53, at 1105. Strictly speaking, agencies sometimes employ constructs other than a NOAEL to identify a safe level of a noncarcinogen, for example a "LOAEL" or "benchmark dose." *See* sources cited *supra* note 151. The basic idea remains the same, though: the NOAEL, LOAEL, or benchmark dose of the toxin, divided by appropriate safety factors, is supposed to be a dose that almost certainly creates zero "individual risk."

the NOAEL dose by a factor of 10, 100, or 1000.[153] The resultant dose, termed the "reference dose" (RfD)[154] by EPA, is the physiologically safe dose, to a high degree of certainty: that dose, and lower doses, do not (it can be said with great confidence) produce an incremental risk of death.

This difference between the procedure for estimating noncancer and cancer risk leads to a difference in how "individual risks" for noncarcinogens are expressed. The "individual risk" incurred by a particular person, given her exposure to a noncarcinogen, is expressed as a ratio of the dose to the RfD—not as a probability number. For example, if the RfD for the toxin is a lifetime dose of 100 grams, and the exposure assessment predicts that the maximally exposed individual will receive a lifetime dose of twenty-five grams, she will be ascribed an "individual risk" index of 1/4. If the exposure assessment instead predicts a maximal exposure of 200 grams, the "individual risk" index to the maximally exposed person is two. These nonprobabilistic indices of "individual risk" are less meaningful than the probabilistic indices employed for carcinogens. All that a nonprobabilistic index number less than one means is this: to a high degree of certainty, the incremental fatality risk incurred by that person is zero. All that a nonprobabilistic index number greater than one means is the negation, namely it cannot be stated with a high degree of confidence that the individual incurs a zero incremental fatality risk.

EPA generally conducts risk assessments for noncarcinogens, and looks to "individual risk" in making regulatory decisions, in the same statutory contexts as for carcinogens.[155] For example, under the Clean

---

[153] *See* Michael L. Dourson et al., *Evolution of Science-Based Uncertainty Factors in Noncancer Risk Assessment*, 24 REG. TOXICOLOGY & PHARMACOLOGY 108 (1996).

[154] For *inhaled* noncarcinogens, EPA uses "reference concentrations" (RfC) instead of "reference doses." *See* Rhomberg, *supra* note 53, at 1106-07.

[155] An important counterexample is air pollution. Although section 112 of the Clean Air Act generally applies to "hazardous air pollutants," including both carcinogens and noncarcinogens, the explicit statutory "individual risk" provision, *see* 42 U.S.C. § 7412(f)(2) (2000), applies only to carcinogens. This provision makes the 1 in 1 million level an "individual risk" trigger, requiring EPA to consider lowering existing technology based standards. *See supra* text accompanying notes 99-102. EPA has announced its intention to use a parallel "individual risk" trigger for noncarcinogens. *See* OFFICE OF AIR & RADIATION, *supra* note 102, at 116-18. But the substantive criteria that EPA will rely upon, once deliberation is triggered, may differ as between carcinogens and noncarcinogens. For carcinogens, EPA will rely upon the pre-1990 test, which is relatively clearly articulated and focuses on "individual risk." The substantive test it will use for noncarcinogens is less clear. *See id.* at 122-25.

The so-called "criteria" air pollutants (currently sulfur dioxide, particulate matter, carbon monoxide, ozone, nitrogen dioxide, and lead), which EPA regulates under sec-

Water Act, EPA encourages states to set water quality standards for noncarcinogenic toxins so that the index of "individual risk" to the above-average individual will be less than one.[156] Under the Safe Drinking Water Act, EPA sets MCLGs and (typically) federal standards so that the index of "individual risk" (again to the above-average individual) is less than one.[157] Under CERCLA, the agency looks to the "individual risk" from noncarcinogens to a person receiving a "reasonable maximum exposure" from the waste site under review, once more seeking to keep that index number below unity.[158] Under FIFRA, the agency takes into consideration the "individual risk" that those exposed to noncarcinogenic pesticides incur.[159]

Despite this procedural parallel between EPA's use of "individual risk" numbers in regulating both carcinogens and noncarcinogens, there is an important, substantive difference. Because EPA's cutoff for noncarcinogens is an "individual risk" index number equaling one, while its cutoff for carcinogens is an "individual risk" probability number ranging from 1 in 10,000 to 1 in 1 million, *"individual risk"-based assessment focusing on the maximally exposed individual tracks "population risk"-based assessment for noncarcinogenic toxins but not carcinogens.* Imagine that a population of 100 million is exposed to a toxin, with five million receiving maximal exposures. If the toxin is carcinogenic and EPA uses the more conservative 1 in 1 million number in regulating the toxin, seeking to bring maximal exposures to that level, it can expect a "population risk" of at least five deaths caused by the toxin after regulation. By contrast, if the toxin is noncarcinogenic and EPA uses the index number of one as its cutoff, it can expect a "population risk" of zero deaths caused by the toxin after regulation.

---

tion 109 of the Clean Air Act, are noncarcinogens. However, EPA has not employed the traditional NOAEL/safety factor methodology in setting standards under section 109 since the criteria pollutants do not have clear thresholds for all their health effects. *See* Cary Coglianese & Gary E. Marchant, *Shifting Sands: The Limits of Science in Setting Risk Standards*, 152 U. PA. L. REV. 1255, 1283-90 (2004); Rhomberg, *supra* note 53, at 1157-58.

[156] *See* Rhomberg, *supra* note 53, at 1162-70.

[157] *See id.* How this practice will be affected by the 1996 amendment to the Act remains to be seen. *See supra* text accompanying notes 112-15 (discussing the amendment).

[158] *See* Rhomberg, *supra* note 53, at 1178-86.

[159] *See id.* at 1111-23. For other discussions of EPA's risk assessment of noncarcinogens, see Donald G. Barnes & Michael Dourson, *Reference Dose (RfD): Description and Use in Health Risk Assessments*, 8 REG. TOXICOLOGY & PHARMACOLOGY 471 (1988); Carole A. Kimmel, *Quantitative Approaches to Human Risk Assessment for Noncancer Health Effects*, 11 NEUROTOXICOLOGY 189 (1990).

In other words, in some contexts involving the regulation of noncarcinogens, the advocate of risk assessment practices sensitive to "population risk" rather than "individual risk" need not be troubled by EPA's focus on "individual risk." For in these contexts the two practices are convergent. But the point should not be pressed too far. Crucially, the convergence depends on whose "individual risk" is being evaluated. Regulatory techniques for noncarcinogens that focus on the "individual risk" borne by some individual other than the maximally exposed person—the average person, say, or the person in the 90th percentile of exposure—and that seek to ensure that *this* individual's index number is below one do *not* correspond to the regulatory goal of zero "population risk." If 100 million are exposed to the toxin, five million receive maximal exposures, and EPA ensures that less exposed individuals are at zero incremental risk of fatality, then (obviously) it remains possible that maximal exposures result in a nonzero "individual risk" and hence a nonzero number of aggregate deaths.

B. *"Individual Risk" and Agency Practice: Other Agencies*

1. The Food and Drug Administration

FDA regulates the safety of foods, drugs, medical devices, and biologics, and employs risk assessment in all these areas.[160] "Individual risk" has long played a key role in FDA decision making, particularly with respect to food safety.[161]

The Food, Drug, and Cosmetic Act[162] requires "food additives" and "color additives" to be licensed by FDA, under a statutory standard

---

[160] Good reviews of FDA risk assessment practices include: D.W. Gaylor et al., *supra* note 91; Ronald W. Moch et al., *Food and Drug Administration Risk Assessment—Process and Toxicologic Pathology*, 25 TOXICOLOGIC PATHOLOGY 61, 61-63 (1997); Rhomberg, *supra* note 53, at 1043-60. An older review is Rodricks et al., *supra* note 47. For state-of-the-art reviews of food risk assessment more generally, see E. Dybing et al., *Hazard Characterisation of Chemicals in Food and Diet: Dose Response, Mechanisms and Extrapolation Issues*, 40 FOOD & CHEM. TOXICOLOGY 237 (2002); L. Edler et al., *Mathematical Modelling and Quantitative Methods*, 40 FOOD & CHEM. TOXICOLOGY 283 (2002).

[161] As already mentioned, the now widespread view that a 1 in 1 million level of "individual risk" constitutes a de minimis risk level—such that toxins or other hazards creating a lower "individual risk" can be ignored by regulators—derives from an FDA decision in the 1970s. *See supra* text accompanying note 47.

[162] 21 U.S.C. §§ 301-397 (2000).

that gives high priority to safety.[163] For food additives and color additives that may be toxic but are not carcinogenic, FDA employs the standard NOAEL/safety factor method that (as we have seen) is used by EPA for noncarcinogens. Experiments are undertaken or epidemiological data is checked, and the dose of the noncarcinogen that produces zero incremental frequency of death in the group of subjects receiving that dose, relative to background, is determined. That dose is then divided by a "safety factor" (typically 100 at FDA) to derive the "safe" level of the noncarcinogenic toxin. This is, in effect, the dose that is highly likely to involve zero "individual risk." FDA then combines that number with information about food consumption patterns to determine the maximum permissible concentration of the noncarcinogenic food or color additive in the foods to which it is added. Specifically, FDA seeks to ensure that the 90th percentile food consumer will not ingest more than the "safe" dose of the noncarcinogenic additive.[164] Because FDA focuses on the 90th percentile of the consumption distribution, not on the maximally exposed individual, this method cannot be justified as assuring zero "population risk." At least in principle, depending on consumption patterns above the 90th percentile and the size of the population, a concentration of some additive that very likely creates zero "individual risk" for the 90th percentile consumer might also be likely to cause more than zero deaths.

FDA apparently employs an analogous method in setting permissible levels of pathogenic microorganisms, including bacteria, molds, yeasts, and viruses, in foods. The degree of microbial contamination that will not sicken a high-end consumer is determined, and then divided by a safety factor.[165]

Carcinogenic food and color additives are governed by the infamous Delaney Clause: a flat ban on any additive if it is "found to induce cancer when ingested by man or animal, or if it is found, after tests which are appropriate for the evaluation of the safety of . . . additives, to induce cancer in man or animal."[166] This Clause bars FDA

---

[163] *See id.* §§ 348, 379e (2000); *see also* 1 JAMES T. O'REILLY, FOOD AND DRUG ADMINISTRATION chs. 11-12 (2d ed. 1995) (describing FDA regulation of food and color additives).

[164] FDA's risk assessment practices for noncarcinogenic food and color additives are discussed in Gaylor et al., *supra* note 91, at 307-09; Moch et al., *supra* note 160, at 61-62; and Rhomberg, *supra* note 53, at 1055-60.

[165] *See* Gaylor et al., *supra* note 91, at 308-09.

[166] 21 U.S.C. § 348(c)(3)(A) (2000) (food additives); *see id.* § 379e(b)(5)(B) (parallel provision for color additives).

from licensing the use of an additive, at any level, if the additive itself has been found to be carcinogenic. But FDA has long interpreted the Clause, without judicial disagreement, to include an important exception: an additive which has not itself been found to be carcinogenic, but has some nonfunctional chemical constituent that *is* carcinogenic, is not governed by the Delaney Clause.[167] Instead, FDA's view is that such carcinogen-containing additives are subject to the background "safety" requirement generally applicable to additives. FDA licenses these additives by setting a maximum permissible concentration of the additive sufficiently low to ensure that the 90th-percentile food consumer incurs an "individual risk" of cancer death no greater than 1 in 1 million.[168]

The Delaney Clause applies not merely to food and color additives, but also to animal drugs and feeds.[169] The concern, of course, is that human carcinogens, if fed to animals, might accumulate in meats and other foods derived from animals. But the Food, Drug and Cosmetic Act includes an explicit exception to the proscription on carcinogenic animal drugs and feeds:

> [The Delaney Clause] shall not apply with respect to the use of a substance as an ingredient of feed for animals which are raised for food production, if [FDA] finds . . . that no residue of the additive will be found . . . in any edible portion of such animal after slaughter or in any food yielded by or derived from this living animal . . . .[170]

FDA's reading of this statutory "no residue" exception to the Delaney Clause has been creative. A literal reading of the exception would make it a virtual nullity, given the miniscule concentrations that are detectable with modern techniques used to analyze food contaminants. Instead, FDA has (until recently) interpreted "no residue" as "safe"; safety, in turn, has been equated with the traditional 1 in 1 million level of "individual risk." Carcinogenic residues of animal drugs

---

[167] *See* Gaylor et al., *supra* note 91, at 309-10; Moch et al., *supra* note 160, at 62; Rhomberg, *supra* note 53, at 1051; Rodricks et al., *supra* note 47, at 535-36.

[168] *See* Gaylor et al., *supra* note 91, at 307-10; Rhomberg, *supra* note 53, at 1052-60.

[169] Carcinogenic animal drugs are directly proscribed by a separate Delaney Clause. *See* 21 U.S.C. § 360b(d)(1)(I) (2000). Animal feeds that fall within the broad statutory definition of food additive, *see id.* § 321(s), are covered by the food additive Delaney Clause.

[170] 21 U.S.C. § 348(c)(3)(A) (2000) (food additives); *see also id.* § 360b(d)(1)(I) (parallel exception for animal drugs); *id.* § 379e(b)(5)(B) (parallel exception for color additives).

and feeds have been permitted at a concentration imposing a 1 in 1 million risk on the 90th-percentile consumer.[171]

This "individual risk"-based interpretation of the Delaney Clause with respect to animal drugs and feeds was revised by FDA in 2002, in response to a determination by the Department of Justice that the "no residue" provision could not be read to countenance carcinogens that actually produced detected residues.[172] FDA's new approach is therefore more roundabout, but still incorporates an "individual risk" test: the 1 in 1 million level is used to determine how sensitive analytical methods for detecting carcinogen residues must be, and the "no residue" requirement is satisfied if none is detected above the "limit of detection" of an approved analytical method.[173]

FDA has general jurisdiction over food safety, subsuming not merely food and color additives and animal drugs, but any toxin in food—for example, environmental contaminants such as PCBs and aflatoxins. The burden of action rests on FDA—foods are not licensed by FDA, but rather are subject to seizure, regulation, and penalties if dangerous—and the underlying statutory standard is more permissive than for additives.[174] FDA decision making here is sensitive to the costs of eliminating food toxins (for example the hedonic and nutritional costs of banning food products containing the toxins, as reflected in the market price of the foods), in contrast to additive regulation, where safety is the sole acknowledged regulatory consideration.[175] Concentrations of food and color additives posing a non-de minimis risk of death or injury are flatly proscribed by the agency; that is not necessarily true for other food toxins. Still, "individual risk" plays a role in FDA regulation of such toxins (carcinogens and noncarcinogens alike) by serving to define the de minimis level. Foods containing toxins below *that* level—a 1 in 1 million "individual risk" for carcinogens, a zero "individual risk" for noncarcinogens—are seen

---

[171] *See* Sponsored Compounds in Food-Producing Animals, 52 Fed. Reg. 49,572, 49,572-74 (Dec. 31, 1987); Gaylor et al., *supra* note 91, at 311-12; Graham, *supra* note 15, at 34-35; Rhomberg, *supra* note 53, at 1048-60; Rodricks et al., *supra* note 47, at 533-35; Sadowitz & Graham, *supra* note 88, at 21-22.

[172] *See* Revision of the Definition of the Term "No Residue" in the New Animal Drug Regulations, 67 Fed. Reg. 78,172, 78,172-73 (Dec. 23, 2002).

[173] *See* 21 C.F.R. §§ 500.80-.92 (2004); 67 Fed. Reg. at 78,172-73.

[174] *See* 21 U.S.C. §§ 331, 332-34, 342(a) (2000). On FDA regulation of food safety, see 1 O'REILLY, *supra* note 163, at ch. 9.

[175] *See, e.g.*, Polychlorinated Biphenyls (PCBs) in Fish and Shellfish, 49 Fed. Reg. 21,514 (May 22, 1984); PETER BARTON HUTT & RICHARD A. MERRILL, FOOD AND DRUG LAW: CASES AND MATERIALS 904-07 (2d ed. 1991) (discussing FDA regulation of aflatoxins).

as safe by FDA and therefore permissible. In effect, this de minimis, "individual risk" level serves as a trigger for regulatory analysis rather than (as with food additives) for an automatic proscription.[176]

Finally, "individual risk" is the linchpin of FDA's so-called "threshold of regulation" policy, which has been in place for nearly a decade.[177] Chemicals that are contained in food packaging materials or other food-contact articles and that migrate into food are technically "food additives," subject to FDA licensure. But the licensing process is expensive; the "threshold of regulation" policy therefore provides that chemicals in food-contact materials not currently known or suspected to be carcinogenic, which leave residues in food at a concentration below 0.5 parts per billion (ppb), are exempt from licensure. FDA arrived at the 0.5 ppb level by looking at the universe of known toxins. Noncarcinogenic chemicals, it emerges, are very unlikely to produce toxicity at that level, while for carcinogens the 0.5 ppb level corresponds to the tried-and-true "individual risk" level of 1 in 1 million. As the agency explained:

> [We] used potency data on a large number of known carcinogens to estimate the likely risk that could be expected if an unstudied compound were later found to be a carcinogen. . . .
>
> . . . FDA further restricted its analysis to the 477 animal carcinogens that were the subject of oral feeding studies. . . .
>
> Based on the range of potencies exhibited by these 477 animal carcinogens, FDA has determined that most known carcinogens pose less

---

[176] *See* Gaylor et al., *supra* note 91, at 308 (noting that "FDA has used the no observed adverse effect level . . . from bioassay data for noncancer effects as the starting point for safety assessment of chemicals," both for food additives and contaminants); Rhomberg, *supra* note 53, at 1054 (discussing the de minimis level for carcinogens).

[177] *See* Food Additives; Threshold of Regulation for Substances Used in Food-Contact Articles, 60 Fed. Reg. 36,582 (July 17, 1995); 58 Fed. Reg. 52,719 (proposed Oct. 12, 1993). For recent discussions, see Robert Kroes et al., *Threshold of Toxicological Concern for Chemical Substances Present in the Diet: A Practical Tool for Assessing the Need for Toxicity Testing*, 38 FOOD & CHEM. TOXICOLOGY 255 (2000); Robert Kroes & Gunhild Kozianowski, *Threshold of Toxicological Concern (TTC) in Food Safety Assessment*, 127 TOXICOLOGY LETTERS 43 (2002). In 1997, subsequent to the adoption of the Threshold of Regulation policy, Congress amended the Food, Drug, and Cosmetic Act to permit food-contact materials that are food additives to be marketed without a license. Manufacturers can instead submit a premarket notification to FDA. However, the agency retains discretion to require the licensure process, and the Threshold of Regulation policy remains in place as an alternative to both premarket notification and licensure. *See* 1 O'REILLY, *supra* note 163, at § 11:9.50 (2004 Fall Cumulative Supp.); Food Additives: Food Contact Substance Notification System, 67 Fed. Reg. 35,724 (May 21, 2002).

> than one in a million lifetime risk if present in the daily diet at 0.5 ppb.[178]

2. The Occupational Safety and Health Administration

OSHA regulates toxic workplace chemicals as well as other occupational hazards. For toxins, the two crucial provisions are sections 6(b)(5) and 3(8) of the Occupational Safety and Health Act.[179] Section 6(b)(5) is specific to toxins and provides:

> The Secretary, in promulgating standards dealing with toxic materials or harmful physical agents . . . shall set the standard which most adequately assures, to the extent feasible, on the basis of the best available evidence, that no employee will suffer material impairment of health or functional capacity even if such employee has regular exposure to the hazard . . . for the period of his working life.[180]

This language, taken alone, might be read as authorizing OSHA to regulate any chemical that is a toxin—that is toxic at some dose—and to require reductions in workplace exposures to the toxin to the lowest feasible level.[181] But section 3(8) of the Act, a generic provision describing the "standards" that OSHA is empowered to issue, defines an "occupational safety and health standard" as "a standard which requires conditions . . . reasonably necessary or appropriate to provide safe or healthful employment and places of employment."[182] In *Industrial Union Department v. American Petroleum Institute*,[183] a plurality of the Supreme Court read section 3(8) as creating a "significant risk" threshold for OSHA regulation of toxins as well as other safety hazards. In effect, the Court determined that OSHA was statutorily required to recognize a de minimis level of risk; workplace toxins creating risks below that level could not be regulated.

> 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

---

[178] Food Additives; Threshold of Regulation for Substances Used in Food-Contact Articles, 58 Fed. Reg. at 52,722.

[179] 29 U.S.C. §§ 651-678 (2000).

[180] Occupational Safety and Health Act of 1970 (OSH Act), Pub. L. No. 91-596, § 6(b)(5), 84 Stat. 1590, 1594 (codified at 29 U.S.C. § 655(b)(5) (2000)).

[181] Indeed, this was OSHA's approach to regulating carcinogens prior to *Industrial Union*. *See* Graham, *supra* note 15, at 36; Rhomberg, *supra* note 53, at 1062.

[182] OSH Act, § 3(8), 84 Stat. at 1591 (codified at 29 U.S.C. § 652(8) (2000)).

[183] 448 U.S. 607, 642 (1980) (plurality opinion).

> 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." . . .
>
> 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 risks are present and can be eliminated or lessened by a change in practices.[184]

The Court then went on to suggest, famously, that the "significant risk" requirement might be understood in terms of "individual risk." A 1 in 1000 "individual risk" was clearly significant, the Court said; a 1 in 1 billion risk was not.[185]

Industrial Union's linkage between significant risk and "individual risk" might have been rejected as dictum. An agency more self-confident than OSHA might have read the case as mandating a de minimis threshold but permitting that to be specified in "population risk" terms—as some number of premature deaths that would be caused by the workplace toxin absent OSHA intervention.[186] Instead, OSHA's practice in regulating workplace carcinogens has been to follow the letter of *Industrial Union*: the agency determines whether the existing concentration of a workplace carcinogen is a "significant risk," warranting OSHA intervention, by determining whether a worker exposed to that concentration for his entire working lifetime (forty-five years of exposure, five days a week, eight hours a day) would incur an "individual risk" of premature death that exceeds, or at least is not too far below, 1 in 1000.[187] Interestingly, OSHA's approach is more eclectic once it has determined that the status quo level of a workplace carcinogen poses a "significant risk." Considerations of "population risk," "individual risk," and economic and technical "fea-

---

[184] *Id.*

[185] *Id.* at 655.

[186] *See* Cross et al., *supra* note 19, at 73-75 (explaining that "significant risk" might be specified in "population risk" or "individual risk" terms).

[187] *See* Occupational Exposure to Methylene Chloride, 62 Fed. Reg. 1494, 1560 (Jan. 10, 1997) ("[A] risk of 1/1000 . . . is clearly significant. It represents the uppermost end of a million-fold range suggested by the Court, somewhere below which the boundary of acceptable versus unacceptable risk must fall." (citing *Industrial Union*, 448 U.S. at 655)); Rhomberg, *supra* note 53, at 1060-70; Sadowitz & Graham, *supra* note 88, at 24-25.

sibility" all seem to bear on the agency's decision as to what the permissible level of the carcinogen should be.[188]

What about noncarcinogens? There is an inherent tension between the standard NOAEL/safety factor method for regulating noncarcinogens—a method that seeks to ensure an "individual risk" level of zero—and *Industrial Union*'s statement that a workplace toxin might pose an "individual risk" above zero but still be too insignificant to trigger OSHA's regulatory authority. This tension came to light in 1992, when the Eleventh Circuit struck down a rulemaking in which OSHA had used the NOAEL/safety factor method to set permissible exposure limits for a variety of noncarcinogens.[189] Since this decision, OSHA has issued only one new exposure limit for a noncarcinogen.[190]

"Individual risk" has, historically, not played a role in OSHA's regulation of workplace safety hazards as opposed to health hazards, such as falls from heights or dangerous machines. To be sure, the "significant risk" threshold created by section 3(8) of the OSHA act applies to all OSHA regulations, whether targeted at illness or injury. But "significant risk" for safety hazards has in the past been understood in "population risk" terms. OSHA officials recently stated that: "Traditionally, OSHA has based its significant risk determination for

---

[188] *See* Occupational Exposure to Methylene Chloride, 62 Fed. Reg. at 1563-69 (Jan. 10, 1997); Occupational Exposure to 1,3-Butadiene, 61 Fed. Reg. 56,746, 56,794-97 (Nov. 4, 1996). For a compact statement of OSHA's approach to setting permissible limits once a toxin has been determined to pose a "significant risk," see 61 Fed. Reg. at 56,791.

[189] AFL-CIO v. OSHA, 965 F.2d 962, 975-80 (11th Cir. 1992).

[190] The one new limit is the lead exposure limit for the construction industry, issued under explicit statutory mandate. *See* Lead Exposure in Construction, 58 Fed. Reg. 26,590, 26,590-91 (May 4, 1993). Subsequent to the Eleventh Circuit decision, OSHA relied upon the NOAEL/safety factor method in proposing exposure limits for glycol ethers. *See* Occupational Exposure to 2-Methoxyethanol, 2-Ethoxyethanol, and Their Acetates (Glycol Ethers), 58 Fed. Reg. 15,526, 15,576-79 (proposed Mar. 23, 1993). That proposal was eventually withdrawn. *See* 68 Fed. Reg. 75,475 (Dec. 31, 2003). OSHA has also issued exposure limits for carcinogens with noncarcinogenic effects. *See, e.g.*, Occupational Exposure to Methylene Chloride, 62 Fed. Reg. at 1494; Occupational Exposure to 1,3-Butadiene, 61 Fed. Reg. at 56,746; Occupational Exposure to Asbestos, 59 Fed. Reg. 40,964 (Aug. 10, 1994); Occupational Exposure to Cadmium, 57 Fed. Reg. 42,102 (Sep. 14, 1992); Occupational Exposure to 4,4-Methylenedianiline (MDA), 57 Fed. Reg. 35,630 (Aug. 10, 1992). In all these cases, OSHA employed its standard approach for carcinogens, namely, satisfying the "significant risk" test by establishing a quantitative "individual risk" of cancer in the vicinity of 1 in 1000. Interestingly, in some of these rulemakings OSHA also quantified the "individual risk" of noncancer toxicity. *See, e.g.*, Occupational Exposure to Asbestos, 59 Fed. Reg. at 40,966 (asbestosis); Occupational Exposure to Cadmium, 57 Fed. Reg. at 42,207-09 (kidney dysfunction).

safety standards on estimates of the annual numbers of injuries and/or fatalities associated with exposure to an occupational injury hazard and the number of injuries and fatalities likely to be prevented with a new standard in place."[191]

This may be changing. As already explained, the "individual risk" of workplace injury is a perfectly coherent concept.[192] Indeed, there is an emerging subliterature, within risk assessment, that seeks to define and measure the "individual risk" of occupational injury.[193] The National Institute for Occupational Safety and Health (NIOSH), the federal government's research institute on occupational health and safety, has encouraged this research.[194] And OSHA, in the huge ergonomics rulemaking a few years ago, relied in substantial part on the "individual risk" construct in arguing that ergonomics hazards were a "significant risk" and therefore fell within OSHA's regulatory jurisdiction.[195] OSHA defined the worker's risk of musculoskeletal disorder as "the probability that a worker will experience at least one work-related musculoskeletal disorder during his or her working lifetime (45 years),"[196] and used data on workplace injuries to determine this risk for different occupations.[197]

OSHA practices in regulating pathogens, covered by the agency's broad statutory authorization to issue workplace health and safety

---

[191] J.F. Martonik et al., *Injury Risk Assessment for Occupational Safety Standards*, 4 HUM. & ECOLOGICAL RISK ASSESSMENT 1259, 1259 (1998).

[192] Just as one can construct a "dosing class" of persons exposed to a particular dose of some toxin, and define the "individual risk" of fatal cancer for that dose as the frequency with which persons in the dosing class develop fatal cancer as a result of the toxin, so too one can construct an "exposure class" of persons exposed in some sense to a workplace hazard—say, all persons who use ladders, or a type of machine—and define the "individual risk" of fatal injury as the frequency with which persons in the class experience a fatal injury as a result of the hazard. *See supra* text accompanying notes 69-72.

[193] *See supra* text accompanying notes 70-72.

[194] *See* Linda Rosenstock & Nancy Stout, *Occupational Injury Risk Assessment: Perspective and Introduction*, 4 HUM. & ECOLOGICAL RISK ASSESSMENT 1255, 1256 (1998).

[195] *See* Ergonomics Program, 65 Fed. Reg. 68,262, 68,538-68, 68,754-55 (Nov. 14, 2000); Ergonomics Program, 64 Fed. Reg. 65,768, 65,926-85 (proposed Nov. 23, 1999). The ergonomics rule was overturned pursuant to the Congressional Review Act. *See* Ergonomics Program, 66 Fed. Reg. 20,403 (Apr. 23, 2001).

[196] 65 Fed. Reg. at 68,556; 64 Fed. Reg. at 65,936.

[197] As OSHA explained in the proposed rulemaking: "Among the 58 industry groups for which BLS provided estimates of the number of MSDs reported in 1996, the median lifetime risk of experiencing at least one LWD [lost workday] MSD is 255 per 1,000 workers, and for only 8 of these industry groups is the estimated lifetime risk below 100 cases per 1,000 workers." 64 Fed. Reg. at 65,980. There is a similar analysis in the final rulemaking. *See* 65 Fed. Reg. at 68,755.

standards, should also be mentioned. In 1991, OSHA promulgated a bloodborne pathogen standard; in 1997, it proposed a standard for tuberculosis, subsequently withdrawn.[198] In both cases the agency quantified "individual risk" and sought to show that its traditional, 1 in 1000 threshold for regulatory significance was satisfied.[199]

3. The Nuclear Regulatory Commission

Reactor safety risk assessment has been central to NRC activities since the Three Mile Island accident. The Commission has employed risk assessment to evaluate the safety of individual reactors, the generic safety of different plant designs, the advisability of new regulations for existing plants, and in other contexts.[200] The Commission formalized its commitment to risk assessment in a 1995 policy statement.[201]

Reactor safety risk assessment need not be keyed to the probability of death and illness. Instead, a probabilistic assessment (be it of an individual plant, a class of plants, a regulation, an inspection protocol, or something else) might focus on the probability of an accident—either the probability of damage to the reactor core, or the probability of what NRC calls a "large early release" of radioactivity into the environment. For example, a risk assessment of an individual reactor might conclude that the annual probability of that reactor experiencing core damage is 2 in 1 million, without tracing or ascribing probabilities to the causal paths leading from core damage to the irradiation of persons near the plant and to their deaths. Or, the annual probabilities that different amounts of radioactivity could be released as a result of core damage plus containment failure might be quantified, without deriving "population risks" or "individual risks" from

---

[198] *See* Occupational Exposure to Tuberculosis, 68 Fed. Reg. 75,768 (Dec. 31, 2003).

[199] *See* Occupational Exposure to Tuberculosis, 62 Fed. Reg. 54,160, 54,211-14 (proposed Oct. 17, 1997); Occupational Exposure to Bloodborne Pathogens, 56 Fed. Reg. 64,004, 64,034-38 (Dec. 6, 1991).

[200] *See* Use of Probabilistic Risk Assessment Methods in Nuclear Regulatory Activities, 60 Fed. Reg. 42,622, 42,622-23 (Aug. 16, 1995); Wall et al., *supra* note 66, at 368-72. Earlier reviews of NRC risk assessment include: Bernero, *supra* note 65; Miller B. Spangler, *A Summary Perspective on NRC's Implicit and Explicit Use of de Minimis Risk Concepts in Regulating for Radiological Protection in the Nuclear Fuel Cycle*, *in* DE MINIMIS RISK 111 (Chris Whipple ed., 1987). An excellent technical overview of reactor risk assessment methodology is Jon C. Helton & Roger J. Breeding, *Calculation of Reactor Accident Safety Goals*, 39 RELIABILITY ENG'G & SYS. SAFETY 129 (1993).

[201] Use of Probabilistic Risk Assessment Methods, 60 Fed. Reg. at 42,628-29.

these release scenarios.[202] However, the Commission in 1986 adopted a series of "safety goals" as the ultimate benchmarks for reactor risk assessment. The goals are an interesting hybrid of "individual risk" and "population risk":

> The Commission has established two qualitative safety goals which are supported by two quantitative objectives. These two supporting objectives are based on the principle that nuclear risks should not be a significant addition to other societal risks. . . .
>
> • The *qualitative safety goals* are as follows:
>
> — Individual members of the public should be provided a level of protection from the consequences of nuclear power plant operation such that individuals bear no significant additional risk to life and health.
>
> — Societal risks to life and health from nuclear power plant operation should be comparable to or less than the risks of generating electricity by viable competing technologies and should not be a significant addition to other societal risks.
>
> • The following *quantitative objectives* are to be used in determining achievement of the above safety goals:
>
> — [Individual risk:] The risk to an average individual in the vicinity of a nuclear power plant of prompt fatalities that might result from reactor accidents should not exceed one-tenth of one percent (0.1 percent) of the sum of prompt fatality risks resulting from other accidents to which members of the U.S. population are generally exposed.
>
> — [Population risk:] The risk to the population in the area near a nuclear power plant of cancer fatalities that might result from nuclear power plant operation should not exceed one-tenth of one percent

---

[202] As one commentator explains:

> [Reactor risk assessment] scopes are categorized as Levels 1, 2, and 3. Level 1 is the systems analysis with an end-state of Core Damage Frequency (CDF), Level 2 analyses the physical processes of the accident including the containment response with an end-state of Large Early Release Frequency (LERF) and the quantities and compositions of the radioactive materials released to the atmosphere (source term), and Level 3 analyses the transport of radioactive material through the environment and estimates the public health and economic consequences of the accident.

Wall et al., *supra* note 66, at 369 n.d; *cf.* Bernero, *supra* note 65, at 287 (distinguishing between these three levels of reactor accident risk assessment).

> (0.1 percent) of the sum of cancer fatality risks resulting from all other causes.[203]

The Commission's position has been that the safety goals are to be used in evaluating NRC regulations, not in measuring the safety of particular plants. Instead, risk assessments for particular plants are to focus on the probability of core damage and "large early release" of radiation; and the general probability benchmarks regarding core damage and "large early release" are, in turn, to be shaped by the safety goals.[204] But there have been recent suggestions by the Commission and staff that the safety goals might be used on a plant specific basis.[205] In any event, the goals make "individual risk" one of several foundational, regulatory criteria for nuclear plants in the United States. Plant designs and procedures are ultimately to be evaluated by considering, inter alia, the "individual risk" of immediate death following a reactor accident incurred by the average individual living "in the vicinity" of a plant.[206] The Commission has clarified that this means "the average individual biologically (in terms of age and other risk factors) and locationally who resides within a mile from the plant

---

[203] Safety Goals for the Operation of Nuclear Power Plants, 51 Fed. Reg. 30,028, 30,028-29 (Aug. 21, 1986).

[204] *See id.* at 30,031-32; Use of Probabilistic Risk Assessment Methods, 60 Fed. Reg. at 42,627-28.

[205] NRC staff recommended various changes to the safety goal policy statement, including a change to the effect that "[t]he Commission approves use of the . . . safety goals . . . in the regulatory decisionmaking process *on both plant-specific and generic bases*." It appears that the Commission tentatively approved this particular change. *See* NRC, SECY-01-0009, Modified Reactor Safety Goal Policy Statement, at pt. V. (Jan. 22, 2001) (emphasis added), *available at* http://www.nrc.gov/reading-rm/doc-collections/commission/secys/2001/secy2001-0009/2001-0009scy.html; *see also* Mark A. Caruso et al., *An Approach for Using Risk Assessment in Risk-Informed Decisions on Plant-Specific Changes to the Licensing Basis*, 63 RELIABILITY ENG'G & SYS. SAFETY 231, 233 (1999) (suggesting, in a staff-authored article, that safety goals may be used on a plant-specific basis). However, the Commission eventually decided not to amend the safety goal policy statement. *See* NRC, Commission Voting Record, Modified Reactor Safety Goal Policy Statement (Apr. 16, 2001), *available at* http://www.nrc.gov/reading-rm/doc-collections/commission/cvr/2001/2001-0009vtr.html.

[206] Indeed, the putative "societal risk" goal, namely that "[t]he risk to the population in the area near a nuclear power plant of cancer fatalities that might result from nuclear power plant operation should not exceed one-tenth of one percent (0.1 percent) of the sum of cancer fatality risks resulting from all other causes," is itself a hybrid of "individual risk" and "population risk," *see* Safety Goals for the Operation of Nuclear Power Plants, 51 Fed. Reg. at 30,028-29, since it divides the number of cancer deaths in the nearby population resulting from an accident by the total number of cancer deaths in that population, rather than using the numerator alone as the relevant measure. *See* Vicki M. Bier, *The U.S. Nuclear Regulatory Commission Safety Goal Policy: A Critical Review*, 8 RISK ANALYSIS 563, 564-65 (1988).

site boundary."[207] The quantitative "individual risk" goal articulated by NRC—0.1 percent of the average "individual risk" of prompt fatality from other causes—translates into an annual risk of 1 in 2 million.[208] In short, reactors are (inter alia) "safe" when average individuals living very near them incur no more than a 1 in 2 million annual chance of dying immediately as a result of a reactor accident. This has been NRC's stated safety policy since a few years after Three Mile Island.

Reactor safety risk assessments focus on the probability of a reactor accident. Considerations of "individual risk" also influence federal regulation of other aspects of the safety of nuclear power generation, for example by shaping NRC criteria for maximum permissible doses of radiation received by the population or plant workers as a result of ordinary plant operation,[209] or by structuring the regulation of nuclear wastes.[210] A salient recent example involves the proposed Yucca Mountain repository for high-level nuclear waste, which is to be built and operated by the Department of Energy, and licensed by NRC, pursuant to safety standards issued by EPA.[211] Congress has enacted a number of relevant statutes, including the Energy Policy Act of 1992,[212] which suggested, without mandating, that the design of the Yucca Mountain site be focused on "individual risk" rather than "population risk." The Act stated that EPA shall promulgate "standards for protection of the public from releases from radioactive materials stored or disposed of in the repository at the Yucca Mountain site,"[213] but instructed EPA to consult first with the National Academy of Sciences on various questions, including: "[W]hether a health-based standard based upon doses to individual members of the public

---

[207] Safety Goals for the Operation of Nuclear Power Plants, 51 Fed. Reg. at 30,030.

[208] *See* Helton & Breeding, *supra* note 200, at 129; *cf.* Spangler, *supra* note 200, at 139 (determining that prompt fatality safety goal translates into an annual "individual risk" of 4×10^{-7}, i.e., 1 in 2.5 million).

[209] *See* Sadowitz & Graham, *supra* note 88, at 18-20; *see also* COMM. ON TECH. BASES FOR YUCCA MOUNTAIN STANDARDS, *supra* note 63, at 41-50 (summarizing different regulatory limitations on radiation exposure and resultant "individual risk").

[210] *See* B. John Garrick, *The Use of Risk Assessment to Evaluate Waste Disposal Facilities in the United States of America*, 40 SAFETY SCI. 135, 135-46 (2002).

[211] The Yucca Mountain program is discussed in *id.* at 139; General Guidelines for the Recommendation of Sites for Nuclear Waste Repositories, 66 Fed. Reg. 57,298, 57,299-57,311 (Nov. 14, 2001); David P. Ross, Note, *Yucca Mountain and Reversing the Irreversible: The Need for Monitored Retrievable Storage in a Permanent Repository*, 25 VT. L. REV. 815 (2001).

[212] Energy Policy Act of 1992, Pub. L. No. 102-486, § 801, 106 Stat. 2776, 2921-23 (1992) (codified as amended in scattered sections of 15, 16, 40, and 42 U.S.C).

[213] *Id.* § 801(a)(1).

[i.e., an individual-risk based standard] . . . will provide a reasonable standard for protection of the health and safety of the general public."[214] The National Academy of Sciences thereupon produced a report, which concluded that the safety of Yucca Mountain should indeed be judged in terms of "individual risk" rather than the total deaths caused by releases from facility—specifically, "individual risk" to maximally exposed persons—and suggested that the safe level of "individual risk" might be set at 1 in 2000 or lower.[215]

EPA responded by promulgating a standard for Yucca Mountain that is framed in terms of radiation dose, not risk: persons maximally exposed to releases from the waste site must not receive more than fifteen millirem of radiation per year.[216] But EPA justified this dose-based standard by adverting to the de minimis "individual risk" imposed on the maximally exposed individual by a radiation dose at or below the level of fifteen millirem. As the agency explained:

> [W]e have based the proposed dose-based standard upon the risk of developing a fatal cancer as a result of that level of exposure based upon a linear, non-threshold, dose-response relationship. . . . Dose and [individual] risk are closely related; one can be converted to the other simply by using the appropriate factor.[217]

A fifteen millirem annual dose, assuming a seventy-year lifetime for the maximally exposed person and the dose-response curve that EPA invoked in the Yucca Mountain rulemaking, translates into an "individual risk" of 6 in 10,000.[218] Both NRC and the Department of

---

[214] *Id.* § 801(a)(2)(A).

[215] *See* COMM. ON TECH. BASES FOR YUCCA MOUNTAIN STANDARDS, *supra* note 63, at 1-14, 33-65. Strictly speaking, the report recommended that the "individual risk" standard focus on risk to members of the small "critical group" of persons at highest risk from the repository, not the single maximally exposed individual. *See id.* at 51-54.

[216] *See* Public Health and Environmental Radiation Protection Standards for Yucca Mountain, NV, 66 Fed. Reg. 32,074, 32,085-95 (June 13, 2001); Ross, *supra* note 211, at 830-36. The relevant exposure is that of the "reasonably maximally exposed" individual.

[217] Environmental Radiation Protection Standards for Yucca Mountain, Nevada, 64 Fed. Reg. 46,976, 46,984 (proposed Aug. 27, 1999). For similar statements in the final rulemaking, see 66 Fed. Reg. at 32,086.

[218] This 6 in 10,000 lifetime "individual risk" is calculated by multiplying the fifteen millirem dose times seventy years times the slope of EPA's linear dose-response model (5.75/10000 cancers per rem). *See* Public Health and Environmental Radiation Protection Standards for Yucca Mountain, NV, 66 Fed. Reg. at 32,086; Environmental Radiation Protection Standards for Yucca Mountain, Nevada, 64 Fed. Reg. at 46,979.

Energy have amended their own Yucca Mountain regulations to conform to EPA's 15 millirem/year dose limit.[219]

In the summer of 2004, the D.C. Circuit vacated EPA's Yucca Mountain regulation insofar as it mandated compliance with the fifteen millirem limit and other requirements for only 10,000 years—a compliance period that the court found to be inconsistent with the Energy Policy Act of 1992.[220] How EPA will respond to this ruling remains to be seen.

4. The Consumer Product Safety Commission

The Consumer Product Safety Commission (CPSC) administers various statutes, including the Federal Hazardous Substances Act,[221] which requires that "hazardous substances" intended for household use bear a specified label.[222] A "hazardous substance" is defined as "[a]ny substance or mixture of substances which (i) is toxic . . . if such substances or mixture of substances may cause substantial personal injury or substantial illness during or as a proximate result of any customary or reasonably foreseeable handling or use."[223] A "hazardous substance" that lacks the required label is "misbranded" and cannot be introduced into interstate commerce.[224]

In 1992, CPSC issued lengthy guidelines that explain when a substance is "toxic" and that employ an "individual risk" test in defining the subset of "toxic" substances that are "hazardous" for statutory purposes.[225] Specifically, a product containing a carcinogen is "hazardous" if a consumer using the product incurs an "individual risk" ex-

---

[219] *See* Disposal of High-Level Radioactive Wastes in a Proposed Geologic Repository at Yucca Mountain, NV, 66 Fed. Reg. 55,732, 55,733 (NRC Nov. 2, 2001); General Guidelines for the Recommendation of Sites for Nuclear Waste Repositories, 66 Fed. Reg. 57,298, 57,310-11 (Dep't of Energy Nov. 14, 2001) .

[220] *See* Nuclear Energy Inst., Inc. v. EPA, 373 F.3d 1251, 1266-73 (D.C. Cir. 2004) *reh'g en banc denied*, 2004 U.S. App. LEXIS 18782, at *1 (Sept. 1, 2004).

[221] 15 U.S.C. §§ 1261-1278 (2000). The Act is summarized in 2 DONALD W. STEVER, LAW OF CHEMICAL REGULATION AND HAZARDOUS WASTE §§ 4:1 to :18 (2003).

[222] *See* 15 U.S.C. §§ 1261(p), 1263 (2000).

[223] *Id.* § 1261(f)(1)(A).

[224] *Id.* § 1263(a).

[225] Labeling Requirements for Art Materials Presenting Chronic Hazards, 57 Fed. Reg. 46,626 (Oct. 9, 1992). On toxic risk assessment at CPSC, see Michael A. Babich, *Risk Assessment of Low-Level Chemical Exposures from Consumer Products Under the U.S. Consumer Product Safety Commission Chronic Hazard Guidelines*, 106 ENVTL. HEALTH PERSP. (SUPP. 1) 387 (1998); Rhomberg, *supra* note 53, at 1070-80.

ceeding 1 in 1 million.[226] The portion of the guidelines covering exposure assessment stipulates that exposure means "anticipated exposure from normal lifetime use," and that "[i]n most cases the best estimate of exposure (average exposure) is acceptable."[227] In short, the relevant "individual risk" level is that of the average, rather than maximal consumer. The "hazardous" cutoff for noncarcinogenic toxins is defined using the standard NOAEL/safety factor method. It is that product concentration resulting in a lifetime exposure to the average consumer which equals the "no observed effect level," divided by a safety factor of 10 or 100.[228]

CPSC has rulemaking authority under the Act. The agency can supplement or vary the statutorily required labeling and, under certain conditions, ban a "hazardous substance."[229] "Population risk" considerations clearly play a role in these rulemakings.[230] For example, in justifying a rule that required retail containers of charcoal to bear a label warning of the carbon monoxide risk from burning charcoal in confined spaces, CPSC explained that "there are approximately 28 deaths and 300 CO-related . . . injuries associated with the use of charcoal each year."[231] From these numbers CPSC inferred both an "individual risk" of dying from charcoal-related CO poisoning[232] and an aggregate monetized cost of death and injury.[233]

C. *Beyond "Individual Risk": "Population Risk" in Agency Practice*

The prior two Sections described a wide range of administrative practices whereby agency choice depends, wholly or partly, on the level of frequentist "individual risk" incurred by some person exposed to a hazard, identified by her place in the exposure distribution—for example, the person with an average exposure, or the person with a

---

[226] Labeling Requirements for Art Materials Presenting Chronic Hazards, 57 Fed. Reg. at 46,656.

[227] *Id.* at 46,647, 46,656.

[228] *Id.* at 46,656.

[229] *See* 15 U.S.C. §§ 1261(q), 1262(b)-(c) (2000).

[230] *See* Rhomberg, *supra* note 53, at 1080.

[231] Requirements for Labeling of Retail Containers of Charcoal, 61 Fed. Reg. 19,818, 19,827 (May 3, 1996) (internal citation omitted).

[232] *See id.* at 19,827 ("[T]he estimated 160 million bags of charcoal briquets sold in 1995 were associated with approximately one death for every 5.7 million charcoal briquet bags . . . .").

[233] *See id.* ("Assuming a statistical value of life of $5 million [and a cost of injury of $10,000], these [300] injuries and [28] deaths cost society about $143 million annually.").

high-end (90th percentile, say) exposure, or the maximally exposed person. I shall argue below that these practices are normatively misguided and should be changed. This normative analysis presupposes that the practices *could* be different—"ought implies can"—and the best way to demonstrate *this* is by showing that health and safety agencies sometimes employ metrics other than "individual risk" to a particular person in evaluating hazards. In particular, federal health and safety agencies sometimes evaluate hazards by quantifying "population risk": the total deaths, illnesses, or injuries caused by the hazards and abated by intervention. What follows is not comprehensive, or even a collection of the major examples, but is rather meant to underscore that the "individual risk"-based methodologies now dominant at EPA and also employed at FDA, OSHA, NRC, and CPSC are by no means inevitable.

Even where carcinogens and other toxins are concerned, regulatory analysis sometimes focuses (at least in part) on "population risk." Various examples were interspersed in my discussion above. Since the 1996 amendments to the Safe Drinking Water Act, EPA has relied both on "individual risk" to the above-average individual, and on cost-benefit analyses incorporating information about aggregate deaths, in setting enforceable drinking water standards.[234] In regulating food contaminants that impose an "individual risk" exceeding what FDA regards as the de minimis level, the agency has taken into consideration "population risk."[235] Similarly for OSHA: once that agency has determined that a carcinogen currently found in the workplace crosses the threshold of regulability because it imposes a sufficiently high (roughly 1 in 1000) "individual risk" on the maximally exposed worker, the agency sets the permissible level of the carcinogen by attending to "population risk" as well as "individual risk" and economic and technical feasibility.[236] CPSC has pointed to the aggregate deaths and injuries caused by toxins in justifying rulemaking under the Fed-

---

[234] *See supra* text accompanying note 115. Another area where EPA has given some weight to "population risk" considerations is the regulation of criteria pollutants. *See, e.g.*, National Ambient Air Quality Standards for Ozone, 62 Fed. Reg. 38,856, 38,863-68 (EPA July 18, 1997); *see also* Coglianese & Marchant, *supra* note 155, at 1290-1323 (arguing that EPA lacks a coherent approach to regulating the criteria pollutants).

[235] *See supra* text accompanying notes 174-76; Polychlorinated Biphenyls (PCBs) in Fish and Shellfish; Reduction of Tolerances, 49 Fed. Reg. 21,514, 21,519 (May 22, 1984).

[236] *See supra* text accompanying notes 187-88.

eral Hazardous Substances Act,[237] as has EPA in its rulemakings under the Toxic Substances Control Act.[238]

A yet more compelling example, perhaps: a recent international survey of carcinogen risk assessment practice found that "European [agencies] have established the estimation of the likely incidence of cancer in the human population as the goal for risk assessment."[239] This is a "population risk" measure—or at least one much closer to "population risk" than the sorts of "individual risk" measures generally employed at EPA.[240]

My survey of agency practice, in Sections A and B above, suggests that "individual risk"-based tests are widely used to evaluate health hazards, such as toxic chemicals, radiation, and pathogens, but are much less often used for safety hazards, where the threat is injury rather than illness. One counterexample is OSHA's recent reliance on "individual risk" levels to assess workplace safety hazards, as in the ergonomics rulemaking.[241] But this is a special case. Agencies administering federal safety programs usually rely on "population risk," not "individual risk," as a measure of regulatory need and success. For example, CPSC has jurisdiction under the Consumer Products Safety Act[242] and other statutes[243] to regulate dangerous consumer products. In recent rulemakings for safety hazards (such as bath seats, which risk infant drownings; bunk beds, which risk entrapment of small children; "dive sticks," which risk impalement; and multi-purpose lighters, which risk fires), the agency has provided historical data about total annual or periodic deaths and injuries, and has used this information,

---

[237] *See supra* text accompanying notes 229-32.

[238] *See supra* text accompanying note 146.

[239] Robert J. Moolenaar, *Carcinogen Risk Assessment: International Comparison*, 20 REG. TOXICOLOGY & PHARMACOLOGY 302, 308 (1994).

[240] It is not clear from the survey whether the Europeans' focus on the "incidence of cancer in the human population" means a focus on the total number of cancer deaths caused by a hazard, or instead on a fraction equaling the total number of cancer deaths from a hazard divided by some measure of the size of the whole population in the country (total population, total cancer deaths from any source). Still, this kind of fraction is more directly related to "population risk" (total cancer deaths caused by a hazard) than the "individual risk" numbers relied upon by EPA, since the fraction equals "population risk" divided by a number that is roughly constant (a measure of the size of the whole population).

[241] *See supra* text accompanying notes 195-97.

[242] 15 U.S.C. §§ 2051-2085 (2001).

[243] *See id.* §§ 1261-78 (Federal Hazardous Substances Act); *id.* §§ 1471-76 (Poison Prevention Packaging Act).

together with a monetary value of lifesaving and injury-reduction, to estimate the monetized benefits of regulation.[244]

The National Highway Traffic Safety Administration (NHTSA) regulates motor vehicle safety. Unlike CPSC, NHTSA has until very recently declined to *monetize* lifesaving and injury-avoidance.[245] But like CPSC (and OSHA), NHTSA regularly *quantifies* aggregate lifesaving and injury-avoidance. In numerous rulemakings conducted over the last decade—concerning power windows; lap/shoulder belts in rear seats; tire performance; tire pressure monitoring systems; airbags; child restraint anchorage systems; reflectors for truck tractors; rear impact guards for trailers; door locks; and head impact protection—the agency has predicted both the total lives and injuries avoided by the rule, and the monetary costs.[246] "Population risk," for NHTSA, has traditionally served as the input to a fuzzy cost-benefit analysis where

---

[244] *See* Bath Seats, 68 Fed. Reg. 74,878, 74,885-86 (Dec. 29, 2003); Dive Sticks, 66 Fed. Reg. 13,645, 13,646-49 (Mar. 7, 2001); Safety Standard for Bunk Beds, 64 Fed. Reg. 71,888, 71,888-90, 71,897-98 (Dec. 22, 1999); Safety Standard for Multi-Purpose Lighters, 64 Fed. Reg. 71,854, 71,855, 71,864-67 (Dec. 22, 1999).

[245] Email from Ro Malik, AEI-Brookings Joint Center, to author (Nov. 18, 2003). NHTSA practice, here, may be changing. In a number of rulemaking analyses published in 2004, the agency did monetize life-saving or injury-avoidance. *See* Federal Motor Vehicle Safety Standards; Head Restraints, 69 Fed. Reg. 74,848, 74,878-79 (Dec. 14, 2004); Federal Motor Vehicle Safety Standards; Occupant Crash Protection, 69 Fed. Reg. 70,904, 70,911-12 & n.10 (Dec. 8, 2004); Federal Motor Vehicle Safety Standards; Tire Pressure Monitoring Systems, 69 Fed. Reg. 55,896, 55,916-18 (proposed Sept. 16, 2004); Federal Motor Vehicle Safety Standards; Side Impact Protection, 69 Fed. Reg. 27,990, 28,013-14 (proposed May 17, 2004).

[246] *See* Federal Motor Vehicle Safety Standards; Power-Operated Window, Partition, and Roof-Panel Systems, 69 Fed. Reg. 55,517, 55,528-29 (Sept. 15, 2004); Federal Motor Vehicle Safety Standards; Occupant Crash Protection, 68 Fed. Reg. 46,546, 46,550 (proposed Aug. 6, 2003) (lap/shoulder belts); Federal Motor Vehicle Safety Standards; Tires, 68 Fed. Reg. 38,116, 38,118 (June 26, 2003) (tire performance); Federal Motor Vehicle Safety Standards; Tire Pressure Monitoring Systems, 67 Fed. Reg. 38,704, 38,739 (June 5, 2002); Federal Motor Vehicle Safety Standards; Occupant Crash Protection, 63 Fed. Reg. 49,958, 49,983 (proposed Sept. 18, 1998) (air bags); Federal Motor Vehicle Safety Standards; Child Restraint Anchorage Systems, 64 Fed. Reg. 10,786, 10,796-97 (Mar. 5, 1999); Federal Motor Vehicle Safety Standards; Lamps, Reflective Devices and Associated Equipment, 61 Fed. Reg. 41,355, 41,359 (Aug. 8, 1996) (reflectors); Federal Motor Vehicle Safety Standards; Rear Impact Guards, 61 Fed. Reg. 2004, 2029 (Jan. 24, 1996); Federal Motor Vehicle Safety Standards; Door Locks and Door Retention Components, 69 Fed. Reg. 75,020, 75,028 (proposed Dec. 15, 2004); Federal Motor Vehicle Safety Standards; Door Locks and Door Retention Components, 60 Fed. Reg. 50,124, 50,132-33 (Sep. 28, 1995); Federal Motor Vehicle Safety Standards; Head Impact Protection, 60 Fed. Reg. 43,031, 43,047-48 (Aug. 18, 1995). The very recent rulemakings in which NHTSA conducts a straightforward monetized cost-benefit analysis also calculate "population risk." *See* sources cited *supra* note 245.

cost per life saved is calculated but no explicit cut-off monetary value for a human life is stated—by contrast with CPSC's more straightforward balancing. But in both cases "population risk" statistics figure centrally in the agencies' analyses.

III. FREQUENTIST RISK AND WELFARIST CONSEQUENTIALISM

"Individual risk" in the frequentist sense should be irrelevant to regulatory practice. The frequentist conceptualizes the "individual risk" of death as the incremental frequency of death, above background frequency, in a "dosing class" (for toxins) or an analogous reference class (for nontoxins). The practices described in Part II, where agencies give weight in some way to frequentist "individual risk," are normatively misguided. This is true across a wide range of moral theories. In this Part, I seek to demonstrate the normative irrelevance of frequentist "individual risk" within welfarist consequentialism—the normative framework now standardly adopted by economists, policy analysts, and many others. Part V examines competing normative frameworks.

So as to avoid extra complexity, and because most of the current examples of regulatory attention to "individual risk" involve toxic chemicals, my normative analysis here and below will focus specifically on toxins.[247] But the results of that analysis are, I believe, fully generalizable to nontoxic hazards. Governmental regulation of workplace accidents, motor vehicle crashes, radiation leaks, mechanically flawed consumer products, and other such targets of regulatory concern—like governmental regulation of toxins—should not be keyed to frequentist "individual risk."

A. *Welfarist Consequentialism: Some Clarifications*

Welfarist consequentialism is the moral view underlying welfare economics[248] and, derivatively, law and economics.[249] It is the view that

---

[247] Similarly, the analysis will focus on the "individual risk" of death but generalizes, quite naturally, to the "individual risk" of nonfatal illness or injury.

[248] More precisely, welfare economics standardly presupposes welfare consequentialism plus a preference-satisfaction view of welfare. *See, e.g.*, ROBIN BOADWAY & NEIL BRUCE, WELFARE ECONOMICS 1-3 (1984).

[249] *See* LOUIS KAPLOW & STEVEN SHAVELL, FAIRNESS VERSUS WELFARE 15-81 (2002). Admittedly, Kaplow and Shavell's claim that normative evaluation of legal rules reduces to a welfare-consequentialist evaluation has proven controversial, even among law-and-economists. *See, e.g.*, Lewis A. Kornhauser, *Preference, Well-Being, and Morality in*

drives cost-benefit analysis,[250] which over the last two decades has become the dominant methodology for normative evaluation employed in the federal government.[251] Welfarist consequentialism, more generally, undergirds much normative talk in our polity, by legislators, presidents, bureaucrats, and citizens—witness the pervasive assumption that policy choices are properly tied to public "goals," "aims," or "purposes" (all consequentialist constructs), and that these goals are ultimately related to human welfare in some way.[252]

The concept of an "outcome" or, more precisely, a "possible world"—a maximally specified outcome[253]—is central to consequentialist thinking. Any consequentialist moral view has two parts: (1) an "ex post" part, namely a criterion for ranking outcomes, a criterion that determines whether an outcome is "better" or "worse" than another; and (2) an "ex ante" part, namely a rule for determining what choice an actor should make in any choice situation, given the range of choices available to him and the possible outcomes of each choice. Further, for a theory to count as consequentialist, both the "ex post" and "ex ante" parts of the theory must satisfy certain formal constraints. For example, the "ex post" criterion for ranking outcomes must be "evaluator neutral." In other words, the ranking must not vary depending on the evaluator's perspective: O_1 will be better or worse than O_2 from each of our viewpoints, not better from your viewpoint but worse from mine.[254]

Social Decisions, 32 J. LEGAL STUD. 303 (2003) (criticizing Kaplow and Shavell's claim that fairness considerations properly influence social choice via preferences for fairness).

[250] More precisely, cost-benefit analysis tracks overall well-being: a welfare-based and consequentialist construct that might be incorporated in a variety of moral views. *See* Adler & Posner, *Implementing Cost-Benefit Analysis when Preferences Are Distorted*, *supra* note 36, at 1108-16; Adler & Posner, *Rethinking Cost-Benefit Analysis*, *supra* note 36, at 194-225.

[251] *See* Adler, *supra* note 20, at 1389-92.

[252] *See, e.g.*, Kadrmas v. Dickinson Pub. Sch., 487 U.S. 450, 457-58 (1988) (holding that the equal protection principle requires that statutory classification be "rationally related to a legitimate governmental purpose"); ANTHONY OGUS, REGULATION: LEGAL FORM AND ECONOMIC THEORY 29-54 (1994) (surveying a range of standard justifications for regulation, all connected to well-being).

[253] *See* JOHN DIVERS, POSSIBLE WORLDS (2002); MICHAEL J. LOUX, METAPHYSICS: A CONTEMPORARY INTRODUCTION 176-214 (2d ed. 2002).

[254] *See* Adler, *supra* note 20, at 1316-21. It might be objected that consequentialism permits evaluator-relative outcome rankings, or moral rules that are coextensive with evaluator-relative rankings. But I would not define "consequentialism" so inclusively. *See infra* note 370.

Welfarist consequentialism is a particular kind of consequentialist view. It offers a particular answer to the ex post question: what determines the moral goodness of outcomes? The welfarist answers: "Welfare, and only welfare."[255] "Welfare" here means the well-being of humans or any other creatures that are welfare subjects, most plausibly certain nonhuman mammals. The moral goodness of outcomes hinges on their welfare goodness, their goodness for welfare subjects. This defining aspect of welfarist consequentialism is often framed as a so-called "supervenience"[256] requirement: one outcome cannot be ranked as better or worse than another outcome, by the ex post criterion, unless the two outcomes differ with respect to at least one person's welfare.[257] Welfare-identical outcomes are seen, by the welfarist consequentialist, to be morally identical.

Just like the Russian dolls that open to reveal smaller ones, the categories here divide into subcategories that themselves divide. Moral views include both consequentialist and nonconsequentialist views; consequentialist views, both welfarist and nonwelfarist variants; welfarist consequentialist views, both utilitarian and nonutilitarian versions.[258] Many assume that welfarist consequentialism just equates with utilitarianism. But this is untrue. Given any set of outcomes, there will typically be different (indeed a multitude of) ranking schemes, all of which satisfy the "evaluator neutrality" requirement definitive of consequentialism and the welfare-supervenience requirement definitive of welfarism.[259] Outcomes can be ranked in light of overall well-being: that is utilitarianism.[260] Or, they can be ranked in a strictly

---

[255] *See* Andrew Moore & Roger Crisp, *Welfarism in Moral Theory*, 74 AUSTRALASIAN J. PHIL. 598 (1996); *cf.* Richard J. Arneson, *Welfare Should Be the Currency of Justice*, 30 CANADIAN J. PHIL. 497, 497-98 (2000) (presenting a welfarist account that incorporates considerations of individual responsibility).

[256] On the general notion of "supervenience," see JAEGWON KIM, SUPERVENIENCE AND MIND 53-78 (1993).

[257] *See* BOADWAY & BRUCE, *supra* note 248, at 143-44; Walter Bossert & John A. Weymark, *Utility in Social Choice*, *in* 2 Handbook of Utility Theory (Salvador Barberà et al. eds. 2004), *available at* http://www.econ.ubc.ca/dp9623.pdf; AMARTYA SEN, CHOICE, WELFARE AND MEASUREMENT 18-19 (1982).

[258] For an excellent overview of these distinctions, albeit with somewhat different terminology, see SHELLY KAGAN, NORMATIVE ETHICS 25-105 (1998).

[259] This point is reflected, within welfare economics, in the recognition that there are multiple "social welfare orderings" consistent with welfarism. *See* BOADWAY & BRUCE, *supra* note 248, at 137-69. It is reflected, within moral philosophy, in the recognition that welfare consequentialism encompasses both utilitarianism and also views more sensitive to the distribution of welfare. *See* KAGAN, *supra* note 258, at 29-54.

[260] *See generally* GEOFFREY SCARRE, UTILITARIANISM (1996); UTILITARIANISM AND BEYOND (Amartya Sen & Bernard Williams eds., 1982).

egalitarian way, such that an outcome with a more equal distribution is better than an outcome with a less equal distribution, even if the first outcome is Pareto-inferior to the second. Or they can be ranked in a "prioritarian" way, such that (1) Pareto-superior outcomes are always ranked better than their inferiors; but (2) in the ranking of Pareto-noncomparable outcomes, special weight is given to those whose welfare levels are lower.[261] Or, they can be ranked in a way that balances utilitarian and egalitarian considerations: this balancing might admit some Pareto-inferior moves, but would also (unlike strict egalitarianism) admit some Pareto-superior moves that purchase large improvements in overall welfare with small decrements in equality.[262] Further possibilities—many of them!—exist. The analysis that follows will be agnostic as between the different versions of welfarist consequentialism. It will not be necessary to take a position in the ongoing, philosophical debates between utilitarians, egalitarians, and prioritarians.

My analysis will, however, take sides in a different debate. It will rely upon a *substantive* rather than *preferentialist* account of well-being.[263] Preferentialists say that P is better off in some outcome O_1, as compared to some other outcome O_2, if and only if P prefers O_1 to O_2 (in some sense). Substantivists offer a list of welfare "values" or "goods." For example, John Finnis claims that these goods are: life itself, knowledge, play, aesthetic experience, sociability, practical reasonableness, and religion.[264] Martha Nussbaum's list includes: life, bodily health, bodily integrity, the use of the "senses, imagination and thought," emotions, practical reason, affiliation, interaction with other species, play, and control over one's environment.[265] Derek Parfit, describing (without endorsing) the substantive view of welfare, writes that "[t]he good things might include moral goodness, rational activity, the development of one's abilities, having children and being a

---

[261] The choice between egalitarianism and prioritarianism is currently a "hot topic" in moral philosophy. *See* 19 ECON. & PHIL. 1-134 (2003).

[262] *See* Adler, *supra* note 35, at 310.

[263] Several of my works provide critical surveys of theories of welfare, most recently Adler, *supra* note 20, at 1303-16. *See also* Adler, *supra* note 35, at 262-67; Adler & Posner, *Rethinking Cost-Benefit Analysis*, *supra* note 36, at 197-204. Hedonism, often contrasted with both substantive (or "objective") and preferentialist views of welfare, might alternatively be seen as a narrow kind of substantive view—one with only two goods, namely pleasure and the avoidance of pain.

[264] JOHN FINNIS, NATURAL LAW AND NATURAL RIGHTS 85-90 (1980).

[265] MARTHA C. NUSSBAUM, WOMEN AND HUMAN DEVELOPMENT: THE CAPABILITIES APPROACH 78-80 (2000).

good parent, knowledge, and the awareness of true beauty."[266] James Griffin lists accomplishment, autonomy, understanding, enjoyment, and deep personal relations.[267] The substantivist claims that *P* is better off in some outcome O_1, as compared to some other outcome O_2, only if O_1 is better for *P* than O_2 with respect to one or more genuine welfare values.

I have elsewhere described at length the grounds in favor of a substantive view of human well-being. To summarize quickly: preferentialists typically end up invoking fully informed and otherwise idealized preferences, rather than actual preferences, as the basis for welfare. Consider the taste for sadism: satisfying this actual preference, intuitively, does not benefit its holder. The construct of idealized preferences licenses that sort of intuition, by providing critical purchase on the welfare subject's actual preferences, and yet still links up with the subject's wants and desires. If someone wouldn't prefer something, even with full information and after full deliberation, then it doesn't improve her welfare.[268] Substantive views go yet further in the direction of idealization. Roughly, substantive goods are those features of lives that *all* idealized agents would prefer. Something is substantively good for a subject if it would be ideally preferred, not merely by the subject herself, but by all persons with sufficient information, deliberation, etc., considering a life with that feature and one without it.[269] In other words, substantive views define welfare in terms of convergent idealized preferences. An important argument in favor of this convergence requirement involves interpersonal comparisons.[270] *Any* welfare-consequentialist moral view (be it a utilitarian view, an egalitarian view, or some other) will need to compare the welfare levels or changes in the welfare levels of different persons.[271] For example, if one person actually or ideally prefers outcome O_1, and another person actually or ideally prefers O_2, what would license the conclusion that O_1 or O_2 is better for their aggregate welfare? Substantive welfare goods can license that sort of conclusion; idealized preferences, alone, cannot.

---

[266] DEREK PARFIT, REASONS AND PERSONS 499 (reprinted with corrections 1987).

[267] JAMES GRIFFIN, VALUE JUDGEMENT: IMPROVING OUR ETHICAL BELIEFS 29-30 (1996).

[268] *See* Adler, *supra* note 20, at 1304-06.

[269] *See id.* at 1353-54.

[270] *See* Adler, *supra* note 35, at 289-302.

[271] *See id.* at 296-97; Adler & Posner, *Rethinking Cost-Benefit Analysis*, *supra* note 36, at 204-09.

B. *The Ex Post Question: Does "Individual Risk" Degrade Outcomes?*

One welfare-consequentialist defense of "individual risk" takes the ex post perspective.[272] The claim, here, is that an outcome or possible world in which some person *P* is subjected to a high frequentist risk of death is worse for him, *ceteris paribus*, than an outcome or possible world in which he is not subjected to that risk. This is not a preposterous idea. Imagine that, while *P* is sleeping, an intruder steps to his bedside, pulls out a revolver with a bullet in one of the six chambers, spins the chambers, pulls the trigger—with no resultant firing—and leaves. At least some people respond to this sort of hypothetical case with the intuition that *P* is made worse off by the game of Russian roulette, even though he is unaware of the game at the time it is played and even if he never learns of it.[273] By analogy, perhaps, the "maximally exposed" or "highly exposed" or representative individual who incurs a 1 in 10,000 or 1 in 1 million risk of death by virtue of his toxic exposure (as calculated relative to the canonical dosing class or some other reference class) suffers a welfare reduction just by virtue of this frequentist risk, even if the toxin does not cause his death or otherwise affect him.

Intuitions are important, but may mislead: for example, in the Russian roulette case, we may have a strong intuition that *P* has been harmed, but this intuition may be grounded in the violation of *P*'s property rights, or in the intruder's contempt for *P*,[274] not the risk im-

---

[272] I argue for the "ex post" irrelevance of "individual risk" at length in Adler, *supra* note 20, at 1340-69. The analysis here reaches the same conclusion, albeit in a different way.

[273] *See* Claire Finkelstein, *Is Risk a Harm?*, 151 U. PA. L. REV. 963, 970-71 (2003); John C.P. Goldberg & Benjamin C. Zipursky, *Unrealized Torts*, 88 VA. L. REV. 1625, 1651 (2002); *cf.* Peter Railton, *Locke, Stock, and Peril: Natural Property Rights, Pollution, and Risk*, *in* TO BREATHE FREELY: RISK, CONSENT, AND AIR 89, 101-07 (Mary Gibson ed., 1985) (suggesting that risk imposition, as in Russian roulette cases, is wrongful); Stephen A. Fogdall, Risks as Harms 1 (unpublished paper, on file with author) (arguing that "it is plausible" that risks are harms). *But see* ARIEL PORAT & ALEX STEIN, TORT LIABILITY UNDER UNCERTAINTY 101-15 (2001) (arguing that risk is not harm); Heidi M. Hurd, *The Deontology of Negligence*, 76 B.U. L. REV. 249, 263-64 (1996) (same); Perry, *supra* note 84, at 330-39 (same); John Oberdiek, The Morality of Risking: On the Normative Foundations of Risk Regulation 30-50 (2003) (unpublished Ph.D. dissertation, University of Pennsylvania) (on file with the University of Pennsylvania Library) (same).

[274] *Cf.* Railton, *supra* note 273, at 101-04 (suggesting that wrongfulness of risk imposition involves a lack of respect for person risked).

position itself.[275] Let's change the hypothetical so that *P*'s neighbor inadvertently fails to service his heating system, increasing the risk of an explosion, which fortunately fails to occur. Assume that *P* never learns of the risk. Now, the response (I submit) is either an intuition that *P* has not been harmed, or at most a weak intuition (weaker than in the Russian roulette case) that he has been. Systematic theories of welfare are crucial in deflating misleading intuitions, and in regimenting our welfare judgments where the intuitions are relatively weak.

Let us turn, then, to the systematic theory of welfare that, I have argued, is correct—namely, the substantive view.[276] Substantivists disagree about the correct list of welfare goods and the correct grounding for this list.[277] I will focus on the list provided by Martha Nussbaum, since it is quite extensive and subsumes most other lists currently in play in the philosophical literature.[278] If frequentist risk is not a welfare setback given Nussbaum's list, then a fortiori it is not a welfare setback given other plausible substantive accounts of welfare.

Nussbaum's list of "the central human functional capabilities" runs as follows. I have included part of her gloss on each item, where that gloss is helpful in determining whether the item should be understood to subsume freedom from fatality risks.

> 1. Life. Being able to live to the end of a human life of normal length; not dying prematurely, or before one's life is so reduced as to be not worth living.
>
> 2. Bodily Health. Being able to have good health, including reproductive health; to be adequately nourished; to have adequate shelter.
>
> 3. Bodily Integrity. Being able to move freely from place to place . . . to be secure against assault, including sexual assault . . . and domestic violence; having opportunities for sexual satisfaction and for choice in matters of reproduction.
>
> 4. Senses, Imagination, and Thought. . . . Being able to have pleasurable experiences, and to avoid non-necessary pain.
>
> 5. Emotions. . . . Not having one's emotional development blighted by overwhelming fear and anxiety . . . .

---

[275] *See* PORAT & STEIN, *supra* note 273, at 113-15 (identifying real harms, such as emotional or property damage, with which risk is often associated).

[276] *See supra* text accompanying notes 263-71.

[277] *See* Adler, *supra* note 20, at 1306-07.

[278] *See id.* at 1306 & n.49, 1307.

> 6. Practical Reason. . . .
>
> 7. Affiliation. . . . Having the social bases of self-respect and non-humiliation; being able to be treated as a dignified being whose worth is equal to that of others. . . .
>
> 8. Other Species. . . .
>
> 9. Play. . . .
>
> 10. Control over One's Environment. . . .[279]

A close analysis of this comprehensive and appealing list suggests that frequentist risk relative to canonical dosing classes, or for that matter relative to noncanonical reference classes, is not a welfare setback in the substantive sense.

Consider two outcomes, O_1 and O_2, differing only in the fact that *P* suffers a toxic exposure in O_1, which imposes a nontrivial frequentist risk on him but does not actually cause his death.[280] Is *P* worse off, in O_1, with respect to one or more of Nussbaum's welfare values? Because the toxin does not cause *P*'s death in O_1, and thus he does not die prematurely in that outcome relative to O_2, the first outcome is not worse with respect to the value of life itself—the first value on Nussbaum's list.

This is pretty obvious: of the multitudes exposed to the risks sufficient to concern EPA, OSHA, etc. (risks as low as perhaps 1 in 1 million, and in any event 1 in 1000), the vast majority will *not* live shorter lives, or have impaired health (Nussbaum's second value), as a result of those risks.[281] But perhaps we might say that everyone exposed to a

---

[279] NUSSBAUM, *supra* note 265, at 78-80 (footnotes omitted).

[280] It might be objected that there are other pairings of outcomes that might be used to test whether frequentist risk is a welfare setback in the ex post sense. Although this is true, using a different pairing wouldn't change my analysis, with respect to "life" or any other of Nussbaum's values. For example, consider the pair O_1* and O_2*: *P* suffers a larger toxic exposure and thus a larger frequentist risk relative to the canonical dosing class in the first outcome, and *P* dies as a result of the toxin in *both* worlds. These two outcomes, like the outcomes compared in the text (O_1 and O_2), are just the same with respect to Nussbaum's "life" value. Longevity, not the risk of truncated longevity, is what is at stake under this rubric.

[281] This is true analytically: if a randomly selected group of individuals is exposed to a toxin that has a 1 in *X* fatality risk relative to the canonical dosing class, then (if the group is large) the frequency of death within the group should approximately equal 1 in *X*. *See* HOWSON & URBACH, *supra* note 75, at 321-22 (stating Mises's "axiom of the existence of limits," a precondition for frequentist risk in the standard, Misean sense). For a dramatic illustration of the divergence between "individual risk" and ac-

frequentist risk of death suffers an infringement of bodily integrity (Nussbaum's third value). After all, a toxic exposure amounts to a kind of physical invasion, at least where the exposure is involuntary: an unconsented-to breach of the body's curtilage by a chemical agent.[282] The mistake, here, is in equating physical invasion with frequentist risk. *Everyone* who unknowingly eats, breathes, or dermally uptakes a toxic (or for that matter nontoxic) chemical suffers a bodily "invasion" of the kind just described. That is true of the person whose exposure produces a nonzero frequentist risk; it is also true of the person whose exposure produces a zero frequentist risk (for example, a person exposed to an amount of a noncarcinogenic toxin falling below the toxicity threshold and therefore resulting in a risk number less than one). One might object that only "small" doses produce subthreshold risks, but this is not true as a matter of toxicology. Physically large doses (doses containing many large molecules) can produce a subthreshold dose, and vice versa—it just depends on the toxin.[283] To be sure, doses below the threshold are necessarily "small" in a nonphysical sense—they don't produce a nonzero frequentist risk—but the notion of physical invasion does not explain why we care about the risk dimension, as opposed to some measure of sheer physical size (such as volume or mass). Note also that, where safety hazards rather than toxins are at issue (for example, unsafe cars, workplaces, or consumer products), a nonzero frequentist risk will typically occur without any physical invasion of the subject's body at all.[284]

Nussbaum's fifth value is "emotion." Fear and related emotional states, such as anxiety, clearly are objective welfare setbacks, for Nussbaum.[285] And there *is* a connection between risk and fear. *P* genuinely fears death only if *P* is, in some sense, at risk of death. But it is

---

tual death, see HAMILTON & VISCUSI, *supra* note 16, at 15. Hamilton and Viscusi examined 150 Superfund sites where EPA—using the test of "individual risk" to the reasonably maximally exposed individual—required remedies and found that "at the majority of sites the expected number of cancer cases averted is less than 0.1" and that "[o]nly ten sites had one or more expected cancer cases."

[282] *Cf.* JUDITH JARVIS THOMSON, THE REALM OF RIGHTS 205-26 (1990) (arguing for a moral right against bodily intrusion).

[283] *See* RODRICKS, *supra* note 150, at 45 ("The *conditions of exposure* under which toxic effects are produced—the size of the dose and the duration of dosing needed—vary greatly among chemicals.").

[284] *See* Railton, *supra* note 273, at 94-95 (distinguishing between risks imposed via an intrusion into the victim's person or property, and pure risk imposition).

[285] NUSSBAUM, *supra* note 265, at 79; *see also* Adler, *supra* note 20, at 1375-85 (arguing that fear is harm); Matthew D. Adler, *Fear Assessment: Cost-Benefit Analysis and the Pricing of Fear and Anxiety*, 79 CHI.-KENT L. REV. 977, 988-89 (2004) (same).

important to be careful in analyzing the fear-risk linkage. The kind of risk conceptually linked to genuine fear is first-party Bayesian risk.[286] *P* genuinely fears death only if *P* believes to a sufficient degree (has a sufficiently high Bayesian probability) that he will die and this belief produces physical arousal and feelings of distress. Frequentist risk, as opposed to Bayesian risk, has no a priori connection to fear. Any link would be contingent and empirical. *P* can be at high frequentist risk of death, but not believe he will die, and not experience attendant arousal and distress.[287] He can be at low frequentist risk from a toxin, yet still fear it—indeed, fear it quite rationally, if the belief-states animating his fear are epistemically rational.

Finally, Nussbaum counts "affiliation" as an objective welfare good. Other substantive welfare theorists propose similar goods: Finnis, "sociability"; Griffin, "deep personal relations."[288] It is here, surprisingly perhaps, that the ex post welfare relevance of frequentist risk becomes most plausible. Return to the Russian roulette game. Imagine that your neighbor intentionally fails to service his boiler, hoping to kill you—or, to wash out the suicide element and make the case purely homicidal, imagine that he builds a bomb in a shed on his lot and sets a timer that will explode the bomb while he's away and kill you. The bomb misfires, and you never learn of the neighbor's plot. So you haven't been physically harmed, nor scared, by his risky activity. Nonetheless, there seems to be a sense in which, qua the objective welfare good of "affiliation," your life history is worse than a life in which he didn't build the bomb. By deliberately imposing a fatality risk on you, without your consent, he failed (to quote Nussbaum) to "treat[] [you] as a dignified being whose worth is equal to that of others."[289]

Here, there is an overlap between welfarist moral theories and nonconsequentialist moral theories, which are considered below. Does the deliberate imposition of a high frequentist fatality risk on *P*, as such, really deny his full human worth and dignity? If so, the non-

---

[286] *See* Wayne A. Davis, *The Varieties of Fear*, 51 PHIL. STUD. 287, 289-302 (1987) (analyzing fear as an involuntary state of arousal and distress produced by the occurrent belief that some harm might occur); *see also* Adler, *supra* note 20, at 1375-76, 1377 & n.229 (presenting Davis's account and citing similar accounts).

[287] *See, e.g.*, Adler, *supra* note 20, at 1350-51, 1358-59, 1363 (arguing that frequentist risks need not be experienced); PORAT & STEIN, *supra* note 273, at 113 (same); Goldberg & Zipursky, *supra* note 273, at 1650-51 (distinguishing between the "intermediate harm" of risk and the "ultimate harm" of emotional distress).

[288] *See* FINNIS, *supra* note 264, at 88; GRIFFIN, *supra* note 267, at 30.

[289] NUSSBAUM, *supra* note 265, at 79.

consequentialist may well say that risk imposition is *wrong*: it violates a "deontological" norm. And the welfare consequentialist who incorporates the substantive welfare good that Nussbaum terms "affiliation," or some similar good, will see risk imposition as a full-blown, ex post, welfare setback. My response, to be fleshed out below, is that the social character of meaning and communication undermines any claim that the imposition of frequentist rather than Bayesian risk is a dignitary harm and thereby a welfare setback or deontological violation.[290]

To sum up: Frequentist fatality risk is not an ex post welfare setback. *P*'s life history is not worse in virtue of the various frequentist risks (relative to canonical dosing classes or any other classes) to which he may be subjected. No pair of possible worlds differing merely in *P*'s level of frequentist risk are different for *P*'s welfare (and obviously, then, they are not for anyone else's either). Therefore, any pair of possible worlds differing merely in someone's level of frequentist risk must be given the same ranking by the welfarist consequentialist. No member of any such pair can be ranked as better or worse without violating the "supervenience" requirement that is fundamental to welfare consequentialism. And so the justification for regulatory attention to individual, frequentist risk cannot be the "ex post" justification. It cannot be that worlds with these features are just worse, *ceteris paribus*, than other worlds—as are, for example, worlds where pain, fear, injury, or premature death occurs.

C. *The Ex Ante Question: Should Frequentist Risk Play a Role in the Choices of the Welfare-Consequentialist Regulator?*

A different welfare-consequentialist defense of "individual risk" invokes the ex ante perspective. Frequentist risk might have a role in determining morally appropriate *choices* even though it doesn't change the goodness of *outcomes*. Consider an actor—for our purposes, a governmental regulator—faced with a choice among various options. Welfare-consequentialist theories are moral theories, and therefore should provide criteria for evaluating these options. But the ex post component of a consequentialist theory does not do that. A ranking scheme for the different possible outcomes of an actor's options is not, yet, a ranking scheme for his options. Should the actor choose the action with the best outcome? The action with the greatest expected value? The action whose worst possible outcome is best? The "ex ante" component of a consequentialist theory provides a

[290] *See infra* text accompanying notes 389-92.

"bridging" rule that derives a ranking of the choices available to any actor, in any choice situation, from the ranking of the possible outcomes of these choices.[291] Perhaps frequentist risk becomes relevant at this stage of welfare consequentialism, rather than at the ex post stage.

Some bridging rules do not invoke probabilities. For example, the maximin rule says to determine the "worst case" outcome associated with each choice (the worst of all of the possible outcomes that might result from the choice) and to rank the choices accordingly.[292] But there is a plausible, indeed quite famous, bridging rule that *does* incorporate probability information: the rule of "expected utility" maximization. This rule says: assign to each outcome (using a technique too complicated to summarize) a "utility" number, which represents the goodness of that outcome; for each action and each of its possible outcomes, determine the *probability* that the action will result in that outcome; using these probability numbers, plus the "utility" numbers attached to outcomes, determine an "expected utility" number for each action equaling the probabilistically weighted average of all its possible outcomes; and then rank the actions available for choice in the order of their "expected utility" numbers.[293]

Perhaps a little formalism will help: Given a choice between $\{A_1 \ldots A_N\}$, where each choice has possible outcomes $\{O_{i,1} \ldots O_{i,M}\}$, there exists a utility function $U(O_{i,j})$ and a probability function $p(A_i, O_{i,j})$—the latter representing the probability that A_i results in $O_{i,j}$. Then the "expected utility" number assigned to each action A_i is: $\Sigma_j\, p(A_i\ ,\ O_{i,j}) \times U(O_{i,j})$. And the appropriate choice, ex ante, is the choice with the highest expected utility.

---

[291] *See* RICHARD A. FUMERTON, REASON AND MORALITY: A DEFENSE OF THE EGOCENTRIC PERSPECTIVE 92-113 (1990); KAGAN, *supra* note 258, at 64-69; Adler, *supra* note 20, at 1318-19.

[292] *See* SIMON FRENCH, DECISION THEORY: AN INTRODUCTION TO THE MATHEMATICS OF RATIONALITY 36 (1988).

[293] The "expected utility" rule has been hugely important within decision theory and economics. For good overviews, see ELLERY EELLS, RATIONAL DECISION AND CAUSALITY 65-86 (1982); FRENCH, *supra* note 292, at 149-209; DAVID M. KREPS, A COURSE IN MICROECONOMIC THEORY 71-131 (1990); DAVID M. KREPS, NOTES ON THE THEORY OF CHOICE (1988); MICHAEL D. RESNIK, CHOICES: AN INTRODUCTION TO DECISION THEORY 81-120 (1987). The seminal works in this area are: FRANK PLUMPTON RAMSEY, *Truth and Probability*, *in* THE FOUNDATIONS OF MATHEMATICS AND OTHER LOGICAL ESSAYS 156, 156-98 (R.B. Braithwaite ed., 1950), *reprinted in* STUDIES IN SUBJECTIVE PROBABILITY 61 (Henry E. Kyburg, Jr. & Howard E. Smokler eds., Henry E. Kyburg, Jr., trans., 1964); LEONARD J. SAVAGE, THE FOUNDATIONS OF STATISTICS (1954); and JOHN VON NEUMANN & OSKAR MORGENSTERN, THEORY OF GAMES AND ECONOMIC BEHAVIOR 15-31, 617-32 (Princeton Univ. Press 2004) (3d ed. 1953).

How is this relevant to risk regulation? The thought is that a regulator's attention to frequentist risk, in making her choices, is just the sort of responsiveness to probability information required by certain ex ante bridging rules, paradigmatically expected utility maximization. But this thought is misguided. Rules for regulatory choice that look to frequentist "individual risk"—be it risk to a particular person or some amalgam of risk to different persons, such as the average "individual risk"—will not track the expected utility rule. Conceptually, the two sorts of rules are quite different. The probabilities relevant for expected utility maximization take possible worlds as their arguments. Each $p(A_i, O_{i,j})$ is a probability of a particular action leading to a given total outcome. By contrast, a frequentist "individual risk," in the regulatory context, is the risk of a particular regulatory action leading to some type of adverse effect suffered by a particular person. These risks take persons and effects, not possible worlds, as their arguments. Their form is $p(A_i, E, P_k)$, where E is the type of effect (here, death) and P_k is the particular person. Probability here is being attributed to a small fragment of a possible world: the little bit of reality that would be changed if some person were to die.

To put the point another way, the probabilities relevant to expected utility maximization are the probabilities of extremely detailed and thorough descriptions of reality: the probabilities of the different possible total consequences of each action.[294] The "individual risk" numbers produced by risk assessment are the probabilities of very sketchy and incomplete descriptions of reality: descriptions that attend only to whether some person dies as a consequence of a toxic exposure or some other hazard. Furthermore, the probabilities relevant to expected utility maximization are *Bayesian*, not frequentist

[294] The expected utility rule tells the decision maker to assign values to acts based upon the possible outcomes of each act and the probability of the outcomes. If the decision maker is unboundedly rational (as she is traditionally assumed to be by economists and decision theorists), then the "outcomes" which are valued using the expected utility rule should be complete possible worlds rather than consequences in a more generic sense (technically, sets of possible worlds). Why? The value that an unboundedly rational decision maker should assign to a generic consequence—a set of possible worlds—is the weighted average of the values assigned to its component possible worlds, as weighted by their probabilities. So the unboundedly rational decision maker either starts with maximally specified outcomes, or starts with less specified outcomes but, in turn, values those by using an expected utility procedure incorporating maximally specified outcomes. *See* RICHARD C. JEFFREY, THE LOGIC OF DECISION 210 (2d ed. 1983) ("[I]t is only the complete consistent novels [i.e., complete possible worlds] that can be said to have nonprobabilistic values: the desirability of a proposition [a set of possible worlds] will be a probability-weighted average of the values of the possible worlds in which it would be true.").

probabilities.[295] The expected utility model builds upon a belief-desire account of rational action. "Utilities" are numbers representing the desirability of different outcomes (how much welfare they contain, how it is spread), while "probabilities" are numbers representing the decisionmaker's degree of belief that different outcomes will occur if different actions are chosen. Imagine, for example, that the regulator is certain that a particular outcome will not occur. Her Bayesian probability of its occurrence is zero. Then, regardless of the frequentist probability of that outcome, she should ignore the outcome in deciding which action to undertake.

It might be objected that the ex ante rule of expected utility maximization presupposes an unrealistically idealized regulator. The regulator is not omniscient, but her logical and conceptual abilities are unbounded. She is able to imagine the totality of possible worlds, each described in complete detail, and to ascribe probabilities to each. Real world regulators, by contrast, are boundedly rational—

---

[295] Admittedly, the expected utility rule *might* be understood as using objective (frequentist) rather than Bayesian (subjective) probabilities. The frequentist version of the rule originates with von Neumann and Morgenstern; the Bayesian version, with Ramsey and Savage. *See* KREPS, A COURSE IN MICROECONOMIC THEORY, *supra* note 293, at 112 ("The von Neumann-Morgenstern model, where probabilities are objective, and the Savage . . . model, where probabilities are subjective, are the chief models of consumer choice under uncertainty in microeconomics."); *see also* sources cited *supra* note 293. But frequentist expected utility maximization has serious flaws. First, the very point of ex ante rules within welfare consequentialism is to reflect the epistemic position of the decision maker. Omniscient decision makers should choose the action with the best outcome; but a nonomniscient decisionmaker might properly follow a different ex ante rule, given her uncertainty about what the outcomes of her choices will be. *See, e.g.*, EELLS, *supra* note 293, at 5 (arguing that Bayesian decision theory seeks to capture the idea that the rationality of a course of action is relative to the decision maker's beliefs and desires); *cf.* ALLAN GIBBARD, WISE CHOICES, APT FEELINGS: A THEORY OF NORMATIVE JUDGMENT 42-43 (1990) (arguing that moral rightness should be understood subjectively, not objectively). If this epistemic view of ex ante rules is rejected, then the rule for an omniscient decision maker might be something other than "choose the action with the best outcome"—which is absurd, since surely the consequentialist would want the omniscient decision maker to choose the action with the best outcome. Second, assigning frequentist probabilities to complete possible worlds is problematic. A complete possible world is a complete description of everything that might occur, and some (all?) of the occurrences therein described are nonrepeatable—paradigmatically, an occurrence that takes place at a particular time. *See infra* text accompanying notes 322-24 (discussing difficulty in assigning frequentist probabilities to nonrepeatable events). For rejections of the von Neumann-Morgenstern approach to expected utility theory, see EELLS, *supra* note 293, at 31-33; HIRSHLEIFER & RILEY, *supra* note 74, at 10; Philippe Mongin & Claude d'Aspremont, *Utility Theory and Ethics*, *in* 1 HANDBOOK OF UTILITY THEORY 371, 404 (Salvador Barberà et al. eds., 1998).

they are humans.[296] Expected utility maximization is an idealization: a plausible ex ante rule for demigods, perhaps, but not for us. Showing that the probabilities relevant to the rule of expected utility maximization are not "individual risks" hardly demonstrates that the ex ante rule appropriate for human regulators dispenses with "individual risk."

What *is* the ex ante rule appropriate for human regulators? What decision-procedure does welfare consequentialism instruct boundedly rational regulators to employ? I have answered this question in other work: cost-benefit analysis (CBA), a rough-and-ready, humanly accessible version of expected utility maximization, in which the salient welfare effects of regulatory choices are priced in dollars and a total dollar cost or benefit is determined for each choice, is generally the appropriate decision-procedure for regulators concerned to maximize overall welfare.[297] To be sure, welfare consequentialism is a broad family of moral views. Utilitarianism, which focuses on overall welfare, is only one member of the family.[298] But it seems plausible that CBA can be adapted to nonutilitarian variants of welfare consequentialism demanding sensitivity to the distribution of welfare (at least to prioritarianism and others that respect the Pareto principle), through the use of appropriate weighting factors—factors that give greater weight to the welfare of those who are worse off.[299]

Let us consider, then, whether "individual risk" in the frequentist sense has a role to play in CBA. Certainly there are forms of CBA that do *not* directly incorporate "individual risk" information. For example, current CBA practice at federal regulatory agencies is to calculate the monetary benefit of regulatory interventions that prevent deaths by predicting the total number of deaths prevented and then multiplying that number by the "value of statistical life," a dollar amount in

---

[296] *See* Matthew D. Adler, *Rational Choice, Rational Agenda-Setting, and Constitutional Law: Does the Constitution Require Basic or Strengthened Public Rationality?*, *in* LINKING POLITICS AND LAW 109, 131-32 (Christoph Engel & Adrienne Héritier eds., 2003). The *locus classicus* for discussions of bounded rationality is, of course, Herbert Simon's work. *See generally* 1 HERBERT A. SIMON, MODELS OF BOUNDED RATIONALITY: ECONOMIC ANALYSIS AND PUBLIC POLICY (1982); 2 *id.*; 3 HERBERT A. SIMON, MODELS OF BOUNDED RATIONALITY: ECONOMIC ANALYSIS AND PUBLIC POLICY (1997).

[297] *See* Adler & Posner, *Rethinking Cost-Benefit Analysis*, *supra* note 36, at 194-225 (conceptualizing cost-benefit analysis as a decision-procedure); Adler & Posner, *Implementing Cost-Benefit Analysis*, *supra* note 36, at 272-80 (same).

[298] *See supra* text accompanying notes 258-62.

[299] *See* Adler & Posner, *Rethinking Cost-Benefit Analysis*, *supra* note 36, at 224 (discussing possibility of distributively-weighted cost-benefit analysis); Adler & Posner, *Implementing Cost-Benefit Analysis*, *supra* note 36, at 300-05 (same).

the range of one to six million dollars.[300] CBA, thus structured, evaluates policy choices by characterizing the impact of those choices on "population risk," not "individual risk." The risk assessor predicts the aggregate deaths that would result from a hazard and be avoided by regulatory intervention; these mortality effects are then monetized at a price of one to six million dollars per death, and added to the non-mortality costs and benefits of the different regulatory options, such as compliance and enforcement costs.

To be sure, the "value of statistical life" number is itself derived from information about willingness to pay to avoid small risks of death, rather than willingness to pay to avoid certain death or willingness to accept in exchange for certain death (the first number will be dramatically skewed by individual wealth, while the second may well be infinite). But "individual risk" has no direct role in the kind of CBA just described. Particular regulatory interventions are evaluated by quantifying and then monetizing their effect in reducing the total number of deaths, not by quantifying and monetizing their effect on the distribution of "individual risk."

Still, it is possible to construct a different variant of CBA—one that does characterize and price the "individual risks" created by particular hazards.[301] For most hazards, even those where the regulator is very confident (ex ante) that one or more deaths will result from the hazard if left unregulated, the regulator will not know (ex ante) the identities of the persons who will die. Imagine, for example, a toxic

---

[300] For general discussions of the "value of statistical life" method, see A. MYRICK FREEMAN III, THE MEASUREMENT OF ENVIRONMENTAL AND RESOURCE VALUES: THEORY AND METHODS 298-321 (2d ed. 2003); W. KIP VISCUSI, FATAL TRADEOFFS: PUBLIC AND PRIVATE RESPONSIBILITIES FOR RISK 34-74 (1992); 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 MANAGEMENT, *supra* note 23, at 345, 345-75. Governmental use of the "value of statistical life" construct is surveyed in Don Kenkel, *Using Estimates of the Value of a Statistical Life in Evaluating Consumer Policy Regulations*, 26 J. CONSUMER POL'Y 1, 4-7 (2003), and W. Kip Viscusi & Joseph E. Aldy, *The Value of Statistical Life: A Critical Review of Market Estimates Throughout the World*, 27 J. RISK & UNCERTAINTY 5, 53-56 (2003).

[301] NRC's approach to CBA, it appears, *is* to price risks rather than deaths. The agency currently monetizes the health costs of radiation exposure at a price of $2000 per person-rem. But because the $2000 figure is used across the board to price exposures, this technique is not sensitive to variation in individual willingness-to-pay to avoid risk. *See* U.S. NUCLEAR REGULATORY COMM'N, REGULATORY ANALYSIS GUIDELINES, NUREG/BR-0058, at 31-32 (Rev. 4, Sept. 2004) (presenting NRC's approach to monetizing radiation exposure), *available at* http://www.nrc.gov/reading-rm/doc-collections/nuregs/brochures/br0058/br0058r4.pdf.

food that produces a 1 in 10,000 fatality risk for those who consume it—or so the regulator believes—and that is consumed by a population of one million. It is therefore overwhelmingly likely, the regulator believes, that the food will cause at least one premature death. Yet, for any given food consumer, the regulator assigns a probability of 1 in 10,000 that that particular individual will die if the food is left on the market. The regulator might monetize the benefits of regulation by pricing the mean number of possible deaths—in this example, one hundred deaths. That is the current practice, as just explained. Alternatively, the regulator might monetize the benefits of regulation by pricing each "individual risk" of death that regulation avoids, and aggregating across the population. In the food example, that would mean: determining how much different categories of food consumers, described in greater or lesser detail, are willing to pay to avoid a 1 in 10,000 fatality risk (for example, willingness to pay to avoid that risk might vary with wealth or age); determining how many individuals in the total population of consumers of the toxic food (one million) fall into the various categories; and then multiplying the numbers in each category by the category-specific willingness to pay to avoid a 1 in 10,000 risk. For example, if willingness to pay is individuated by age, wealth, and gender, the regulator would calculate the total risk-reduction benefit of banning the toxic food as illustrated in the table below:

Gender	Age	Wealth	Mean WTP	Total Number	Total WTP in Category
M	< 21	Poor	$150	25,000	$3,750,000
M	21 - 40	Poor	$100	30,000	$3,000,000
M	41 - 70	Poor	$175	40,000	$7,000,000
M	> 70	Poor	$250	10,000	$2,500,000
F	< 21	Poor	$150	25,000	$3,750,000
F	21 - 40	Poor	$100	35,000	$3,500,000
F	41 - 70	Poor	$250	45,000	$11,250,000
F	> 70	Poor	$300	15,000	$4,500,000
. . .	. . .	. . .	. . .	. . .	. . .

In short, CBA's methodology for monetizing the benefits of avoided fatalities might take two different forms: the "population risk" form, where mean total fatalities avoided are predicted and

priced;[302] or the "individual risk" form, where the "individual risks" of fatality imposed on various groups of individuals by hazards are characterized, priced, and aggregated. Choosing between these variants of CBA is a complicated matter, which I cannot pursue here. The choice depends on the degree of variability of willingness-to-pay to avoid risk across the population,[303] and the deliberation costs of disaggregating the population at risk from a hazard into different willingness-to-pay categories.[304] But both represent plausible procedures.

It should be stressed that the "individual risk" version of CBA does not presuppose the mistaken view, criticized above, that risk is itself an ex post harm. Like the "population risk" version, it assumes that *death* is the welfare setback targeted by risk regulation: it is death, not risk, that makes a human life worse and that legitimately influences the welfare-consequentialist ranking of possible worlds. "Individual risk" becomes relevant at the ex ante stage by virtue of the regulator's uncertainty about who will die.[305] Consider this analogy: where the regulator knows that one person in a population of one million will experience great pain as a result of a hazard, but doesn't know who the sufferer will be, that "pain cost" might be priced either as the amount a single person would pay to avoid certain pain, or as the aggregate amount that one million persons would pay to avoid a 1 in 1 million risk of pain. Neither approach to valuing pain presupposes that the *risk* of pain, as opposed to pain, is an ex post welfare setback. Ditto for the "individual risk" approach to valuing death: it *denies*, rather than presupposes, that the unwitting victim of a game of Russian roulette has been harmed as such.

Have we finally, then, identified a legitimate role for "individual risk" within regulatory choice governed by welfare consequentialism?

---

[302] More precisely, pricing the mean number of deaths is the simplest, and currently practiced, variant of the "population risk" approach to CBA, but not the only possibility. The well-informed analyst will have a Bayesian probability distribution over the possible number of deaths avoided by regulatory intervention. That Bayesian distribution can then be used, in a variety of different ways, to monetize the benefit of intervention.

[303] On heterogeneity in willingness-to-pay to avoid fatality risks, see FREEMAN, *supra* note 300, at 318; VISCUSI, FATAL TRADEOFFS, *supra* note 300, at 42-47.

[304] *See* Adler & Posner, *Rethinking Cost-Benefit Analysis*, *supra* note 36, at 216-43 (arguing that the design of CBA should be sensitive to deliberation costs); Adler, *supra* note 285, at 998-1001 (same).

[305] *See* FREEMAN, *supra* note 300, at 301 ("[One] way of characterizing the economic approach [to the valuation of life] is by saying that the economic value is derived by focusing on choices ex ante, that is, before the uncertainty about the individual's death during a specified period is resolved.").

No, because the "individual risks" relevant within CBA, if at all, are Bayesian not frequentist.[306] CBA, again, is a rough and ready version of expected utility maximization—one that does not require the regulator to imagine possible outcomes in complete detail, but instead asks her to focus on the welfare-relevant features of outcomes likely to be substantially affected by the choices before her, and to price those. Still, just as the probabilities that would properly drive the regulatory choices of an unboundedly rational regulator maximizing expected social utility are Bayesian probabilities—numbers measuring her own beliefs about the likelihood of different possible worlds—so the probability numbers figuring in CBA and driving regulatory choice under bounded rationality are also, appropriately, Bayesian. Consider once more the toxic food case. The food has a frequentist risk, relative to the canonical dosing class, of 1 in 10,000. One of every 10,000 persons who eat a standard serving of the food dies as a result. But (let us imagine) there is a subpopulation that (the regulator believes) is resistant to the toxin. Ethnic Albanians have a gene that, current scientific theory predicts, should almost always render the toxin harmless. Then the "individual risk" number that the regulator should use in determining the benefit that individual Albanians would gain from a food ban should be less than 1 in 10,000. If, for example, the regulator believes to degree 1 in 1 million (not 1 in 10,000) that the food will cause the death of a given Albanian consumer, then the monetized benefit for that consumer of a food ban should equal his willingness-to-pay to avoid a 1 in 1 million risk of death.

---

[306] *See id.* at 210 (suggesting that the relevant probabilities for purposes of CBA under uncertainty are Bayesian); Viscusi & Aldy, *supra* note 300, at 10 (same). Whether the Bayesian probabilities relevant to CBA should be first-party probabilities (Bayesian probabilities measuring the beliefs of the persons affected by the governmental choice being evaluated) or third-party probabilities (Bayesian probabilities measuring the beliefs of the CBA analyst or governmental official) is a further question. Given the status of CBA as a rough-and-ready approximation to expected utility maximization, and the fact that the relevant Bayesian probabilities for expected utility maximization are third party, not first party, it follows (I believe) that the probabilities relevant for CBA are also third party. *See* Matthew D. Adler, *The Puzzle of "Ex Ante Efficiency": Does Rational Approvability Have Moral Weight?*, 151 U. PA. L. REV. 1255, 1276-79 (2003) (arguing that government decision makers who are maximizing expected utility should not attend to first-party Bayesian probabilities); Mongin & d'Aspremont, *supra* note 295, at 437-44 (developing social choice theory using Bayesian probabilities, and discussing relevance of first-party probabilities). This is not a point I will argue at length here, since it is tangential to my main claim in this Article: namely that the probabilities that properly figure in CBA, and more generally in plausible governmental decision procedures for risk regulation, are Bayesian measures of someone's beliefs, not frequencies.

The claim I am making here—that the probabilities relevant to CBA are Bayesian, not frequentist—rests on a view about bounded rationality. For actors with unbounded cognitive resources, the best developed and most plausible model of rationality meshes beliefs and desires, and is formalized through the expected utility model. It is conceivable, I suppose, that the shift from unbounded to bounded cognitive resources radically changes the role of beliefs in rational choice.[307] Beliefs, and the numbers measuring them, might drop out of the picture. Where the boundedly rational regulator strongly believes that an outcome will (or will not) occur, but the frequency of the outcome relative to some reference class is low (or high), it *might* be rational for the regulator to scrap her beliefs and rely instead on frequencies as the ultimate determinants of her choice.[308]

But why think that? The topic of bounded rationality is undertheorized,[309] and I certainly don't have a full-blown account to offer here. If regulators were so epistemically flawed that frequentist probabilities, relative to certain stipulated classes, were generally closer to the "idealized" Bayesian probabilities (i.e., the Bayesian probabilities that the regulators would have, if unboundedly rational) than the regulators' actual Bayesian probabilities, then perhaps it might be rational for the regulators to ignore their actual Bayesian probabilities and substitute frequentist probabilities calculated using the stipulated classes. But there are many serious problems with this frequentist picture of rational choice under conditions of bounded rationality. First,

---

[307] It has been argued that the shift from unbounded to bounded rationality produces other large changes in the structure of rational choice—most famously, that boundedly rational decision makers should "satisfice" rather than maximizing expected value. *See* Michael Byron, *Satisficing and Optimality*, 109 ETHICS 67, 70-75 (1998).

[308] Even on this view, beliefs would play a role. It is difficult to imagine an account of rational choice where they wouldn't. (For any proposed rule of rationality, the chooser would have to be guided by her beliefs about how the rule applied to a given choice situation.) But on the view now under consideration, beliefs about the frequencies of outcomes relative to stipulated reference classes, rather than probabilistic beliefs about the outcomes themselves, would determine rational choice.

[309] To be sure, a vast empirical literature now exists documenting the extent to which individuals deviate from expected utility maximization. *See generally* REID HASTIE & ROBYN M. DAWES, RATIONAL CHOICE IN AN UNCERTAIN WORLD: THE PSYCHOLOGY OF JUDGMENT AND DECISION MAKING (2001) (surveying this literature). But bounded rationality as a topic in *normative* decision theory remains relatively neglected, *cf.* ARIEL RUBINSTEIN, MODELING BOUNDED RATIONALITY 3 (1998) ("[A]ttempts to model bounded rationality have yet to find the right track."), despite Simon's early efforts, *see* SIMON, *supra* note 296. For some recent attempts to develop norms for bounded decision makers, see GERD GIGERENZER ET AL., SIMPLE HEURISTICS THAT MAKE US SMART (1999).

there are actual humans who are pretty good at reasoning probabilistically along Bayesian lines—for example, risk assessors or other statisticians trained in Bayesian techniques[310]—and regulators can defer to them, just as they now defer to traditional, frequentist risk assessors.[311] Second, frequencies relative to reference classes which are finely specified to take account of the individual characteristics that would influence the degrees of belief of an idealized regulator cannot be observed; they must be inferred, and making that inference in a reliable way assumes a high degree of epistemic competence on the part of the regulator or those experts to whom she defers.[312] In our food example, the regulator will not have the option of determining the frequency of premature death relative to a reference class of "Albanian humans who eat a standard dose of the food" by actually feeding the food to Albanians and a control group, given the ethical limits on human experimentation.[313] Instead, she will have to infer that frequency from scientific theory, animal tests, and epidemiological data. If the regulator, or the expert risk assessor to whom she defers, has enough epistemic competence to do that, why not just use the expert's Bayesian probabilities? Conversely, if the frequentist probabilities putatively determinative of rational choice by boundedly rational regulators are stipulated to be directly observable—in the case of toxic risks, the frequencies directly observed in animal tests—then those frequencies will have to be defined relative to fairly crude reference classes, given the costs of experiments and other constraints such as the limited number of direct analogues for human characteristics in the animals that are the subject of toxicity experiments.[314] The fre-

---

[310] *See, e.g.*, sources cited *infra* note 330 (presenting Bayesian approaches to risk assessment).

[311] *See generally* ROGER M. COOKE, EXPERTS IN UNCERTAINTY: OPINION AND SUBJECTIVE PROBABILITY IN SCIENCE (1991) (discussing use of expert subjective probabilities in governmental decision making as well as other contexts); M. GRANGER MORGAN & MAX HENRION, UNCERTAINTY: A GUIDE TO DEALING WITH UNCERTAINTY IN QUANTITATIVE RISK AND POLICY ANALYSIS (1990) (same); M. Granger Morgan, *Uncertainty Analysis in Risk Assessment*, 4 HUM. & ECOLOGICAL RISK ASSESSMENT 25, 30-34 (1998) (discussing applicability of this approach to risk assessment).

[312] *Cf.* Finkel, *supra* note 86, at 305-08 (discussing different ways to infer the variation in individual cancer risk that results from variation in individual "susceptibility" to cancer).

[313] *See* COVELLO & MERKHOFER, *supra* note 17, at 149 ("Due to ethical and practical problems associated with experimenting with human subjects, controlled human exposure studies typically have been limited to the study of mild and reversible health effects (e.g., allergic skin reactions).").

[314] Indeed, even where frequencies are defined relative to canonical dosing classes rather than more finely specified classes, and estimated using dose-response experi-

quency of toxicity in a class of rats fed a ten-gram dose of the dangerous food might be directly extrapolated to predict the frequency of toxicity in a class of humans fed a ten-gram dose; but what rat experiment could be undertaken to directly predict the frequency of death in a group of Albanian humans fed a ten-gram dose?

* * *

To summarize the analysis to this point: Consequentialism is the class of moral views that specify the actions that regulators should choose with reference to the outcomes that those actions might lead to. Consequentialist theories have both an "ex post" part, describing the features of outcomes that make them better or worse, and an "ex ante" part, providing rules (perhaps probabilistic in form) for deriving a ranking of the choices available to the regulator from the "ex post" ranking of outcomes. *Welfare* consequentialism says that the only morally relevant feature of outcomes is human welfare. "Individual risk" in the frequentist sense has no place within welfare consequentialism. First, "individual risk" lacks ex post significance. A life where the subject at some point consumes a toxin that imposes a high frequentist fatality risk on her—relative to the canonical dosing class or any other reference class—is not worse, *ceteris paribus*, than a life where she is not thus put at risk. In other words, the frequentist risk of death, unlike death itself, or the pain or bodily harms that accompany disease and injury, or the fear of death, is not a "cost" for purposes of cost-benefit analysis.

Second, "individual risk" in the frequentist sense lacks ex ante significance. If regulators are as traditional economic theory assumes them to be—if they have unbounded cognitive resources—then ex ante they should maximize expected social utility. The probabilities relevant to regulatory expected utility maximization are Bayesian probabilities over complete possible worlds, not frequentist probabilities that particular persons will die prematurely. To be sure, regulators with bounded cognitive resources should follow some procedure more tractable than expected utility maximization. Specifically, at least where implementing a variant of welfare consequentialism that is not radically egalitarian, regulators should engage in cost-benefit analysis. One plausible version of CBA uses "population risk" as an input, pricing aggregate predicted deaths; another plausible version

---

ments in rats, those frequencies are not literally directly observed. Rather, inference procedures must be used to control for experimental error and to extrapolate from high to low dose and from animals to humans. *See id.* at 137-43.

seeks instead to describe the distribution of "individual risk" across the population and prices these risks individually. But the kind of risk relevant to this latter sort of CBA is *Bayesian* rather than frequentist. Some subject's "individual risk" of dying prematurely as a result of a toxin is (in this context) the regulator's degree of belief that the subject will die prematurely, not the frequency of death relative to the canonical dosing class or any other reference class.

One possible objection to my analysis is that I haven't adequately considered egalitarian variants of welfare consequentialism. I have suggested that some of these variants might be implemented through CBA, by using "distributive weights" to adjust costs and benefits; but other variants would warrant non-CBA procedures. As a clear example, consider strict welfare egalitarianism, which rejects the Pareto principle,[315] enjoins regulators to aim at equalizing welfare levels, and therefore considers an outcome in which everyone's welfare level is the same to be better than an outcome in which some persons have that level and others have a higher level. This radical view would not justify CBA, but neither would it justify regulatory attention to frequentist risk—certainly not if regulators have unbounded cognitive resources,[316] and plausibly not if their rationality is bounded.[317]

A different general objection to my welfare-consequentialist analysis is that my persistent distinction between frequentist and Bayesian probabilities is illusory. This objection merits a longer response, which I shall now attempt to provide.

---

[315] *See supra* text accompanying note 261.

[316] Strictly egalitarian regulators with unbounded cognitive resources would maximize expected utility, with the crucial twist that the "utility" numbers ranking worlds would track the degree to which welfare is equally distributed in the worlds, not the total amount of welfare. Probabilities would be relevant ex ante, not ex post, but—as with utilitarianism, prioritarianism, or other welfare-consequentialist views—these would be the regulator's Bayesian probabilities over total outcomes.

[317] Plausibly, the boundedly rational egalitarian regulator would compare, in a rough way, the expected welfare of those at risk from hazards with the expected welfare of the general population or some other control population. If those at risk have lower expected welfare, shutting down the hazard increases equality. If those at risk have higher expected welfare, then the hazard should not be shut down. And, plausibly, the boundedly rational regulator's assessment of the expected welfare of the exposed population versus controls would hinge on his own Bayesian probabilities—since (as before) it is hard to see why bounding rationality would change the kind of probabilities relevant to regulatory choice.

IV. FREQUENTIST RISK AND BAYESIAN RISK: ARE THEY REALLY DIFFERENT?

I have tried to demonstrate that "individual risk" in the frequentist sense would have no role in the decision procedures of regulators choosing in accordance with welfare consequentialism. Probabilities may play a role, at various points, but these are Bayesian rather than frequentist probabilities, or so I have argued. The reader is surely entitled to wonder whether these two kinds of probabilities are genuinely distinct. Don't Bayesian probabilities reduce, ultimately, to relative frequencies? Doesn't the rational regulator's degree of belief that some person will die prematurely as a result of a toxin equal the frequency with which persons exposed to the toxin die prematurely—if not the frequency relative to the canonical dosing class, then surely the frequency relative to some other reference class? How, concretely, would Bayesian risk assessment of toxins differ from the frequentist risk assessment that is now the norm?[318]

The Bayesian account of probability, generically, sees probabilities as numbers that measure a person's beliefs in various propositions and that satisfy the probability calculus. The numbers must not be greater than 1 or less than 0; if the probability of some proposition is p, then the probability of its negation must be $1\text{-}p$; if two propositions are mutually exclusive, then the probability of their disjunction must equal the sum of their probabilities; and so on.[319] Various Bayesian accounts for attaching numbers to beliefs have been offered by philosophers and mathematicians. To give one example: if an individual has comparative beliefs, beliefs in the relative likelihood of different propositions, that are well-behaved, then (it turns out) those beliefs can be represented by a scheme of numbers conforming to the probability calculus. Well-behaved, in this context, means something like this: (1) *completeness* (for given propositions *X* and *Y*, either our believer *B* believes *X* to be more likely than *Y*, or she believes *Y* to be more likely than *X*, or she believes them to be equally likely); (2) *transitivity* (if *B* believes *X* to be more likely than *Y* and *Y* to be more likely

---

[318] For an excellent discussion of the difference between Bayesian and frequentist approaches to environmental policy, published after this Article was drafted, see David E. Adelman, *Scientific Activism and Restraint: The Interplay of Statistics, Judgment, and Procedure in Environmental Law*, 79 NOTRE DAME L. REV. 497 (2004).

[319] *See* Adler, *supra* note 20, at 1312 n.73 (citing overviews of Bayesian theories of probability). On Bayesian statistical techniques, see BAYESIAN BIOSTATISTICS (Donald A. Berry & Dalene K. Stangl eds., 1996); PETER M. LEE, BAYESIAN STATISTICS: AN INTRODUCTION (2d ed. 1997).

than *Z*, then she believes *X* to be more likely than *Z*); and (3) *additivity* (if *B* believes *X* to be more likely than *Y* and both *X* and *Y* to be mutually exclusive of *W*, then she believes the disjunction *X*-or-*W* to be more likely than *Y*-or-*W*). If *B*'s beliefs satisfy these axioms and several other technical ones, then there exists a unique probability function mapping each belief onto a number between 0 and 1.[320] These numbers *represent* the beliefs of *B* in the sense that *p(X)* is greater than *p(Y)* if and only if *B* is more certain in *X* than in *Y*.

The specific Bayesian approach just described, and Bayesian accounts more generally,[321] do *not* draw a foundational link between probabilities and relative frequencies. The fact that *B* is comparatively more certain in *X* than in *Y* might be taken as a primitive, or analyzed in behavioral terms: if *B* is comparatively more certain in *X*, then if given the choice between a lottery which makes some attractive payoff contingent on *X* and one that makes the very same payoff contingent on *Y*, she will choose the first. In any event, there is no insistence that *X* and *Y* must be disaggregated into their component, repeatable events and that *B*'s comparative belief in *X* and *Y* must be reduced to the comparative frequency of these events relative to stipulated, canonical reference classes. Indeed, one virtue of Bayesian approaches, both this approach and others, is that they permit the attachment of probability numbers to propositions where ascribing frequentist probabilities is unthinkable or at least extremely silly.[322] Consider, for example, a proposition that makes some temporal assertion, claiming an event to have occurred at a particular time: "Nixon resigned in 1974." We might imagine an infinite series of political experiments where Presidents "like Nixon" are placed in similar political contexts and either resign or not; but this would be the frequency with which Presidents like Nixon resign at some time, not the frequency with which they resign at a particular time (1974). Particular times are necessar-

---

[320] *See* Patrick Suppes, *Qualitative Theory of Subjective Probability*, *in* SUBJECTIVE PROBABILITY 17, 27 (George Wright & Peter Ayton eds., 1994).

[321] Another approach sees Bayesian probabilities as primitives. *See* FRENCH, *supra* note 292, at 224-32. Yet another, pioneered by de Finetti, reduces them to odds for small money bets. *See* Bruno de Finetti, *Foresight: Its Logical Laws, Its Subjective Sources* (1937), *reprinted in* STUDIES IN SUBJECTIVE PROBABILITY, *supra* note 293, at 93. And some (most famously Ramsey, Savage, and Jeffrey) have proposed techniques for using a rational agent's choices to derive both a "utility" measure of her preference for outcomes and a "probability" measure of her beliefs. *See* JEFFREY, *supra* note 294, at 59-163; RAMSEY, *supra* note 293; SAVAGE, *supra* note 293, at 6-104.

[322] *See, e.g.*, COVELLO & MERKHOFER, *supra* note 17, at 210 ("Bayes's theory is more general than [frequentist] probability in that probabilities may be assigned to any meaningful hypothesis or proposition . . . .").

ily unrepeatable.[323] Or, more relevant to risk regulation, imagine a proposition that asserts some scientific model of toxicity to be true. What could the frequentist probability of a given scientific model possibly consist in?[324] Could it consist in the frequency with which a very large or infinite class of "similar" models—models that fall in a reference class characterized by some attribute that this particular model has—turn out to be true? How could we possibly have grounds for believing what that frequency is, and why on earth would we care?

Thus far I have discussed Bayesian approaches at a high level of abstraction: as theories stipulating some criterion for mapping individual beliefs onto numbers satisfying the probability calculus. More concretely, how would Bayesian approaches be employed to guide regulatory choice? These approaches have been widely used in certain areas of regulation, such as predicting the probability of a nuclear reactor accident.[325] The risk assessor consults with experts in nuclear engineering to identify the different series of occurrences within the reactor that might lead to core damage. Then, various techniques are used to ensure that each expert has a coherent set of beliefs with respect to the possible occurrences, sufficient to allow the ascription of Bayesian probabilities and to elicit those probabilities. (For example, the expert might be asked the odds at which he'd be willing to make a small monetary bet on each occurrence or its subevents; if those odds are incoherent, in the sense that they fail to satisfy the probability calculus, then the expert might be encouraged to revise his assessment.)[326] Finally, the different experts' probabilities are averaged or otherwise amalgamated by the risk assessor, so as to arrive at a single,

---

[323] *See id.* at 209-10 (discussing difficulty in assigning frequentist probabilities to nonrepeatable events); Adler, *supra* note 20, at 1345 n.158 (citing sources that describe this difficulty).

[324] *Cf.* HOWSON & URBACH, *supra* note 75, at 203 (noting that classical statistics refuses to assign probabilities to scientific theories).

[325] For a general discussion of how to integrate Bayesian probabilities into policymaking, *see* FUMIKA OUCHI, A LITERATURE REVIEW ON THE USE OF EXPERT OPINION IN PROBABILISTIC RISK ANALYSIS (Feb. 2004) (World Bank Policy Research, Working Paper 3201), *available at* http://econ.worldbank.org/files/32873_wps3201.pdf; Elisabeth Pate-Cornell, *Risk and Uncertainty Analysis in Government Safety Decisions*, 22 RISK ANALYSIS 633 (2002); sources cited *supra* note 311. On the use of expert subjective probabilities in nuclear reactor risk assessment, see COOKE, *supra* note 311, at 27-41. For a recent example, see Truong V. Vo et al., *Estimates of Rupture Probabilities for Nuclear Power Plant Components: Expert Judgment Elicitation*, 96 NUCLEAR TECH. 259 (1991).

[326] *See generally* COOKE, *supra* note 311, at 121-57 (discussing elicitation of expert probabilities); MORGAN & HENRION, *supra* note 311, at 102-71 (same); DETLOF VON WINTERFELDT & WARD EDWARDS, DECISION ANALYSIS AND BEHAVIORAL RESEARCH 112-33 (1986) (same).

pooled set of Bayesian probabilities of the various possible occurrences leading to reactor failure.[327]

In the area of toxic risk assessment, and more generally in ascribing probabilities of death, illness, and injury, Bayesian techniques have not been much used to date by U.S. regulators.[328] Nuclear regulators, for example, have been Bayesians in predicting the probability of reactor failure, but not in predicting the deaths that would result from a failure. And frequentism has dominated toxics regulation, as discussed.[329] Still, there is a growing scholarly literature that explores and defends the Bayesian variant of toxic risk assessment.[330] The basic

---

[327] *See* COOKE, *supra* note 311, at 171-220 (discussing combination of expert probabilities); WINTERFELDT & EDWARDS, *supra* note 326, at 133-36 (same).

[328] One important exception is EPA's use of expert subjective probabilities in regulating "criteria" air pollutants. *See* MORGAN & HENRION, *supra* note 311, at 9-11, 154-57.

[329] *See supra* Parts II.A-II.B. The problem of "uncertainty" is now reasonably high on the agenda of governmental risk assessors. *See, e.g.*, Fred Hansen (Deputy Administrator), EPA, *Policy for Use of Probabilistic Analysis in Risk Assessment* (May 15, 1997), *available at* http://www.epa.gov/osa/spc/htm/probpol.htm. The focus here, however, is exposure assessment rather than dose-response assessment. *See* CULLEN & FREY, *supra* note 58, at 2. And, in any event (oddly enough), it seems to be a widespread view among risk assessors that the probabilities quantifying uncertainty need not be Bayesian. *See* COVELLO & MERKHOFER, *supra* note 17, at 207-24; CULLEN & FREY, *supra* note 58, at 15-21.

[330] *See* COVELLO & MERKHOFER, *supra* note 17, at 213-16; CULLEN & FREY, *supra* note 58, at 18-21; Catherine Petito Boyce, *Comparison of Approaches for Developing Distributions for Carcinogenic Slope Factors*, 4 HUM. & ECOLOGICAL RISK ASSESSMENT 527, 535-41 (1998); James D. Englehardt, *Predictive Bayesian Dose-Response Assessment for Appraising Absolute Health Risk from Available Information*, 10 HUM. & ECOLOGICAL RISK ASSESSMENT 69 (2004); John S. Evans et al., *A Distributional Approach to Characterizing Low-Dose Cancer Risk*, *in* LOW-DOSE EXTRAPOLATION OF CANCER RISKS, *supra* note 86, at 253; John S. Evans et al., *Use of Probabilistic Expert Judgment in Uncertainty Analysis of Carcinogenic Potency*, 20 REG. TOXICOLOGY & PHARMACOLOGY 15 (1994); William E. Fayerweather et al., *Quantifying Uncertainty in a Risk Assessment Using Human Data*, 19 RISK ANALYSIS 1077 (1999); Vic Hasselblad & Annie M. Jarabek, *Dose-Response Analysis of Toxic Chemicals*, *in* BAYESIAN BIOSTATISTICS, *supra* note 319, at 235; Ryan A. Hill, *From Science to Decision-Making: The Applicability of Bayesian Methods to Risk Assessment*, 2 HUM. & ECOLOGICAL RISK ASSESSMENT 636 (1996); Gary Koop & Lise Tole, *Measuring the Health Effects of Air Pollution: To What Extent Can We Really Say That People Are Dying from Bad Air*, 47 J. ENVTL. ECON. & MGMT. 30 (2004); Jane C. Lindsey & Louise M. Ryan, *Analyzing Rodent Tumorigenicity Experiments Using Expert Knowledge*, *in* BAYESIAN BIOSTATISTICS, *supra* note 319, at 503; Morgan, *supra* note 311; Tapan K. Nayak & Subrata Kundu, *Calculating and Describing Uncertainty in Risk Assessment: The Bayesian Approach*, 7 HUM. & ECOLOGICAL RISK ASSESSMENT 307 (2001); D. Warner North, *Use of Expert Judgment on Cancer Dose-Response: Probabilistic Assessment and Plans for Application to Dieldrin*, *in* LOW-DOSE EXTRAPOLATION OF CANCER RISKS, *supra* note 86, at 275; Anthony O'Hagan, *Uncertainty in Toxicological Predictions: The Bayesian Approach to Statistics*, *in* FORECASTING THE ENVIRONMENTAL FATE AND EFFECTS OF CHEMICALS 25 (Philip S. Rainbow et al. eds., 2001).

idea, I propose, is this. Different models of toxicity exist. Some are deterministic, positing that under certain conditions an exposed person will definitely suffer the toxic effect, and that under other conditions she definitely will not. The "conditions" might include the size of the dose, the genetic makeup of the person, her health history, and so on. Other models are stochastic: they attach an irreducible risk of toxicity, greater than zero and less than one, to different kinds of exposures. For a given toxic exposure of a person, characterized in terms of the dose she received and perhaps other characteristics, the risk assessor's and therewith the regulator's Bayesian probability that the individual will die prematurely as a result of that exposure can be determined by: (1) consulting with toxicologists, biologists, and other experts to determine the different plausible models of toxicity, be they deterministic or stochastic; (2) for each expert, eliciting her Bayesian probabilities that the different possible models are true; (3) for each expert and each deterministic model, determining her Bayesian probability that the exposure being analyzed satisfies the conditions requisite for toxicity (here, death); (4) for each expert and each stochastic model, determining her Bayesian probability that the conditions requisite for an irreducible risk of toxicity are satisfied, and discounting that risk by that Bayesian probability; (5) for each expert, determining an all-in Bayesian probability that the exposure will result in death by discounting the probability conditional on each model by the probability that the model is true, and then summing these amounts across all the models; and (6) averaging or otherwise pooling the experts' all-in Bayesian probabilities.[331]

"Individual risks" that are ascribed using this Bayesian approach need not correspond to frequencies relative to reference classes, be they canonical dosing classes or other classes. Assume, first, that the relevant models are all deterministic. This assumption is not that

---

[331] Admittedly, much of the Bayesian risk assessment literature cited *supra* note 330 employs the sort of Bayesian "probability tree" sketched above, where experts assign subjective probabilities to nested sets of mutually exclusive alternatives, to determine a pooled expert estimate of frequentist risk relative to a canonical dosing class—more precisely, a pooled second-order Bayesian probability distribution over possible first-order frequentist probabilities—and not a Bayesian probability that some person will die. In this sense, the existing literature is incompletely Bayesian. *See* T. Aven & J.T. Kvaloy, *Implementing the Bayesian Paradigm in Risk Analysis*, 78 RELIABILITY ENG'G & SYS. SAFETY 195, 197 (2002). Still, it seems that a "probability tree" approach allowing for model uncertainty, parameter uncertainty, and other kinds of uncertainty would be the natural way to arrive at first-order Bayesian probabilities of some person's death. *See* COVELLO & MERKHOFER, *supra* note 17, at 220 (discussing Bayesian probability trees); LEE, *supra* note 319, at 35-36 (discussing Bayesian prediction).

strange: toxicologists have traditionally relied upon deterministic models for noncarcinogenic toxins, positing that each person has an individual threshold such that doses above the threshold definitely result in the toxic effect and doses below it do not.[332] Imagine, now, that the risk assessor, relying upon experts, is ascribing a Bayesian risk of death to a person (the Subject) individuated in terms of his dose (one gram), rough age (adult), and ethnicity (Danish).[333] Danes typically have some gene which may confer resistance to the toxin. The expert believes to degree 0.4 that the gene confers resistance, and to degree 0.6 that it does not. She believes to degree 0.9, based on published genetic studies of Danes, that this individual has the gene. Other work suggests very strongly that the frequency with which adults generally die as a result of the one-gram dose is 1 in 20,000. (This work consists, say, in dose-response studies on adult rats plus epidemiological evidence.) For simplicity, assume that the expert is virtually certain that the frequency of mortality from one gram dosings in the general adult population is 1 in 20,000. Based on this frequency, the expert determines that 1 in 20,000 adults who lack the (possibly) resistant gene have a threshold below one gram, and 19,999 have a threshold exceeding or equaling one gram.[334]

With all this information, the expert's Bayesian probability that the adult, Danish Subject will die prematurely as result of the one-gram dose is:

$$
\begin{aligned}
&[.9 \times .4 \times 0] + \\
&[.9 \times .6 \times (1/20{,}000) \times 1] + \\
&[.9 \times .6 \times (19{,}999/20{,}000) \times 0] + \\
&[.1 \times (1/20{,}000) \times 1] + \\
&[.1 \times (19{,}999/20{,}000) \times 0] \quad = .64 \times (1/20{,}000).
\end{aligned}
$$

---

[332] *See* CRAWFORD-BROWN, *supra* note 28, at 128; RODRICKS, *supra* note 150, at 166-70.

[333] Individual susceptibility to deterministic or "threshold" toxins, as to carcinogens, can certainly vary. *See* Dale Hattis et al., *Human Interindividual Variability in Parameters Related to Health Risks*, 19 RISK ANALYSIS 711, 716-22 (1999) (non-cancer effects); sources cited *supra* note 86 (cancer). Indeed interindividual variability in susceptibility to noncarcinogens is part of the rationale for using "safety factors" to derive a reference dose from a NOAEL. *See* Dourson et al., *supra* note 153, at 110-11. For a discussion of quantitative models that presuppose interindividual variability in toxicity thresholds, see CRAWFORD-BROWN, *supra* note 28, at 105-6.

[334] I am additionally assuming, here, that the expert is quite sure that only a small proportion of the general population has the gene. She is therefore able to infer from a 1 in 20,000 frequency of mortality in the general population that there is approximately a 1 in 20,000 frequency of mortality in the population lacking the gene.

The different terms in this equation represent the different, mutually exclusive possibilities as the expert sees them. First, the Subject might have the gene and the gene might confer resistance, in which case the Subject determinately will not die. Second, the Subject might have the gene, the gene might fail to confer resistance, and the Subject might have a toxicity threshold below one gram, in which case he determinately will die. Third, the Subject might have the gene and the gene might fail to confer resistance, yet the Subject might (by virtue of the rest of his genetic makeup, his health history, etc.) have a toxicity threshold greater than or equal to one gram, in which case he determinately will not die. Fourth, the Subject might lack the gene and have a toxicity threshold below one gram, in which case he determinately will die. Finally, the Subject might lack the gene and have a toxicity threshold greater than or equal to one gram, in which case he determinately will not die.

The expert's Bayesian probability that the Subject will die (in this example .64 × (1/20,000)) clearly does not equal the Subject's frequentist risk of death relative to the canonical dosing class. The proportion of all individuals (children as well as adults) who die prematurely as a result of a one-gram dose might be 1 in 10,000, not 1 in 20,000 or 1/20,000 discounted by a factor of 0.64. Nor does the expert's Bayesian probability of the Subject's death equal the Subject's frequentist risk of death relative to a maximally specified reference class—a reference class specifying all of his characteristics, including the dose amount, his genetic code, and the current physiological condition of his body.[335] The frequency relative to this class will either be zero (if in fact individuals with that genetic code and physiological condition have a threshold equal to or exceeding one gram) or one (if in fact someone with that genetic code, etc., has a threshold below it).[336] Nor, finally, does the expert's Bayesian probability equal the Subject's frequentist risk relative to the reference class of "adult Danes consuming a one gram dose": a reference class defined by the characteristics known to the regulator.

This last point is subtle, and to see it we should distinguish between (1) the actual frequentist risk, (2) the frequentist risk directly observed in an experiment, and (3) the frequentist risk that the ex-

---

[335] *See* Adler, *supra* note 20, at 1352, 1357-62 (discussing maximally specified reference classes).

[336] *See id.* at 1360; Perry, *supra* note 84, at 336.

pert believes to obtain. In the example at hand, the actual frequentist risk relative to the reference class of "one gram dosings by adult Danes" might well be 1 in 20,000. But the expert doesn't know that; instead, she believes to degree 0.4 that Danes often possess a gene that confers genetic resistance, and therefore ascribes a lower Bayesian probability to the Subject's death. Nor would it be ethical to actually give a group of Danish adults a one-gram dose and observe how many die.[337] Even if such a grisly experiment had been performed with a result of one in *X* deaths, the expert might not believe to degree one in *X* that the Subject will die. To give a stark example along these horrific lines: if twenty adult Danes in an experimental group of 1000 die prematurely as a result of a one-gram dose of the toxin, the expert would not necessarily ascribe a Bayesian probability of 1 in 50 to the Subject's death. Instead, assuming that the expert remains confident that ninety percent of Danes have the particular gene posited by one possible genomic model of toxicity to confer resistance, this experimental result would simply induce the expert to revise downwards her "prior" probability that the model is true, and perhaps to change her estimate of the frequency of a one-gram threshold in the general population.[338] What the revised numbers would be is a complicated question in Bayesian statistics, but it should be clear that the expert's revised (or "posterior") Bayesian probability of the Subject's death need not equal 1 in 50. Assume, at the extreme, that the grisly experimental results lead the expert to reject the gene-resistance model entirely. She now assigns that model a probability of zero and believes that adult Danes have the same distribution of thresholds as the general adult population.[339] But many prior studies have suggested that

---

[337] *See, e.g.*, COVELLO & MERKHOFER, *supra* note 17, at 149-50 (discussing the ethical limitations on the use of human experiments in risk assessment).

[338] It is a mainstay of Bayesian statistics that an experimental observation should not directly determine an observer's degree of belief in a model, a model parameter, or some other object of belief, but instead should be integrated with the observer's "prior" (pre-experimental) beliefs about the model, parameter, or other object. For a general discussion, see LEE, *supra* note 319, at 33-73. For good discussions with reference to risk assessment, see O'Hagan, *supra* note 330; Nathan O. Siu & Dana L. Kelly, *Bayesian Parameter Estimation in Probabilistic Risk Assessment*, 62 RELIABILITY ENGINEERING & SYS. SAFETY 89 (1998).

[339] Another possibility, suggested by these results, is that Danes are particularly vulnerable—in other words, that (1) the gene-resistance model is untrue and (2) something in the genetic makeup or health history of Danes produces *lower* thresholds than the general population. Even taking this possibility into account, the expert's posterior probability of the Subject's death, given the grisly experimental data, need not equal 1 in 50—because the expert will have prior beliefs about how likely it is that

only 1 in 20,000 adults have a threshold lower than one gram. The current study suggests that the number is 1 in 50. Integrating all this information, the expert will ascribe a nonzero probability to the possibility that the fraction of adults with a threshold lower than one gram is 1 in 50. Yet she need not be certain of that possibility, for again there is considerable, conflicting, prior data; she need not assign it a Bayesian probability of one. Therefore, despite the grisly experimental result, she need not ascribe a 1 in 50 Bayesian probability to the Subject's premature death.

So the expert's Bayesian probability that the Subject, characterized as an adult Dane consuming a one-gram dose, will die as a result of that exposure need not equal (1) the actual frequency with which adult Danes die as a result of a one-gram dose nor (2) the observed frequency in an experiment on adult Danes. Finally, it need not equal (3) the frequency with which adult Danes die as a result of a one-gram dose that the expert believes to obtain. Return to the original example sans grisly experiment. The expert believes to degree 0.4 that a gene (which she believes ninety percent of Danes possess) confers resistance; to degree 0.6 that it does not; and that 1 in 20,000 adults lacking the gene have a toxicity threshold below one gram. She therefore believes to degree 0.4 that the frequency of death among adult Danes consuming one-gram doses is 1 in 200,000 (1 in 20,000 times 0.1, the proportion of Danes lacking the gene) and to degree 0.6 that it is 1 in 20,000. There is no single frequency that the expert wholeheartedly believes to be the true frequency; rather, she believes to some degree (less than one) that one frequency is the true one, and to some degree (also less than one) that a different, incompatible frequency is the true one. We might stipulate that the expert "believes" in a frequency if she thinks it more likely than not—if she ascribes it a Bayesian probability greater than 0.5. In this example, that would be 1 in 20,000: more likely than not, thinks the expert, the model conferring resistance on the gene is untrue. But the expert's Bayesian probability that the Subject will die prematurely is *less* than 1 in 20,000, since it also incorporates the expert's probabilistic belief (here, 0.4) that the less likely model is correct.

This example, I hope, dispels the thought that the difference between Bayesian and frequentist risk assessment is insubstantial. A regulator's degree of belief that some individual will die prematurely

---

Danes have lower thresholds and, if so, what proportion of the Danish population has a threshold below one gram.

as a result of a toxic exposure, formulated in consultation with risk assessors and the experts that the assessors in turn consult, certainly need not equal the actual, experimental, or believed frequency of death relative to the canonical dosing class. After all, regulators and assessors will sometimes have more information about the individual than her dose, and this information may influence their Bayesian probability ascriptions. Nor (more subtly) need the regulator's degree of belief equal the actual, experimental, or believed frequency of death relative to some reference class more finely specified than the canonical dosing class, be it a maximally or non-maximally specified class. Regulators are not omniscient and may well not know what the actual frequency is; on a Bayesian view, frequencies observed in any given experiment need not determine the expert's degree of belief, but must be integrated with prior experiments and scientific theory; and, whatever the reference class, there is no single frequency relative to that class which the expert need believe wholeheartedly to obtain.[340]

To be sure, information about frequencies and proportions, so central to traditional risk assessment, will continue to have an important role in Bayesian practices. Animal dose-response experiments (demonstrating that various fractions of rat populations die after consuming various toxins at various doses) and epidemiological data (showing the proportions of groups of humans at various exposure levels that have died in the past) will shape the Bayesian probabilities of experts, assessors, and regulators.[341] They will do so by influencing the experts', assessors', and regulators' degrees of belief in different models of toxicity and, for a given model, in different model parameters and in the distribution of causal features relevant to the model

---

[340] In the example at hand, it *is* true that the expert's Bayesian probability of the Subject's death (.64 × (1/20,000)) equals the weighted average of the possible frequencies relative to the Subject's characteristics that the expert knows about (here his dose, age, and ethnicity), summed across possible models and discounted by the expert's belief in the different models—in this instance .4 × (1/200,000) + .6 × (1/20,000). But even this sort of connection between Bayesian and frequentist probabilities need not hold true—for it presupposes that the expert knows for sure that the Subject has some characteristics (which are used to define the reference class) and otherwise believes him to have causally relevant characteristics just to the degree that those exist in the reference class. In any event, I am not denying that there is some connection between Bayesian and frequentist probabilities. The point, rather, is that Bayesian probabilities cannot be reduced to frequencies—they are a different kind of entity—and thus none of the straightforward connections, mooted in the text, obtain. *See infra* note 346 (acknowledging one connection between frequentist and Bayesian probability in the case of stochastic models).

[341] *See, e.g.,* O'Hagan, *supra* note 330, at 30-32 (describing the Bayesian approach to dose-response data).

across the population. I say "influencing" rather than "determining" because the experts', assessors', and regulators' Bayesian probabilities will also be influenced by general scientific theory and by experiments that do not produce frequency data (for example, molecular structure or mutagenicity tests).[342] Rational belief-formation, even by boundedly rational toxicologists, assessors, and regulators, is holistic; it is the amalgamation of frequency data with the other (theoretical or observational) grounds of belief that produces the experts', assessors', and regulators' Bayesian probabilities.

The example I have been relying upon to illustrate the divergence of "individual risk" in the Bayesian sense (the regulator's Bayesian probability that some person will die) from "individual risk" in the frequentist sense assumed that the causal models to which the regulator ascribed a nonzero probability were all deterministic. This assumption eased exposition, but was not essential. The divergence still obtains even if some or all of the models are stochastic—as are, for example, some prominent models of carcinogenesis.[343]

A stochastic model posits some irreducible, causal randomness. It posits some exhaustive set of causal features such that individual exposures identical with respect to these features need not all result in, or fail to result in, the toxic effect.[344] Instead, the long-run frequency with which identical exposures result in the effect is greater than 0 and less than 1. This frequency is the stochastic "risk" that the model attaches to various, maximally-specified toxic exposures. Frequencies are therefore integral to stochastic models in a way that is not true of deterministic models. The deterministic threshold model used in the example above says that each individual's genetic makeup and current physiology produces a threshold for her such that exposures above the threshold definitely kill her and exposures at or below the threshold do not. A different, stochastic model might say that each individual's genetic makeup M and current physiology H determine a unique probability function $p(D, M, H)$, such that for any dose D the frequency of fatalities in a large group of individuals consuming just that

---

[342] *See* COVELLO & MERKHOFER, *supra* note 17, at 134-37 (describing these tests).

[343] *See* CRAWFORD-BROWN, *supra* note 28, at 128; RODRICKS, *supra* note 150, at 156, 168.

[344] *See* CRAWFORD-BROWN, *supra* note 28, at 106 (contrasting deterministic and stochastic models of toxicity); HOWSON & URBACH, *supra* note 75, at 7 ("[M]any scientific theories are explicitly probabilistic [such as] Mendel's theory of inheritance. This states the probabilities with which certain combinations of genes occur during reproduction; but, strictly speaking, the theory does not categorically rule out, nor predict, any particular genetic configuration.").

dose and possessing just that genetic makeup and physiology is $p(D, M, H)$.

But even with stochastic models in the picture, the expert's Bayesian probability that some individual will die prematurely as a result of an exposure need not equal the actual, observed, or most likely[345] frequentist risk relative to the canonical dosing class, the maximally specified class, or the class specified with reference to the features of the Subject that the expert knows about. If the expert knows that the Subject is an adult Dane, and believes that adult Danes have a different distribution of genetic makeups and physiologies than the general population, then (clearly) her Bayesian probability of the Subject's death can differ from the actual, observed, or most likely canonical, frequentist risk. Further, the expert might have nonzero degrees of belief in two or more incompatible stochastic models, and within each model in different possible characteristics that the Subject might have. $p_1(D, M, H)$ might be the correct model, and the Subject might actually be characterized by genetic makeup M^* and physiology H^*, but the expert might be unsure whether the function is $p_1(D, M, H)$ or $p_2(D, M, H)$, and what precisely the Subject's genetic makeup and physiology is. Given these uncertainties, the expert's Bayesian probability of the Subject's death certainly need not track the actual frequency relative to the maximally specified class (given by p_1, by the Subject's actual dose, and by M^* and H^*),[346] nor the most likely frequency[347] relative to the maximally specified class.[348] Finally, the actual

---

[345] By "most likely," I mean that frequency to which the expert ascribes the highest Bayesian probability. In the earlier example, *see supra* text accompanying notes 333-40, the expert believed to degree 0.4 in a frequency of 1 in 200,000 and to degree 0.6 in a frequency of 1 in 20,000; thus the "most likely" frequency was 1 in 20,000.

[346] Note that my analysis *does* acknowledge that the expert's beliefs about the stochastic fatality risks to which a Subject may be vulnerable shape the expert's Bayesian probability assessment that the Subject will die. If a given model/parameter combination ascribes a given stochastic risk to the Subject, then that stochastic risk, *discounted by the expert's Bayesian probability that the model/parameter combination truly characterizes the Subject*, is added to other such discounted risks or deterministic outcomes (if deterministic models are also in the picture) to produce the expert's all-things-considered assessment. But, except in the limiting case where the expert is certain of a particular model parameter combination, her Bayesian assessment that the Subject will die doesn't reduce to any particular stochastic probability. On the linkages between frequentist and Bayesian risk, see HOWSON & URBACH, *supra* note 75, at 344-47.

[347] Think of the point this way: concatenating possible models (here, p_1 and p_2) and model parameters (here, different values of M and H) produces a set of possible frequencies relative to maximally specified classes. The most likely frequency is that single frequency, in this set, to which the expert ascribes the highest Bayesian probability. But the expert's Bayesian probability that the Subject will die is a weighted average of *all* the frequencies in the set, as weighted by her Bayesian probability for each.

frequency of premature death among adult Danes is determined by the true model (p_1), while the expert's Bayesian probability of the Subject's death is a function both of that model and of others to which she ascribes a nonzero probability. The observed frequency in a grisly experiment with adult Danes must be integrated with the expert's prior beliefs about models and about the distributions of gene types and physiologies across the adult Danish population. And the frequency relative to the class of adult Danes that the expert believes to be most likely is *not* necessarily the same as her Bayesian probability in the Subject's death, since that latter probability is a weighted average of all possible models, gene types, and physiologies that might characterize an adult Dane (as weighted by the expert's beliefs in these various model/parameter combinations), while the former is not a weighted average but is instead the single most likely frequency that emerges from combining the different possible models with different possible distributions of genes and physiologies across the population of adult Danes.

Part of my argument for the divergence between the regulator's Bayesian probability that some person will die prematurely as a result of a toxic exposure, and that person's "individual risk" in the frequentist sense, has been that regulators may know more about the individuals exposed to toxins than their dosing amounts. This, among other things, drives a wedge between Bayesian risk and the kind of frequentist risk standardly used in current risk assessment practices—namely frequencies relative to canonical dosing classes. The regulator, in my example, knows that the person receiving a one-gram dose is an adult Dane—not merely that he is a person receiving a one-gram dose. It might be objected that regulators lack this information. But certainly some characterizing information, beyond dose amount, is *accessible* to regulators. The age, gender, and ethnic makeup of a community exposed to a toxic waste dump, a group of workers at risk from some occupational toxin, or the consumers currently enjoying some risky food, could be established (to some degree of certainty) by risk assessors evaluating these hazards.[349] Indeed, the influential 1993 Na-

[348] What about *observed* frequencies relative to maximally specified classes? Where the models invoke characteristics such as genetic makeup or health history, experiments holding constant these characteristics will be extremely difficult to perform, and thus observations of frequencies relative to different maximal classes will likely be unavailable. In any event the general Bayesian point about the need to incorporate observations produced by any experiments with prior beliefs will still hold.

[349] Those who argue that risk regulation should be sensitive to considerations of "environmental justice"—specifically, to skews in death or risk across gender, racial, or

tional Academy of Sciences report on risk assessment encourages agencies to characterize individuals in greater detail than is currently the norm:

> EPA, NIH, and other federal agencies should sponsor molecular epidemiologic and other research on the extent of interindividual variability in various factors that affect susceptibility and cancer, on the relationships between variability in each factor and in the health end point, and on the possible correlations between susceptibility and such covariates as age, race, ethnicity, and sex. Results of the research should be used to adjust and refine estimates of risks to individuals (identified, identifiable, or unidentifiable) and estimates of expected incidence in the general population.[350]

A different objection, from within welfare consequentialism and CBA, is that characterizing individuals and individual exposures is costly. Perhaps, given the deliberation costs of fuller characterization, the practice of describing only the distribution of doses that toxic hazards deliver to the population, not the age, gender, ethnic, or other casually relevant features of that population, is optimal.[351] But it is very hard to see why a blanket rule against fuller characterization

---

socioeconomic lines—often suggest, not surprisingly, that risk assessors should characterize exposed populations in gender, racial, or socioeconomic terms. *See, e.g.*, John D. Graham, *Making Sense of Risk: An Agenda for Congress*, *in* RISKS, COSTS, AND LIVES SAVED: GETTING BETTER RESULTS FROM REGULATION 183, 190-91 (Robert W. Hahn, ed., 1996) ("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." (citation omitted)); Robert R. Kuehn, *The Environmental Justice Implications of Quantitative Risk Assessment*, 1996 U. ILL. L. REV. 103, 151 ("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."). Indeed, Executive Order 12,898 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 . . . ."). Exec. Order No. 12,898, § 1-101, 3 C.F.R. 859, 859 (1995). *See generally* Adler, *supra* note 20, at 1423-31 (discussing environmental justice).

[350] NAT'L RESEARCH COUNCIL, *supra* note 17, at 219. For scholarly work advocating greater regulatory attention to individual variability in nondose characteristics relevant to toxicity, see Finkel, *supra* note 86, at 313-20; Jean A. Grassman et al., *Accounting for Variability in Responsiveness in Human Health Risk Assessment*, *in* HUMAN VARIABILITY IN RESPONSE TO CHEMICAL EXPOSURES, *supra* note 86, at 1-26; Hattis, *supra* note 87, at 7-11. Governmental guidelines or proposed guidelines encouraging attention to individual variability are cited in Grassman, *supra*, at 15-16.

[351] *See supra* text accompanying note 304 (discussing deliberation costs).

would indeed be optimal. What of cases where non-dose information is cheaply obtainable or already known by the assessor?

V. BEYOND WELFARISM: FREQUENTIST RISK AND NONWELFARIST VIEWS

Regulatory attention to frequentist risk is problematic not merely for welfarist consequentialism but for other moral views too. In this Part, I consider the leading competitors: nonwelfarist consequentialism, deontological or rights-based views, contractualism, and democratic views.

A. *Safety-Focused Consequentialism*

Welfarist consequentialism, whether in the form of utilitarianism or in a more egalitarian form, presumes that welfare is the "currency" of morality.[352] What makes one outcome better or worse than another is the welfare that different persons enjoy—nothing else. Various moral theorists have objected to this focus on welfare, arguing that moral requirements may depend (at least in part) on something other than human well-being[353]—for example, on the distribution of "primary goods" (Rawls's idea),[354] "resources" (Dworkin's),[355] or the satisfaction of needs.[356]

Skepticism about the moral primacy of welfare might impel the skeptic to abandon consequentialism entirely—for example, to adopt deontology[357] or contractualism.[358] But it need not do so. Skepticism about the moral primacy of welfare can also be accommodated within consequentialism. How? By making the ex post ranking of outcomes depend upon something other than welfare. For example, one can imagine a kind of "safety-focused consequentialism" that makes the ex post ranking hinge on human health and longevity, either exclusively

---

[352] *See* Arneson, *supra* note 255.

[353] *See* Adler & Posner, *Rethinking Cost-Benefit Analysis*, *supra* note 36, at 211-16 (surveying and criticizing "resourcist" views).

[354] *See* JOHN RAWLS, A THEORY OF JUSTICE 90-95 (1971).

[355] *See* Ronald Dworkin, *What is Equality? Part 2: Equality of Resources*, 10 PHIL. & PUB. AFF. 283, 283 (1981).

[356] *See* Thomas Nagel, *Autonomy and Deontology*, *in* CONSEQUENTIALISM AND ITS CRITICS 142, 145-50 (Samuel Scheffler ed., 1988).

[357] *See infra* Part V.B.

[358] *See infra* Part V.C.

or predominantly.[359] For simplicity, I'll confine my discussion to the first kind of safety-focused consequentialism: where the world-ranking depends solely on human health and longevity. On this view, the "primary good" or "resource" of bodily integrity becomes the sole, bedrock moral factor for adjudicating between outcomes. This view, because simplified, is implausible: presumably any plausible non-welfarist consequentialism will incorporate a range of "primary goods" or "resources,"[360] and not merely health and longevity. But my analysis of the simplified view carries over, I believe, to more complicated non-welfarist accounts.

Note that safety-focused consequentialism (the simple kind) maps nicely onto certain decision procedures that, it is sometimes proposed, regulators should employ instead of CBA: namely, "risk-risk" analysis[361] and "QALY-based" analysis.[362] Risk-risk analysis aims to maximize longevity, that is, the cumulative life-span of the entire population; QALY-based analysis takes account of health as well as longevity, by giving greater weight to years that persons live in good health, and less weight to years that they live in poor health. Safety-focused consequentialism, in effect, is the moral view that directly

---

[359] *See, e.g.*, Amartya Sen, *Consequential Evaluation and Practical Reason*, 97 J. PHIL. 477, 487-92 (2000) (emphasizing that consequentialism does not entail welfarism).

[360] *See* RAWLS, *supra* note 354, at 92 (listing "rights and liberties, opportunities and powers, income and wealth" as well as self-respect—"a sense of one's own worth"—as primary goods).

[361] For general discussions of risk-risk analysis, see LESTER B. LAVE, THE STRATEGY OF SOCIAL REGULATION: DECISION FRAMEWORKS FOR POLICY 15-17 (1981); 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 (John D. Graham & Jonathan Baert Wiener eds., 1995); Samuel J. Rascoff & Richard L. Revesz, *The Biases of Risk Tradeoff Analysis: Towards Parity in Environmental and Health-and-Safety Regulation*, 69 U. CHI. L. REV. 1763 (2002); 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).

[362] For overviews of the QALY (quality-adjusted life years) method for measuring health states, see DOUGLAS MCCULLOCH, VALUING HEALTH IN PRACTICE: PRIORITIES, QALYS, AND CHOICE (2003); ERIK NORD, COST-VALUE ANALYSIS IN HEALTH CARE: MAKING SENSE OUT OF QALYS (1999); Paul Dolan, *The Measurement of Health-Related Quality of Life for Use in Resource Allocation Decisions in Health Care*, *in* 1B HANDBOOK OF HEALTH ECONOMICS 1723 (Anthony J. Culyer & Joseph P. Newhouse eds., 2000); Robert Fabian, *The Qualy Approach*, *in* VALUING HEALTH FOR POLICY: AN ECONOMIC APPROACH 118 (George Tolley et al. eds., 1994); Robert M. Kaplan, *Utility Assessment for Estimating Quality-Adjusted Life Years*, *in* VALUING HEALTH CARE: COSTS, BENEFITS, AND EFFECTIVENESS OF PHARMACEUTICALS AND OTHER MEDICAL TECHNOLOGIES 31 (Frank A. Sloan ed., 1995). For a proposal that regulators employ QALYs, see Richard H. Pildes & Cass R. Sunstein, *Reinventing the Regulatory State*, 62 U. CHI. L. REV. 1, 83-85 (1995).

maps onto these regulatory techniques.[363] Relatedly, safety-focused consequentialism is one way to justify statutory provisions that preclude cost-benefit analysis and give greater weight to health and longevity than welfarist views would countenance.[364]

Safety-focused consequentialism has the same outcome-oriented structure as welfare consequentialism. It contains: (1) an "ex post" part, that is, a criterion for ranking possible worlds; and (2) an "ex ante" part, that is, a bridging rule for deriving "oughts" over choices from the ranking of outcomes. It also includes a "supervenience" principle analogous to the welfarist's—namely, two outcomes identical with respect to human longevity and health must be ranked the same.[365] Think of the scheme this way: for each outcome, each person is assigned a QALY number that measures her total lifespan, as adjusted for her health at different points.[366] Then the supervenience principle demands that any two outcomes identical with respect to the QALY numbers for all persons must be accorded the same ex post ranking.

Would "individual risk" in the frequentist sense play a role in regulatory practice grounded on safety-focused consequentialism? I suggest not. First, "individual risk" lacks ex post relevance. Interestingly, it is much easier to show this than to demonstrate ex post irrelevance within welfare consequentialism. *P*'s risk of death, in the frequentist or for that matter the Bayesian sense, doesn't change her longevity or health. As between two outcomes, differing only in that a toxic exposure (or something else) occurs and subjects *P* to a frequentist or Bayesian fatality risk, her QALY number in both outcomes will be the same. Welfarists can't stop here, but must consider (as we did above) whether the risk of death is an ex post setback on non-safety grounds: because it is dispreferred; because it gives rise to fear; because it disrupts affiliational goods.[367] Safety consequentialists need

---

[363] Safety-focused consequentialism has distributively insensitive and distributively sensitive variants, just as welfare consequentialism does. *See supra* text accompanying notes 258-62. The distributively sensitive variants of safety consequentialism would require a kind of risk-risk analysis or QALY-based analysis which attends to the distribution of longevity or health-adjusted longevity, rather than simply maximizing longevity or health-adjusted longevity.

[364] *See* Adler, *supra* note 20, at 1391-92, 1414-15 (describing statutes that clearly or arguably give priority to safety and preclude CBA).

[365] *See supra* text accompanying notes 255-57.

[366] For example, if she lives 60 years, 40 of those in excellent health, and 20 during which her health is midway between excellent health and a health state no better than death, her QALY measure is $(40 \times 1) + (20 \times (1/2)) = 50$.

[367] *See supra* text accompanying notes 280-90.

not consider these more difficult questions; the obvious point that the risk to some person of a serious or fatal accident or illness doesn't entail her actual illness, injury, or premature death suffices for the ex post analysis.

Second, "individual risk" lacks ex ante relevance within safety consequentialism. Here, the analysis *is* quite parallel to the welfarist analysis.[368] Various bridging rules are possible: some will be probabilistic, others will eschew reference to probabilities. Expected utility maximization is the most plausible example of a probabilistic bridging rule for cognitively unbounded regulators. The probabilities, here, are Bayesian probabilities over complete outcomes, not frequentist probabilities of individuals dying. Boundedly rational regulators would employ a more tractable decision procedure: for example, minimizing the number of premature deaths, maximizing the population's longevity, or maximizing the population's health-adjusted longevity. The "individual risk" of death incurred by different persons is not directly relevant to the decision procedures just described, but it is possible (I suppose) to imagine a safety-consequentialist decision procedure where it would be.[369] One such possible procedure instructs the regulator to determine the "individual risk" of death, illness, or injury that members of different groups face as a result of hazards, and therewith their expected loss of longevity or health-adjusted longevity. Note, however, that the risks here would be third-party Bayesian risks: the regulator's degrees of belief that the various individuals will die prematurely or suffer bad health or bodily harm.

B. *Deontological Views*

Deontological views deny that the moral ordering of a set of possible actions necessarily derives from the outcomes those actions might produce, as ranked in some impartial (evaluator-neutral) way. Deontologists propose that morality includes certain "agent-relative constraints"—prohibitory rules that, in identifying the actions that are impermissible for a given actor, point to some nexus between those

---

[368] *See supra* Part III.C.

[369] *Cf.* Adam M. Finkel, *Comparing Risks Thoughtfully*, 7 RISK: HEALTH, SAFETY & ENVIRONMENT 325, 342-44 (1996) (enumerating possible dimensions for evaluating hazards, including "population-based measures" and "individual risk measures"); Graham & Wiener, *supra* note 361, at 30-31 (arguing that risk-risk analysis should be sensitive to both "population risk" and "individual risk").

actions and that particular person.[370] A classic example is the rule "do not kill." This rule, which purports to apply even to perfectly informed actors, does *not* instruct the actor to choose the action leading to the outcome in which the fewest killings occur, in which the fewest deaths occur, or in which aggregate longevity or health-adjusted longevity is greatest. Rather, he is instructed not to perform an action that would constitute a killing *by him*, even if his not killing has as its consequence more total killings.[371] Because deontological rules are "agent relative" in this sense, they do not correspond to any impartial, consequentialist ranking of outcomes. The deontological injunction "do not kill" is structurally distinct, in a deep way, from the consequentialist injunction "minimize the number of killings."

What would a deontological practice of risk assessment look like? At the threshold, one needs to confront a generic problem for the view that government should enforce some deontological constraints. Agent *A* is about to breach some constraint, harming or otherwise affecting the victim *P*. Why should government intervene to prevent the breach, absent some consequential justification for doing so? And why should it punish *A* after the fact (as proposed by deontological theorists of criminal law),[372] or coerce him to compensate *P* (as proposed by deontological theorists of tort law),[373] if prevention fails? After all, the constraint is agent relative: it is worse from *A*'s perspective that he act in a certain way, not worse from ours. Does the deontologist want to say that it *is* deontologically worse from our perspective—

---

[370] On the distinction between deontological and consequentialist moral views, with particular reference to deontological "constraints" (my focus here), see KAGAN, *supra* note 258, at 70-78; ROBERT NOZICK, ANARCHY, STATE, AND UTOPIA 26-53 (1974); SAMUEL SCHEFFLER, THE REJECTION OF CONSEQENTIALISM 80-114 (rev. ed. 1994); Nagel, *supra* note 356, at 156-72. My understanding of consequentialism as a family of moral views that necessarily eschew agent-relative moral rules is quite standard. *See, e.g.*, Douglas W. Portmore, *Can an Act-Consequentialist Theory Be Agent Relative?*, 38 AM. PHIL. Q. 363, 363-64 (2001) ("A theory is agent neutral if it gives every agent the exact same set of aims and agent relative otherwise. . . . [M]any philosophers would deny that a consequentialist theory can be agent relative." (footnotes and paragraph structure omitted)). Consequentialism might be defined more inclusively, *see id.* at 364, but it strikes me that this is less useful than the more traditional and narrower definition keyed to the salient distinction between agent-neutral and agent-relative rules.

[371] *See, e.g.*, Portmore, *supra* note 370, at 364 ("[A] theory would have to be agent relative in order to accommodate a constraint against intentionally killing the innocent . . . since such a constraint prohibits the commission of murder even for the sake of preventing numerous others from committing comparable murders.").

[372] *See* Michael S. Moore, *The Moral Worth of Retribution*, *in* RESPONSIBILITY, CHARACTER, AND THE EMOTIONS: NEW ESSAYS IN MORAL PSYCHOLOGY 179, 179 (Ferdinand Schoeman ed., 1987).

[373] *See* JULES L. COLEMAN, RISKS AND WRONGS 303-60 (1992).

as fellow citizens of *A* and *P*—not to prevent *A*'s breach of the constraint or, failing that, to punish *A* or coerce compensation from him after the breach occurs? If so, how does this square with the general deontological aversion to affirmative duties, duties of intervention?[374] Again, this is a very general problem for deontological accounts of law and governmental practice, and so I will place it to one side. I will assume that *if* there are deontological prohibitions against killing or risk imposition, then these constraints are properly enforced by regulators in some way and therefore influence their evaluation of regulatory options. But it bears remembering that the assumption is questionable.

Assuming this very general problem for a deontological account of governmental practice can be overcome, what are the implications of deontology for risk assessment? There are subtle issues here, which need to be carefully unpacked. One needs to distinguish between: (1) a deontological constraint against *killing*, which governmental officials operating under uncertainty attempt to enforce; and (2) a deontological constraint against *risk imposition* (even risk impositions that do not cause deaths), which governmental officials attempt to enforce.

As for the first possibility, the set of deontological constraints, if non-empty, surely includes a prohibition against killing.[375] No deontologist suggests otherwise. Providing a precise characterization of this constraint is tricky. "Killing" is not the same as "causing death." One needs to distinguish, perhaps, between active and passive causation (since merely failing to prevent death may not breach the constraint); between direct and indirect causation (since actions that cause death only when conjoined with certain "intervening" causes, such as the victim's own voluntary action, may not breach the constraint); between death-causing actions motivated by a sufficiently wrongful mental state (perhaps purpose or knowledge) and death-causing actions motivated by an innocent mental state; and so on.[376] But all of this could, in principle, be incorporated within risk assess-

---

[374] *See* KAGAN, *supra* note 258, at 94-100.

[375] *See, e.g.*, *id.* at 71 ("Intuitively, at least, most of us have little doubt that it is morally forbidden to chop up an innocent person, even if this is the only way to save five other innocent people from death."); Nagel, *supra* note 356, at 157 (listing deontological constraints supported by "[c]ommon moral intuition," including "rights not to be killed, injured, imprisoned, threatened, tortured, coerced, robbed").

[376] *See, e.g.*, KAGAN, *supra* note 258, at 84-105 (discussing content of constraint against harming); 2 F.M. KAMM, MORALITY, MORTALITY 17-120 (1996) (discussing distinction between killing and letting die); THOMSON, *supra* note 282, at 205-48 (discussing content of moral rights not to suffer bodily intrusion or be harmed).

ment. One might ask, for example, whether a "killing" is likely to result if a given toxic source is left unregulated—where that determination will be sensitive, if necessary, to the directness of the causal route between the toxin and the victims, the mens rea of the persons who control the source, the responsibility of the victims for their own deaths (did they "come to the hazard" despite warnings?), and other such considerations that may distinguish killings from those deaths that, albeit unfortunate, wouldn't breach deontological rules.

Risk assessment practices grounded on the no-killing constraint would need to take into account the limited information of regulators.[377] Regulators aren't omniscient; they can't simply be instructed to "prevent killings," since they may well be unsure whether some toxin or other hazard, if unregulated, would indeed produce a killing. Instead, I suggest, the generic decision rule for nonomniscient deontological regulators would be something like the following: "*ceteris paribus*, intervene to prevent any hazardous activity that you believe, to degree p or higher, would constitute a killing." In other words, deontologically grounded risk regulation would incorporate a probability threshold p (perhaps zero) for one or more killings; once a hazard reached that threshold, the regulator would intervene, absent sufficient conflicting considerations, to enforce the no-killing rule. Conflicting considerations could be the cost of enforcement or a liberty interest on the part of the actor.[378] Any full-blown deontological account of risk regulation would need to specify these considerations, since the notion of an absolute governmental obligation to prevent killings, even where prevention has huge costs, is implausible. More plausible is that the risk regulator has a prima facie obligation to intervene when she believes, to some degree p or higher, that a killing would otherwise occur.

The crucial point to see here is that the probability threshold p invokes a Bayesian, not frequentist, probability. The decision procedure for enforcing deontological rules, such as the no-killing rule, will appeal—I am suggesting—to the regulator's knowledge. Intuitively, if

---

[377] Moral theorists haven't focused much on the problem of choice under uncertainty. *See* TED LOCKHART, MORAL UNCERTAINTY AND ITS CONSEQUENCES 16-21 (2000). But moral theories, including deontological theories, are normative (act-guiding) and therefore should contain instructions for actors as they (virtually always) find themselves, namely in an epistemically limited state. Whether "wrongness" and "rightness" itself is subjective or objective is a different issue. *Cf.* GIBBARD, *supra* note 295, at 42-43 (arguing that moral wrongness is subjective).

[378] *See* KAGAN, *supra* note 258, at 78-82 (explaining that deontological constraints might be defeasible by countervailing considerations).

the regulator is certain that a hazardous activity would constitute a killing, then (*ceteris paribus*) she should intervene to stop the activity. Given the divergence between frequentist and Bayesian probabilities,[379] there can be cases in which the regulator, even an omniscient regulator, is certain that a killing will result, yet the frequentist probability of the hazardous activity constituting a killing is low. Intuitively, the regulator in such cases has a prima facie obligation to intervene (assuming the more basic premise that deontological rules are enforceable by government). Relatedly, where the regulator is nonomnisicent and uncertain, the relevant kind of probability for a deontological decision procedure enforcing the no-killing constraint is a Bayesian measure of her belief that killings may ensue, not a relative frequency.

What about the deontological rule against risk imposition? Some deontologists, most visibly the philosopher David McCarthy, argue that *A* infringes *P*'s moral rights by imposing a risk of death (or other bodily harm) on *P*, even if *A*'s action doesn't actually cause *P*'s death or bodily harm.[380] Assuming that deontological rules exist and are properly enforced by regulators, might a no-risking rule account for regulatory attention to frequentist risk, even if the no-killing rule doesn't?

McCarthy's view is highly controversial, even among deontologists. While no deontologist worth her salt would deny that a killing (suitably characterized) violates the victim's rights, many would deny that a risking (without more) does.[381] Consider, specifically, the imposition of a frequentist death risk, the central concern of risk assessors. Where *A* acts and thereby imposes a high frequentist death risk on *P*, relative to some reference class, it does not follow that *P* suffers death, injury, fear,[382] or an invasion of her property[383] or body,[384] as a result of

---

[379] *See supra* Part IV.

[380] David McCarthy, *Rights, Explanation, and Risks*, 107 ETHICS 205, 205-06 (1997).

[381] *See id.* at 205-06 (stating that "many writers have been pessimistic about whether the morality of risk imposition can be accommodated in a plausible theory of rights" and citing Robert Nozick, Peter Railton, and Dennis McKerlie); *id.* at 208 (citing Judith Thomson, Dennis McKerlie, and Samuel Scheffler as denying the "Risk Thesis," i.e., that "we have the right that other people not impose risks of harm upon us").

[382] On rights not to be frightened, see RICHARD A. EPSTEIN, A THEORY OF STRICT LIABILITY 29-32 (1980). *But see* THOMSON, *supra* note 282, at 250-59 (arguing for moral right against non-belief-mediated distress, not fear or other kinds of belief-mediated distress).

[383] *See, e.g.*, Railton, *supra* note 273, at 90-92 (summarizing Lockean view of moral rights to property).

this action. For example, if *A* builds a bomb on his own property, which fails to explode and which *P* fails to learn about, then *P* may well have suffered a very high frequentist risk (relative to the class of all persons who are as proximate as *P*, for a similar duration, to similar explosives), but there has been no physical or emotional harm to *P*; setback to her property interests in land, goods, or person; or manipulation[385] of her—the typical disjunctive preconditions for rights violations. And in the case of toxic exposures, which admittedly are physically invasive, the rights-theorist could not justify her attention to frequentist risk by adducing a concern with physical invasion—since *every* exposure, risky or not, intrudes on the subject's body.

One line of argument in defense of a no-risking constraint (although McCarthy doesn't pursue this) is that risk imposition amounts to a dignitary harm. Risking is not (as such) tangibly or emotionally harmful, invasive, manipulative, or a breach of property rights, but it *does* amount to a kind of insult to the victim. The homicidal neighbor who builds a bomb so as to kill *P*, thereby risking *P*'s death, has wronged *P* in virtue of the contemptuousness of this action—or so the intuition goes. Certainly there are plausible moral theories in the neighborhood (as it were) of this intuition. Jean Hampton has proposed a general "expressive" theory of deontological constraints;[386] Martha Nussbaum, as noted above, includes affiliational goods (subsuming the good of respectful treatment) on her list of welfare goods;[387] and Peter Railton has, specifically, sketched a dignitary defense of a no-risking constraint:

> Would it make any difference to the wrongness of my playing Russian roulette on my sleeping roommate that he is someone who constitutionally feels no fear? . . . Kantians, at least, would presumably deny this. They have held that the wrong done to an individual by (for example) fraud or coercion is not just a matter of the discomfort such an act, if known, would cause him . . . . Rather, Kantians have argued that such acts fail to show adequate respect for the individual as an autonomous being, discomfort apart. Therefore, [those] who would employ a Kantian interpretation of the notion of respect for persons and rights must

---

[384] *See, e.g.*, THOMSON, *supra* note 282, at 205-26 (arguing for moral right against bodily intrusion).

[385] *See, e.g.*, Nagel, *supra* note 356, at 157 (stating that deontological constraints include "the special obligations created by promises and agreements [and] the restrictions against lying and betrayal").

[386] *See* Jean Hampton, *Correcting Harms Versus Righting Wrongs: The Goal of Retribution*, 39 UCLA L. REV. 1659, 1671-85 (1992) (developing an expressive account of the distinction between wrongdoing and mere harm-doing).

[387] *See supra* text accompanying notes 288-90.

> affirm that what makes endangering [risk-imposing] acts wrong is not merely the uncomfortable psychological states they may cause in others.[388]

For the sake of argument, I will accept the premise that *A* harms and wrongs *P* by "expressing," in some sense, the proposition that *P* is a lesser being, not worthy of full respect, not a full person, and so on—even if *P* is unaware of *A*'s communication (and therefore suffers no emotional harm).[389] But how does this premise explain a constraint against risking? Doesn't it explain a constraint against *insults*: namely, those actions with a particular linguistic content, either in virtue of the actor's subjective intentions, or in virtue of the linguistic conventions existing in his society?[390] One answer, perhaps, is that risk imposition that is not linguistically insulting can still be "insulting"—can still constitute a wrongful "expression"—in a broader sense. For this explanation to work, we need a broad view of meaning that ascribes propositional content to actions independent of actors' intentions or linguistic rules.[391] Such a view is problematic on its own terms. And even bracketing that point, what specifically about *A*'s imposition of a risk on *P* "expresses" *P*'s inferiority where *A* neither means to say that *P* is inferior, nor conventionally says that *P* is inferior in the way that an inscription or verbal performance would? The meaning of *A*'s imposition, if not a matter of his intentions or of linguistic conventions, would have to arise from the responses of suitably characterized onlookers.[392] Anyone characterized in the right way would see *A*'s action as contemptuous of *P*. But wouldn't that, in turn, depend on whether this onlooker sees the action as a risky one? If the action imposes a high frequentist risk, relative to the canonical dosing class, but

---

[388] Railton, *supra* note 273, at 106-07.

[389] *But see* THOMSON, *supra* note 282, at 210 ("[W]hile lack of respect for one's claims can insult, it cannot be the insult itself that makes an act be an infringement of a claim.").

[390] *See* Matthew D. Adler, *Expressive Theories of Law: A Skeptical Overview*, 148 U. PA. L. REV. 1363, 1387-88, 1393-96 (2000) (summarizing speaker's-meaning and sentence-meaning accounts of linguistic meaning).

[391] *See* Elizabeth S. Anderson & Richard H. Pildes, *Expressive Theories of Law: A General Restatement*, 148 U. PA. L. REV. 1503, 1523-27 (2000) (arguing that governmental action can be expressive even if it lacks a linguistic meaning); Matthew D. Adler, *Linguistic Meaning, Nonlinguistic "Expression," and the Multiple Variants of Expressivism: A Reply to Professors Anderson and Pildes*, 148 U. PA. L. REV. 1577 (2000) (analyzing the Anderson and Pildes account).

[392] *See* Anderson & Pildes, *supra* note 391, at 1525 ("Expressive meanings are socially constructed. . . . Although these meanings do not actually have to be recognized by the community, they have to be *recognizable* by it, if people were to exercise enough interpretive self-scrutiny.").

the onlooker is better informed and believes that the risk is lower, why would she read contempt into the action (where the riskiness, rather than *P*'s subjective contemptuousness or language rules, licenses the reading)? In short, here as elsewhere, the relevant probabilities would seem to be Bayesian—in this case, a measure of the onlooker's beliefs—not frequentist.

McCarthy does not pursue the dignitary idea. He proposes a generic constraint against risk imposition, expressive or not, apparently on the grounds that such a constraint explains our intuitions in Russian roulette and other cases where no one has been physically harmed and yet, we intuit, a wrongdoing has occurred.[393] Thus he avoids my objection that the risks relevant for the expressivist would be Bayesian, not frequentist. But other problems await McCarthy. First, any account which claims that the imposition of frequentist risk is a rights infringement must identify the relevant reference class (or groups of reference classes).[394] *A*'s action will have a high risk of causing *P*'s death relative to some descriptions, a low risk relative to others, and perhaps a zero risk relative to yet others. The description identified as relevant by the deontological account will either be conventional (it will take into consideration social facts) or natural. But there is only one nonarbitrary natural class—namely the class identifying all of the causally relevant features of *A* and *P*. Incorporating *that* reference class in the deontological account will have the upshot that, where the applicable causal laws are deterministic, *A*'s action imposes a nonzero risk of death on *P* if and only if the action will actually cause *P*'s death—so that the constraint reduces to a no-killing constraint.[395] What about variants of the no-risking constraint that define risk relative to conventional classes?[396] One wonders why risk, thus defined, would be morally problematic apart from its expressive significance. In any event, the most straightforward way to conceptualize conventional reference classes—namely, as those classes that are socially salient—leads us back to Bayesian probabilities. *A*'s action will have a high conventional risk of causing *P*'s death if the "ordinary" or "reasonable" member of *P* and *A*'s society would view the action as falling in a class of actions that frequently cause death. But we can then

---

[393] *See* McCarthy, *supra* note 380, at 208.

[394] *See* Stephen Perry, Imposing Risk 55-58 (unpublished paper, on file with author) (criticizing McCarthy and arguing that the imposition of frequentist risk is not a rights-infringement, given the reference-class relativity of frequentist risk).

[395] *See id.* at 57-58.

[396] *Cf. id.* at 59-67 (considering "social facts" version of McCarthy's view).

imagine cases in which the "ordinary observer" will ascribe a low Bayesian probability to *P*'s death, even though its conventional frequency (as just defined) is high. Why, in such a case, would we care morally about the conventional frequency? Why look partly but not completely to the "ordinary observer"—to determine what reference class to apply to *A*'s action, but not to determine what the observer believes about the chances of that action killing *P*?

A different problem for McCarthy is the threshold problem. Whatever the reference class, a constraint against all actions that impose some nonzero frequentist risk of death relative to that class will presumably be sweeping in its coverage.[397] For example, I impose a death risk on my neighbor (relative to the conventionally salient class) by turning on my light switch, driving my car down the driveway, or starting the stove. Events in the classes of all light-switch-turnings, car-down-driveway-drivings, or stove-startings result in disastrous fires, crashes, or explosions with a nonzero frequency. One route around this problem is to restrict the no-risking constraint to the imposition of *high* risks. But McCarthy rejects that alternative for the following reasons: any particular threshold would be arbitrary; purposeful riskings below the threshold would be permissible, which is counterintuitive; imposing below-threshold risks on many people would be unconstrained while imposing above-threshold risks on a few would be constrained, again counterintuitive; and the problem of the pervasiveness of risking can be handled another way.[398] McCarthy's solution to that problem is complex but, roughly, involves the idea that *A* can *permissibly infringe P*'s right not to be risked if (1) the benefits of risk imposition to *A* exceed the costs to *P*, (2) it is difficult to secure *P*'s consent to the risk imposition, and (3) *A* compensates *P* for the risk.[399] This solution would suggest that enforcement of the no-risking constraint—if such a constraint does exist, and government is morally obliged to enforce it, both doubtful points—is a matter for the tort system, not regulatory agencies.[400] Imagine that a toxic exposure imposes a 1 in 10,000 frequentist death risk on the maximally exposed

---

[397] *See* THOMSON, *supra* note 282, at 244-45; Railton, *supra* note 273, at 107-08.

[398] *See* McCarthy, *supra* note 380, at 212-14.

[399] *See id.* at 210-12, 215-16, 218-19; *see also* David McCarthy, *Liability and Risk*, 25 PHIL. & PUB. AFF. 238, 259 (1996) (arguing that "if X performs an action that she knows or ought to know will impose a risk of harm to Y, then, roughly speaking, . . . X is liable to Y for the risk").

[400] *See also* Christopher H. Schroeder, *Corrective Justice and Liability for Increasing Risks*, 37 UCLA L. REV. 439, 473-77 (1990) (arguing that risk imposition should be tortious if the administrative costs of adjudicating risk imposition claims are low).

individual, relative to the canonical dosing class, and that this class is (somehow) the relevant one for McCarthy's purposes. Then (on McCarthy's account) this risk is not grounds for banning the exposure, even prima facie grounds, if the exposure passes a cost-benefit analysis. Rather, it would warrant a money judgment against the person who releases the toxin, in favor of the risked individual, to compensate him for bearing the risk.[401]

C. *Contractualist Views*

The contractualist insists on *unanimous hypothetical agreement.* It must be the case that each and every individual, suitably informed, would agree, hypothetically, to the policy choice or to a set of underlying principles authorizing that choice.[402]

Naive contractualism says this: where a governmental official is choosing among options $\{A_i\}$ at time T, she should choose the option A^* such that everyone in the population, given her actual beliefs at time T, would agree to A^* as opposed to the other options in $\{A_i\}$.[403] In effect, the official's choice must be ex ante Pareto-superior to her other possible choices. This view would routinely create moral dilemmas for governmental officials (since, for many choice situations, no option is Pareto-superior to the others given actual citizen beliefs) and thus contractualist theorists have typically proposed other, more sophisticated variants of contractualism. One variant was offered by the economists Nicholas Kaldor and John Hicks (in the days when welfare economists were deeply uncomfortable with interpersonal comparisons and thus advanced views that were closer to contractualism than consequentialism in structure). Kaldor and Hicks's proposal, of course, was that governmental choices are required if poten-

---

[401] It might be argued that McCarthy's account warrants a hybrid institutional structure: (1) regulatory prohibition of frequentist risks that aren't permissible (because the benefits don't outweigh the costs or consent ought to have been secured), and (2) tort liability for permissible risks. If so, frequentist risk *would* function as a decisional trigger for regulatory consideration of prohibition, but might not have any additional regulatory role—since the probabilities relevant for cost-benefit analysis, at least the kind of CBA I discuss above, are Bayesian not frequentist. *See supra* text accompanying notes 306-14.

[402] *See* KAGAN, *supra* note 258, at 240-56 (surveying contractualist views).

[403] *Cf.* Adler, *supra* note 306, at 1272-74 (arguing that governmental officials do not have consent-based moral grounds for choosing the "ex ante efficient" option, i.e., the option that maximizes each individual's subjective expected utility given her actual beliefs at the time of choice).

tially Pareto-superior.[404] What this means in the case at hand (roughly) is that the official should choose A^* just in case there exists a scheme of transfer payments such that everyone in the population, given her actual beliefs at T plus the belief that the scheme would be implemented, would agree to A^* over the other options in $\{A_i\}$.

A different route around the difficulties of naive contractualism, suggested by Rawls's[405] and Scanlon's[406] work, is to change the informational or motivational structure of the citizens whose hypothetical consent to governmental choices is at issue.[407] One might stipulate that the official should choose A^* over the other options in $\{A_i\}$ just in case every citizen operating under some kind of "veil of ignorance" and advancing her own interests would choose A^*; or, alternatively, just in case every citizen basing her choice on her actual beliefs but choosing less selfishly would agree to A^*.

All the variants of contractualism I have thus far described, both naive contractualism and more sophisticated versions, are *act* based.[408] One imagines citizens, with actual or counterfactual information, and with actual or counterfactual motivational sets, agreeing or not to particular governmental choices (either the actual choices at hand, or those choices paired with transfer payment schemes). What sort of regulatory decision procedure would implement these *act-contractualist* views? Presumably the regulator should attempt to determine how various members of the population, characterized in greater or lesser detail, would evaluate the choices at hand if the choices were suitably modified, the individuals were suitably motivated and informed, and so on. That determination, in turn, partly depends on what the individuals *believe* about the choices. Choices are a product of beliefs and preferences[409] (be they self-interested preferences or morally motivated preferences). So "individual risk" might be part of the act-contractualist decision procedure. But it would be "individual risk" in

---

[404] *See, e.g.*, Adler & Posner, *Rethinking Cost-Benefit Analysis*, *supra* note 36, at 190-91 (explaining the Kaldor-Hicks criterion).

[405] *See* RAWLS, *supra* note 354.

[406] *See* T.M. SCANLON, WHAT WE OWE TO EACH OTHER (1998).

[407] *See* KAGAN, *supra* note 258, at 243-45 (explaining that the rationality, knowledge, and motivation of the contractors might be specified in different ways); Adler, *supra* note 35, at 272-73 (summarizing Rawls's and Scanlon's specifications of the contracting scenario).

[408] *See* KAGAN, *supra* note 258, at 242 (distinguishing between act and rule contractualism).

[409] *See* DANIEL M. HAUSMAN & MICHAEL S. MCPHERSON, ECONOMIC ANALYSIS & MORAL PHILOSOPHY 38-41 (1996).

the first-party Bayesian sense: probability numbers measuring the beliefs of the individuals in the contracting scenario, not numbers measuring the beliefs of the regulator, or frequentist risks.

Perhaps a concrete example might help. Imagine that some natural contaminant is poisoning a community's drinking water. Ten thousand people live, and will remain for their lifetimes, in the community. An agency official is determining whether to spend $500,000 to correct the contamination affecting this cohort, raised through a special levy of $50 per head. If the correct moral view is welfare consequentialism, and the correct decision procedure within this view is cost-benefit analysis, then the official's decision depends on her beliefs about the toxicity of the contaminant. Third-party, rather than first-party, Bayesian risk comes into play.[410] For example, if the official after consulting a risk assessor believes to degree 1 in 1 million that any individual drinking water from the source for his lifetime will die as a result, while each individual believes to degree 1 in 10,000 that he will die, it is the 1 in 1 million Bayesian risk, *not* the 1 in 10,000 Bayesian risk, that should figure in the CBA analysis of the cleanup.[411] Presumably it won't pass the CBA test.

By contrast, an act-contractualist view would instruct the official to think about a hypothetical referendum.[412] Would the citizenry, hypothetically, agree to the cleanup? The official might answer that question by surveying a sample of the community, providing them the kind of information stipulated by the underlying act-contractualist view,

---

[410] *See supra* note 306.

[411] Third-party and first-party Bayesian risks surely can differ. *See, e.g.*, Ted Gayer et al., *Private Values of Risk Tradeoffs at Superfund Sites: Housing Market Evidence on Learning about Risk*, 82 REV. ECON. & STAT. 439 (2000) (finding that individuals' perceived probabilities of cancer from Superfund sites are changed by EPA's release of information from its remedial investigations of the sites); Mary Riddel et al., *Environmental Risk and Uncertainty: Insights from Yucca Mountain*, 43 J. REGIONAL SCI. 435, 437 (2003) (finding that the perceived risk of an accident during the transportation of high-level radioactive waste, among residents of Southern Nevada, is generally much higher than the Department of Energy's risk estimate). *See generally* HASTIE & DAWES, *supra* note 309 (discussing lay probability mistakes).

[412] *Cf.* Kevin J. Boyle & John C. Bergstrom, *Doubt, Doubts, and Doubters: The Genesis of a New Research Agenda?*, *in* VALUING ENVIRONMENTAL PREFERENCES: THEORY AND PRACTICE OF THE CONTINGENT VALUATION METHOD IN THE US, EU, AND DEVELOPING COUNTRIES 183, 195-96 (Ian J. Bateman & Kenneth J. Willis eds., 1999) (discussing "referendum format" for contingent-valuation surveys); Alan Randall & Warren Kriesel, *Evaluating National Policy Proposals by Contingent Valuation*, *in* ECONOMIC VALUATION OF NATURAL RESOURCES: ISSUES, THEORY, AND APPLICATIONS 153 (Rebecca L. Johnson & Gary V. Johnson eds., 1990) (discussing design of contingent valuation survey to elicit citizen valuations of national environmental policies).

and inducing the requisite motivation.[413] Or, if the view stipulates that the contractors are rational, the official might think about whether the citizens would rationally be willing to pay for the cleanup, given their probability beliefs and the possible costs and benefits. *That* inquiry would focus the official on the citizens' Bayesian probabilities—what I am calling "first-party" Bayesian probabilities. If each person in the population believes to degree 1 in 10,000 that he will die as a result of the contamination, and is willing to pay $600 to avoid a 1 in 10,000 risk of death, then the cleanup passes this act-contractualist test[414] even though it fails a CBA test. Further, and quite obviously, first-party Bayesian probabilities are not the same as frequentist probabilities. A contractor might believe to degree 1 in 10,000 that he will die from the contamination even if his frequentist risk, relative both to the canonical dosing class and to a more finely specified class incorporating the features that the official knows about him, or that he knows about himself, is much lower.

So much for act-contractualism. I said earlier that the work of Rawls and Scanlon suggests a modification to naive contractualism, namely placing citizens under a veil of ignorance or assuming them to be other regarding (to some degree) rather than purely self interested. The actual theories advanced by Rawls and Scanlon do incorporate these informational or motivational assumptions, but they also depart from naive contractualism in another way. Rawls and Scanlon are *rule-contractualists*.[415] Rawls does not imagine a new "behind the veil" bargaining session for each governmental choice; rather, Rawlsian contractors agree once-and-for-all to certain basic principles that will regulate the structure of government and society.[416] Similarly for Scanlon: his account uses the contractualist device to identify moral principles (specifically, those that no person reasonably could reject) and then evaluates choices by applying the winning set of principles.[417]

My argument above that contractualist risk assessment would depend upon Bayesian, not frequentist, probabilities does not work for rule-contractualists. For even though it is an individual's beliefs

---

[413] *Cf.* VALUING ENVIRONMENTAL PREFERENCES, *supra* note 412 (generally discussing the use of contingent-valuation surveys to determine individuals' monetary valuations of nonmarket goods, for purposes of cost-benefit analysis).

[414] On the premise that contractors are appropriately self-interested.

[415] *See* KAGAN, *supra* note 258, at 242 (noting that philosophers commonly think of contractualism in rule-contractualist terms).

[416] *See* RAWLS, *supra* note 354, at 60-65.

[417] *See* SCANLON, *supra* note 406, at 197-202.

(hence Bayesian probabilities) and preferences that rationalize and explain her choices, it is possible that contractors would choose rules that, in turn, somehow incorporate frequentist probabilities. For example, it is possible that *P*, given his Bayesian probabilities about the results of different possible rule choices, would choose a moral rule incorporating a prohibition on risk imposition in the frequentist sense. Or, he might choose a rule that looked to the "fair" distribution of frequentist risk, not Bayesian risk, across the population.[418]

But would such a possibility be realized? This is much too large a question for me to address here, because it entails both a specification of the correct contracting scenario, and an analysis of what rules would be chosen in that scenario. Rawls does not address the morality of risk imposition—his focus in *A Theory of Justice* is the "basic structure" of society, not moral rules governing individuals[419]—and Scanlon addresses it only briefly. Since Rawlsian contractors are self interested—guided, behind the veil, by their own welfare[420]—a plausible conjecture is that the rules of individual morality emergent from a Rawlsian contract would not regulate frequentist risk. Frequentist risk is welfare-irrelevant, or so I have tried to show.[421]

Scanlonian contractors are not constrained to focus on welfare. For example, Scanlon notes, the pretheoretical "unfairness" of a possible rule might motivate a contractor to reject the rule even though the rule does not harm him.[422] Specifically, a contractor might see a rule permitting the imposition of frequentist risks on him as somehow unfair, albeit welfare neutral. This makes Scanlon's view murkier than Rawls's and extremely hard to apply.[423] Scanlon's very brief comments on the issue of risk suggest skepticism about a broad rule against risk imposition—he notes that "the cost of avoiding all behavior that involves risk of harm would be unacceptable"[424]—and willingness to entertain narrower rules, precluding certain kinds of risks of serious harms.[425] Whether Scanlon means, here, to suggest that contractors

---

[418] *Cf.* Adler, *supra* note 20, at 1423-31 (discussing view, within risk-regulation scholarship, that distributive justice concerns the distribution of "individual risk").

[419] *See* RAWLS, *supra* note 354, at 60-65, 109-18.

[420] *See id.* at 136-50.

[421] *See supra* notes 272-90 and accompanying text.

[422] *See* SCANLON, *supra* note 406, at 191-97, 213-18.

[423] *See generally* Oberdiek, *supra* note 273, at 91-193 (employing framework similar to Scanlon's to evaluate morality of risk imposition).

[424] SCANLON, *supra* note 406, at 209.

[425] *See id.* (noting that "it is intuitively obvious that the likelihood that a form of behavior will lead to harm is an important factor in determining its permissibility").

might converge on a rule prohibiting frequentist (rather than Bayesian) risks of serious harms is unclear.

D. *Democratic Views*

All plausible moral views, certainly consequentialist ones, recognize the benefits of democratic legal structures, including popularly elected legislative bodies and some version of legislative supremacy vis-a-vis regulatory agencies. Utilitarians, for example, can certainly approve the framing of a democratic constitution, given the predictable benefits of these structures as opposed to monarchy or other alternatives with respect to overall welfare. But there is no *bedrock* role for democratic responsiveness within utilitarianism, the other consequentialist views considered to this point, or classic deontology.[426] By contrast, a foundationally democratic account of risk regulation would give regulators an independent (perhaps overriding) moral reason to respond to citizen views[427]—a reason distinct from the welfare costs of nonresponsiveness.[428] Such an account could, of course, specify "democratic responsiveness" in a multitude of ways. Must citizens be public-regarding?[429] Is the responsiveness of administrative agencies appropriately direct, or mediated by statutes? Does the view care about responsiveness to majorities or supermajorities?[430] How are Arrow problems resolved?[431]

---

[426] *See generally* Adler, *supra* note 35, at 267-88 (arguing against "proceduralist" theories of regulation, which accord intrinsic moral significance to certain democratic, governmental procedures).

[427] Under the rubric of "foundationally democratic" accounts of governmental choice (including, presumably, risk regulation), I include recent scholarship on "deliberative democracy" and related scholarship on "civic republicanism." *See* DELIBERATIVE DEMOCRACY: ESSAYS ON REASON AND POLITICS (James Bohman & William Rehg eds., 1997); AMY GUTMANN & DENNIS THOMPSON, DEMOCRACY AND DISAGREEMENT (1996); Steven G. Gey, *The Unfortunate Revival of Civic Republicanism*, 141 U. PA. L. REV. 801 (1993) (surveying civic-republican scholarship). For a deliberative-democratic view of risk regulation, see RICHARD P. HISKES, DEMOCRACY, RISK, AND COMMUNITY: TECHNOLOGICAL HAZARDS AND THE EVOLUTION OF LIBERALISM 132-58 (1998).

[428] *See* Paul J. Weithman, *Contractualist Liberalism and Deliberative Democracy*, 24 PHIL. & PUB. AFF. 314, 314-17 (1995) (emphasizing deliberative democrats' aim to present a noninstrumental justification for democracy).

[429] *See Introduction* to DELIBERATIVE DEMOCRACY: ESSAYS ON REASON AND POLITICS, *supra* note 427, at x-xii (contrasting economic and pluralist models of democracy, which see citizens as advancing their interests, with deliberative-democratic views, which envision them deliberating about the public good)

[430] *Compare* ALEXANDER BICKEL, THE LEAST DANGEROUS BRANCH: THE SUPREME COURT AT THE BAR OF POLITICS 1-33 (1962) (defending a majoritarian conception of

Recent work within legal and political theory suggests that the most plausible democratic views look to citizen judgments rather than preferences—suitably impartial judgments about the "public good" or, more precisely, about which governmental choices are morally required.[432] Given the intractability of moral disagreement, governmental responsiveness to judgments is (arguably) an independent moral requirement.[433] Within this judgment-focused class of views, we might further distinguish between idealized and actualized views. Democratic accounts that demand responsiveness to idealized citizen judgments are not likely, I suggest, to warrant regulatory attention to frequentist risk. Idealized judgments should be relatively good (if not necessarily perfect) in tracking true moral requirements. But my argument in preceding Sections has (I hope) provided strong reason to believe that morality doesn't truly attend to frequentist risk.

We are left, then, with accounts that demand some kind of direct or indirect regulatory responsiveness to actual citizen judgments. If citizens widely judge (incorrectly) that the level of frequentist risk is a morally relevant feature of toxins and other possible targets for regulatory intervention, then, according to the kind of democratic view now under consideration, regulators have a moral reason to attend to frequentist risk, or at least legislators have a moral reason to direct regulators to do that. But *do* citizens, in fact, typically reach this judgment? This is an empirical question—one that implicates the psychological literature on "risk perception," a literature much too large for me systematically to engage here.[434] Let me merely note that the best-known and most influential scholarship in this area, that of Paul Slovic and his collaborators,[435] calls into question whether the level of frequentist "individual risk" is judged by citizens to be an im-

---

democracy), *with* 1 BRUCE ACKERMAN, WE THE PEOPLE 54-55, 266-94 (1991) (defending a supermajoritarian conception).

[431] *See* KENNETH J. ARROW, SOCIAL CHOICE AND INDIVIDUAL VALUES (1951). For a recent survey of the literature on "cycling" and related features of majority rule, see DENNIS C. MUELLER, PUBLIC CHOICE III 79-127, 147-58 (2003).

[432] That politics should (in some way) be public-regarding is a key claim advanced by deliberative democrats and civic republicans. *See Introduction* to DELIBERATIVE DEMOCRACY, *supra* note 427, at xiii-xiv; Cass R. Sunstein, *Beyond the Republican Revival*, 97 YALE L.J. 1539, 1548-51, 1554-55 (1988).

[433] *See* JEREMY WALDRON, LAW AND DISAGREEMENT 10-17 (1999).

[434] For an introduction to this literature, see Nick F. Pidgeon & Jane Beattie, *The Psychology of Risk and Uncertainty*, *in* HANDBOOK OF ENVIRONMENTAL RISK ASSESSMENT AND MANAGEMENT, *supra* note 23, at 289, 296-301.

[435] This work is collected in PAUL SLOVIC ET AL., THE PERCEPTION OF RISK (Ragnar E Löfstedt ed., 2000).

portant determinant of appropriate regulatory choice. Slovic's standard technique is to elicit, from lay survey respondents, quantitative judgments of the overall "riskiness" of various hazards, plus assessments of each hazard with respect to a host of different dimensions, and then to determine which dimensions best explain perceived riskiness.[436] For example, in one study, respondents were asked to rank ninety hazards (such as home gas furnaces, home appliances, microwave ovens, or nuclear weapons) on a 0-100 scale from "not risky" to "extremely risky," and to rank each hazard on a 1-7 scale for eighteen different dimensions: the hazard's observability; the degree to which those exposed are aware of the hazard; the number of people exposed to the hazard; the degree to which the hazard is widely feared; the degree to which the hazard is catastrophic; the threat to future generations; and so on.[437]

Slovic's general finding is that lay judgments of riskiness are not solely determined by lay beliefs about the number of people exposed to some hazard or the aggregate deaths resulting from some hazard. Rather, two other broad factors (each combining a host of correlated dimensions) also are important in explaining perceived riskiness. These factors are termed, by Slovic, "dread" and "familiarity". "Dread" subsumes the risk's controllability, the fear it provokes, whether it has catastrophic consequences, whether it has fatal consequences, the equity of its distribution, the threat to future generations, the ease of reduction, whether the risk is increasing over time, and whether it's involuntary. The "familiarity" factor subsumes the risk's observability, the awareness of those exposed, whether its effects are delayed, its novelty, and whether the risk is known to science.[438] Although a few of the dimensions underlying "dread" and "familiarity" may, in turn, depend on "individual risk"—for example, this may be true of the equity dimension—many seemingly do not. "Individual risk" to the maximal, highly exposed, or representative individual, so crucial to agency risk assessments, does not (if Slovic's results are accurate)[439] play much of a role in determining actual citizen judgments about the riskiness of hazards and the need for government intervention.

---

[436] *See* Paul Slovic, *Perception of Risk*, in SLOVIC ET AL., *supra* note 435, at 220-22.

[437] Paul Slovic et al., *Facts and Fears: Understanding Perceived Risk*, in SLOVIC ET AL., *supra* note 435, at 137-39.

[438] *See id.* at 139-41.

[439] *See* Pidgeon & Beattie, *supra* note 434, at 299-301 (noting that Slovic's factor structure is confirmed by some, but not all, subsequent work on psychology of risk perception).

VI. RISK ASSESSMENT AND POPULATION SIZE

My analysis, to this point, has sought to demonstrate that "individual risk" in the frequentist sense should be irrelevant to regulatory evaluation of toxic and other possibly fatal hazards. "Individual risk" in the Bayesian sense, be it first-party or third-party Bayesian risk, may well have an appropriate role in regulatory choice. But frequentist risk and Bayesian risk are different, and neither welfare consequentialism, nor competing moral views—safety consequentialism, deontology, contractualism, or democratic views—focus the regulator on frequentist risk.

I now want to argue for a second and distinct claim: that risk assessment should be sensitive to population size.[440] Regulatory techniques for evaluating hazards should incorporate information about the numbers of persons at risk of dying from the hazards. The regulatory practices described in Part II—where agencies such as EPA, OSHA, FDA, NRC, and CPSC employ the frequentist risk to the maximally exposed, highly exposed, or average exposed individual as a partial determinant of regulatory choice—are *doubly* misguided. These practices are misguided, first, because they focus on frequentist rather than Bayesian risk and, second, because a regulatory criterion that looks to the level of "individual risk" incurred by the maximal, high-end, or average individual is insensitive to population size.[441]

---

[440] I am certainly not the first to argue that risk regulation should be sensitive to population size. *See* CASS R. SUNSTEIN, RISK AND REASON: SAFETY, LAW, AND THE ENVIRONMENT 214-16 (2002); HAMILTON & VISCUSI, *supra* note 16, at 21-23. However, my demonstration that the demand for size-sensitivity is robust across plausible moral views, both consequentialist and nonconsequentialist, is (I believe) novel.

[441] By "size insensitive," I mean a regulatory criterion that indicates the very same regulatory response to hazards that have identical distributions of risk across the exposed populations and that are otherwise identical, but differ in the size of those populations. *Ceteris paribus*, changes in the number of persons exposed to a hazard do not change what the criterion instructs regulators to do.

To be sure, a population-size-insensitive criterion might be embedded in a broader regulatory practice that is sensitive, in some way, to population size. For example, satisfying an "individual risk" test focused on the risk level incurred by the average, high-end, or maximal individual might be a necessary condition for regulating some hazard; "population risk" considerations might then inform the level of regulation for hazards that satisfy the initial condition. Some of the regulatory practices described in Part II are population-size sensitive in this broader sense. But they still incorporate a size-insensitive, "individual risk"-based criterion—and that itself is hard to justify.

It might also be objected that looking to the level of "individual risk" incurred by some person in the exposure distribution *is* population-size sensitive where the entire U.S. population is seen as "exposed" to the hazard. *Cf.* Rhomberg, *supra* note 53, at

To put the point another way: the practices described in Part II cannot be salvaged by substituting Bayesian for frequentist risks. Regulating foods, waste dumps, pesticides, air pollutants, nuclear plants, or other hazards by ensuring that the maximally exposed, highly exposed, or average exposed individual has a first-party or third-party Bayesian risk below some threshold (be it 1 in 1 million, 1 in 10,000, or something else) would be morally arbitrary. This claim, like the claimed irrelevance of frequentist risk, is (with a few exceptions) robust across plausible moral views. Both welfare consequentialism and alternative moral views generally demand that regulatory criteria for addressing hazards attend to the *number* of persons incurring various levels of (Bayesian) risk from the hazards.

My argument for this claim will be relatively brief. Because I have already characterized the moral views, and the regulatory decision procedures they support, in some detail, it should be pretty clear that population-size insensitive procedures are generally problematic. *Welfare consequentialism* employs cost-benefit analysis (CBA) as its decision procedure.[442] CBA, in turn, has both "population risk" and "individual risk" versions.[443] The "population risk" version predicts the mean number of deaths resulting from toxic or other hazards, and ascribes a monetary price to each death; these are aggregated, counted as costs of the hazard and benefits of intervention, and then added to the overall sum of monetized costs and benefits to determine whether regulatory intervention has net costs or benefits. The "individual risk" version describes the distribution of (third-party) Bayesian "individual risk" resulting from the hazard, across the population, and ascribes a monetary cost to each risk, based on individual willingness to pay to avoid the risk of death; these monetized risks are aggregated and then added, once more, to the overall sum of monetized costs and benefits.

1055-59 (describing FDA exposure assessments, which look at the entire U.S. food-consuming population). But this point, too, is incorrect. Imagine that FDA will approve an additive if the "individual risk" to the average U.S. citizen is less than 1 in 1 million. Then this criterion requires the very same result regardless of the total number of U.S. citizens.

[442] *See supra* text accompanying notes 297-99. Admittedly, this is not true for certain egalitarian variants of welfare consequentialism. One imagines, however, that even these variants would require size-sensitive procedures. Compare two cases in which a hazard lowers the expected welfare of an exposed population that is already worse off than nonexposed persons: in one case the exposed population has many more members than the non-exposed group; in the other case, it has many fewer. Presumably the egalitarian would (or at least might) treat the two cases as differentially serious departures from strict equality.

[443] *See supra* text accompanying notes 300-04.

Both versions of CBA are sensitive to population size. The "population risk" version is size sensitive, obviously, because the number of deaths predicted to result from a hazard depends on the number of exposed individuals. The "individual risk" version is size sensitive because the aggregate willingness to pay to avoid the risks created by a hazard depends on the number of persons incurring those risks as well as the distribution of personal characteristics such as wealth and age that influence willingness to pay.

Safety consequentialism focuses on longevity and health, not the other sources of welfare. Thus, it counsels regulators to employ "risk-risk" analysis, or some such procedure, rather than CBA.[444] Most straightforwardly: regulators would evaluate a proposed intervention by determining whether intervention is likely to produce a net decrease or increase in the number of premature deaths, or a net increase or decrease in the population's aggregate life years or health-adjusted life years. All these procedures would require risk assessment techniques that are sensitive to the numbers of persons at various levels of risk from hazards. A workplace toxin, say, which is used in one hundred specialized workplaces, with an average risk to exposed workers of 1 in 1000 (OSHA's level of clear significance), will likely produce many fewer deaths than a toxin used in one hundred times more workplaces at the same average risk.

The *deontological* view of risk regulation asserts that regulators have a (prima facie) obligation to enforce deontological rules, including the classic no-killing constraint and perhaps a separate no-risking constraint. As for the no-killing constraint, I have suggested that this constraint, if enforceable by government, would warrant a decision-procedure that enjoins the regulator: "*ceteris paribus*, intervene to prevent any hazardous activity that you believe, to degree *p* or higher, would kill at least one person."[445] This decision procedure *is* sensitive to population size. That point is subtle, but crucial. There is a crucial distinction between the probability of a killing, and the probability that some particular individual will be killed. My suggestion is that a risk assessment practice enforcing the deontological rule against killing would focus on the former probability, not the latter, and thus would be sensitive to the numbers of individuals at various (Bayesian) levels of risk of being killed.

---

[444] *See supra* text accompanying notes 361-63.

[445] *See supra* text accompanying notes 377-78.

Imagine, for example, that the Bayesian risk threshold set forth by the deontological decision procedure is 1 in 1 million. The enforcement rule says: "intervene, *ceteris paribus*, to stop some hazardous activity when you believe to a degree greater than 1 in 1 million that at least one person will be killed." Crucially, this rule is *not* the same as a rule that instructs the regulator to intervene if the "individual risk" to the maximally at-risk person, or any other person, exceeds 1 in 1 million. If a hazardous activity exposes everyone in a population of two hundred million to a toxin, and the regulator's degree of belief that any given person will die is 1 in 10 million, then the Bayesian risk to the maximally exposed person (and also the average and high-end person) is less than 1 in 1 million, but the regulator's Bayesian probability that the toxin will kill at least one person is virtually one.[446] What the deontological probability threshold for the no-killing constraint concerns, again, is the probability that *some* person will be killed, and that probability *increases* as increasing numbers of individual persons are each exposed to a nonzero risk of being killed.

What about the putative deontological constraint against risk imposition? I have evinced skepticism that any such constraint exists. Risk imposition need not be harmful, invasive, or manipulative, the hallmarks of deontological violations. It need not reduce the physical or emotional well-being of the person put at risk, invade her body or property, or manipulate her.[447] But perhaps risk imposition involves a more intangible, dignitary harm. Thus the "expressive" view of the no-risking constraint: risk imposition is a kind of insult. Yet is it, necessarily? A concern with insults could justify a constraint against actions that the actor intends to be insulting or that are conventionally insulting given existing linguistic rules. More puzzling is why risking, as such, should be constrained. Perhaps social meaning comes into play: any action that the ordinary member of the relevant society believes, to a degree greater than *q*, will be a killing of a given person, has the "social meaning" of insulting that person.[448] So perhaps there *is* a no-risking constraint that invokes the Bayesian probabilities of ordinary observers—reasonable persons—and a matching decision procedure that tells regulators, "intervene *ceteris paribus* to prevent some

---

[446] At the limiting point where the probability threshold *p* equals zero, the no-killing rule *does* generate a size-insensitive procedure: the regulator's Bayesian probability that some person will be killed is nonzero just in case the maximally exposed individual has a nonzero probability of being killed.

[447] *See supra* text accompanying notes 381-85.

[448] *See supra* text accompanying notes 386-92.

hazard if the 'reasonable person' would believe, to a degree greater than *q*, that the hazard will kill any given individual." This deontological decision procedure *would* license population–size insensitive risk assessment practices. Specifically, it would warrant regulatory attention to the Bayesian "individual risk" imposed on the maximally exposed individual. If a hazardous activity creates a Bayesian risk greater than *q* (as judged from the epistemic standpoint of the reasonable person) that the maximally exposed individual will be killed, then the hazard "expresses" contempt for that person and, *ceteris paribus*, should be stopped; if, conversely, the Bayesian risk to the maximally exposed individual is less than *q*, then no one suffers a risk level exceeding *q*, and no one has been insulted.

The deontological no-risking constraint thus creates an exception to my claim about population-size sensitivity. But it is a dubious exception, since even those sympathetic to deontology and the classic, no-killing constraint have good cause to wonder whether a no-risking constraint really exists; and even those sympathetic to the view that government should enforce some deontological constraints might wonder whether a no-insult rule is important enough to fall in the category of enforceable constraints.

Contractualism, at least in the act-contractualist variants, asks the regulator to imagine a hypothetical referendum among the citizenry: as between various options, perhaps paired with hypothetical transfer payments, which option would the citizens unanimously agree upon if suitably informed and motivated?[449] It might seem that act-contractualism is insensitive to population size, because any person—even a single one—has a veto in the hypothetical referendum. But the issue is more subtle than that. First, the Kaldor-Hicks variant of act-contractualism is actually size sensitive. The fewer persons there are at risk from some hazard, the smaller the hypothetical transfer payment needed to compensate them for the risk. Indeed, economists have suggested that the Kaldor-Hicks criterion should be implemented through CBA (presumably a CBA procedure employing first-person rather than third-person Bayesian probabilities).[450] The intuition, here, is that the aggregate monetized benefits associated with a regulatory choice track the transfer payment that the beneficiaries from that choice could afford to pay, and the aggregate

---

[449] *See supra* text accompanying notes 402-14.

[450] *See* Adler & Posner, *Rethinking Cost-Benefit Analysis*, *supra* note 36, at 190-91 (describing Kaldor-Hicks defense of CBA); FREEMAN, *supra* note 300, at 222-29 (discussing implications of Kaldor-Hicks approach for CBA under uncertainty).

monetized costs associated with the choice track the transfer payment that the losers would require as compensation.

Second, and more generally, act-contractualist views do *not* justify the kinds of regulatory practices described in Part II. In determining whether any given person would exercise her "veto" over some regulatory choice, we would need to consider all the effects of the choice on her, not simply its effect on her risk of premature death. The fact that the average (Bayesian) risk imposed by some hazard, and a fortiori the risk imposed on the maximally exposed person, exceeds some threshold, does not imply that these particular persons would "veto" the regulatory option of inaction. That depends on whether the hazard has countervailing benefits for these persons, and whether the other options have countervailing costs.[451] Workers exposed to a risky toxin might be willing to bear the risk, if the alternative is lower risk but lower wages (given the cost of containing the toxin or finding a substitute); food consumers exposed to a risky food might be prepared to eat it, if they like the way it tastes and lower-risk substitutes are pricier.

Consider, finally, *democratic* views of risk regulation, which posit a bedrock moral obligation on the part of regulators to respond to citizen judgments.[452] If the citizenry is sufficiently concerned about risk imposition itself, apart from death, injury, pain, fear, or catastrophe, then risk regulation practices focusing (in part) on "individual risk" to the maximally exposed, highly exposed, or average individual could be justified. But if, instead, citizen judgments about risk regulation tend to focus on particular kinds of deaths, for example "dreaded" or "unfamiliar" deaths—as seems more likely[453]—then risk regulation should be keyed to a function of "population risk." The requisite practice would not be a straightforward minimization of premature death or maximization of longevity, as per safety-consequentialism, but rather a more democratically attuned procedure where different types of deaths are given different weights in line with their popular characterization. Preventing the 1000 deaths caused by a radioactive release from a faulty nuclear reactor would take priority over preventing 1000 deaths from poorly designed consumer products. Still, the likely number of deaths from a reactor release depends on the size of the population adjacent to the plant, and the likely number of con-

---

[451] Here, I'm assuming the self-interested variant of act-contractualism. In the case where contractors are more public-regarding, insofar as the moral views motivating them are population-size sensitive, the act-contractualist criterion also is.

[452] *See supra* text accompanying notes 432-33.

[453] *See supra* text accompanying notes 434-39.

sumer deaths depends on the total number purchasing the flawed goods.

CONCLUSION

"Individual risk" has come to play a large role in federal health and safety regulation. This is particularly true for toxic chemicals, both carcinogens and noncarcinogens, regulated at the federal level by EPA, FDA, and OSHA, as well as by other agencies (for example CPSC). EPA employs an "individual risk" test as a decisional trigger in the air pollution context; as a standard-setting criterion in the water pollution context, for both ambient and drinking water toxins; in determining when a substance containing toxic chemicals is a "hazardous waste" and therefore subject to various stringent regulatory constraints; in licensing and setting tolerances for toxic pesticides; in evaluating whether and, if so, how aggressively to clean up waste dumps; in policing "environmental justice"; and elsewhere.[454] FDA relies heavily on an "individual risk" approach in licensing toxic food additives and, to a lesser extent, in regulating food safety more generally.[455] OSHA, since the Supreme Court's seminal decision in *Industrial Union* (1980),[456] has taken the position that workplace toxins cannot be regulated unless they impose a sufficiently high "individual risk" on workers.[457] These three agencies also have a role in regulating pathogens: "individual risk" comes into play there,[458] as it does for NRC and other bodies in regulating radiation[459] and, recently, for OSHA in addressing safety hazards.[460]

This Article has argued that the range of regulatory practices just mentioned, and described at much greater length in Part II, are doubly misguided—indeed, doubly misguided across a range of moral theories, both consequentialist theories (including but not limited to the welfare-centered consequentialism that grounds cost-benefit analysis and dominates modern economics) and nonconsequentialist theories (specifically deontological, contractualist, and democratic theories). First, the practices just mentioned focus on "individual

---

[454] *See supra* Part II.A.

[455] *See supra* Part II.B.1.

[456] Industrial Union Dep't v. Am. Petroleum Inst., 448 U.S. 607 (1980) (plurality opinion).

[457] *See supra* text accompanying notes 186-88.

[458] *See supra* note 93; *supra* text accompanying notes 165, 199.

[459] *See supra* Part II.B.3.

[460] *See supra* text accompanying notes 193-97.

risk" in the frequentist rather than Bayesian sense; second, they are insensitive to population size. Neither feature is morally supportable.

Risk assessment practices that invoke frequentist risk, and are insensitive to population size, are now pervasive. So this Article *is* a critique of risk assessment. But it is a sympathetic critique—a critique of the methodology as it's now (mis)used, not generically. Nothing in the laudable ambition to quantify health and safety threats entails frequentism, or a focus on the risk to some individual in the exposure distribution rather than, say, aggregate deaths, aggregate killings, aggregate loss of life years (quality adjusted or not), or the total *numbers* of individuals incurring various "individual risk" levels.[461]

My analysis has been critical, not constructive. I have not proposed a specific decision procedure that health and safety regulators should employ in evaluating regulatory measures or, at the threshold, in deciding whether to regulate at all or to deliberate about regulation. It is ludicrous to think that different moral views would converge on the same, specific procedure. Indeed they would not: welfare consequentialists care about deaths and much else; safety consequentialists, about deaths and little more; deontologists, about killings; democrats, about whatever citizens care about.[462] What is *not* ludicrous is to imagine that thoughtful observers, rooted in different moral traditions, might all converge in *rejecting* some particular practice. My argument, in effect, has been for a negative version of what Rawls calls an "overlapping consensus":[463] we may not be able to come to overlapping consensus on how health and safety threats should be regulated, but we can (I think) come to agree that the regulatory status quo should change. Relatedly, we can come to consensus, in a very general way, on what the change should be. Whatever the moral view, probabilities do have a role to play in regulatory choice, but those should be measures of belief, not proportions or frequencies. More radically: the level of "individual risk" (Bayesian or frequentist) imposed on the average, high-end, or maximal individual should have no function (even a prima facie one) in regulatory decision procedures.

It might be objected that my analysis has been moral, not legal. This is a function of the generality of the analysis. There is a forest/trees problem: an analysis focused (say) on the Clean Air Act, or

---

[461] *See supra* text accompanying notes 442-46 (discussing these alternative possible procedures).

[462] *See supra* Parts III.B., V.

[463] *See* JOHN RAWLS, POLITICAL LIBERALISM 133-72 (1993).

even on the environmental statutes generally, would miss the broader point that current risk regulation practices in many different statutory contexts have common, problematic features. Surely the objection can't be that moral and legal considerations are wholly separate.[464] Moral argument is one legally appropriate source—although not the only source—of statutory interpretation.[465] So think of the analysis, here, as a starting point in interpreting a wide range of statutes: the Clean Air Act, the Clean Water Act, the Safe Drinking Water Act, RCRA, CERCLA, FIFRA, the Toxic Substances Control Act, the Food, Drug and Cosmetic Act, the Occupational Safety and Health Act, the Atomic Energy Act, the Consumer Product Safety Act, the Federal Hazardous Substances Act, the motor vehicle safety laws, and others.[466] Some of these statutes give high priority to safety; others do not. But the statutory priority given to safety is *orthogonal* to the issues analyzed here. The Food, Drug and Cosmetic Act, for example, demands that food additives be "safe,"[467] and the Occupational Safety and Health Act tells the agency to regulate toxins so as "to assure[], to the extent feasible, . . . that no employee will suffer material impairment of health or functional capacity."[468] By contrast, FIFRA instructs EPA to license pesticides that will not cause "any unreasonable risk to man or the environment, taking into account the economic, social, and environmental costs and benefits"[469] of the pesticides; and the Consumer Product Safety Act employs a similar "unreasonable risk of injury" standard.[470] Yet "population risk" tests are no less legally permissible—

---

[464] *See generally* RONALD DWORKIN, LAW'S EMPIRE (1986) (arguing that law and morality are necessarily connected); Kenneth Einer Himma, *Inclusive Legal Positivism*, *in* THE OXFORD HANDBOOK OF JURISPRUDENCE AND PHILOSOPHY OF LAW 125-66 (2002) (summarizing debate about "inclusive legal positivism," the view that law possibly incorporates moral considerations).

[465] *See* CASS R. SUNSTEIN, AFTER THE RIGHTS REVOLUTION: RECONCEIVING THE REGULATORY STATE 111-59 (1990).

[466] *See* 42 U.S.C. §§ 7401-7671q (2000) (Clean Air Act); 33 U.S.C. §§ 1251-1387 (2000) (Clean Water Act); 42 U.S.C. §§ 300f to 300j-26 (2000) (Safe Drinking Water Act); 42 U.S.C. §§ 6901-6992k (2000) (RCRA); 42 U.S.C. §§ 9601-9675 (2000) (CERCLA); 7 U.S.C. §§ 136-136y (2000) (FIFRA); 15 U.S.C. §§ 2601-2692 (2000) (Toxic Substances Control Act); 21 U.S.C. §§ 301-397 (2000) (Food, Drug and Cosmetic Act); 29 US.C. §§ 651-678 (OSH Act); 42 U.S.C. §§ 2011 to 2297h-13 (2000) (Atomic Energy Act); 15 U.S.C. §§ 2051-2085 (2000) (Consumer Product Safety Act); 15 U.S.C. §§ 1261-1278 (2000) (Federal Hazardous Substances Act); 49 U.S.C. §§ 30101-30170 (2000) (motor vehicle safety provisions).

[467] 21 U.S.C. § 348(c)(3)(A) (2000).

[468] 29 U.S.C. § 655(b)(5) (2000).

[469] 7 U.S.C. §§ 136(bb), 136a(c)(5) (2000).

[470] 15 U.S.C. §2056(a) (2000).

I have suggested—under the first pair of statutes than under the second pair. It is perfectly plausible for FDA and OSHA to take the position, as a matter of statutory interpretation, that an additive or workplace toxin causing fewer than some threshold number of deaths is legally safe.[471]

Safety priority does not moot my analysis, legally speaking. Other statutory language would. If health and safety statutes explicitly mandated population-size insensitive regulatory procedures for evaluating hazards, and instructed agencies to focus on frequentist rather than Bayesian risk, then agencies would be *legally* obliged to do both regardless of the moral difficulties. I have not suggested that current practices are so deeply immoral that a statutory requirement undergirding them would be unconstitutional, or justify civil disobedience by regulators. But in fact these practices typically do not flow from explicit statutory mandate. The Clean Air Act is an exception, and even here what is mandated is EPA's use of the 1 in 1 million risk test as a decisional trigger and a criterion for exempting sources from regulation, without any requirement that the risk be frequentist rather than Bayesian.[472]

In any event, and more importantly, the Clean Air Act is virtually the only federal health and safety regulatory statute that I'm aware of where an agency is required by clear statutory text to attend to "individual risk" in reaching some decisions and further to employ population-size insensitive procedures or a frequentist conception of "individual risk."[473] In general, then, a shift to a Bayesian and population-size sensitive approach to risk regulation is both morally *and* legally warranted. And it is a shift that regulators, legislators, legal scholars, policy analysts, and the risk assessment community need to start contemplating. Given the size and impact of the federal regulatory estab-

---

[471] *See supra* note 19 and accompanying text.

[472] *See* 42 U.S.C. §§ 7412(c)(9)(B)(i), (f)(2)(A) (2000); *supra* text accompanying notes 99-102 (explaining these provisions).

[473] EPA is instructed to use a 1 in 1 million "individual risk" test for setting pesticide tolerances by a congressional committee report, but not by the statutory text of the Food Quality Protection Act of 1996. *See supra* text accompanying notes 139-41. A different provision of that statute, permitting higher tolerances under special conditions, does mandate a size-insensitive, "individual risk" test. *See* 21 U.S.C. § 346a(b)(2)(B) (2000); Madigan, *supra* note 137, at 202-04 (describing this provision). Other possible examples of statutory language requiring agencies to employ population-size insensitive, "individual risk" tests for evaluating hazards are 21 U.S.C. § 360*ll*(a)(2) (2000) (radiation risks from electronic products); and 42 U.S.C. § 263b(h)(2) (2000) (mammography facilities). I have identified no statutory language mandating risk assessment procedures keyed to frequentist rather than Bayesian risk.

lishment, how these agencies should go about their business is no trivial matter. The choices between "individual risk" and "population risk," and between frequentism and Bayesianism, are too significant to remain the esoterica of risk assessors and toxicologists.