Reasonable Precaution for the Individual

Dov Waisman[*]

     This Article offers a new answer to a hotly-debated question in normative ethics, one that has important implications for accident law. When a firm engages in a socially beneficial activity that exposes workers, consumers, or bystanders to a risk of death or serious bodily harm—e.g., building a skyscraper, manufacturing prescription drugs, or operating a nuclear power plant—what level of precaution is the firm morally required to exercise? How, in other words, should the <u>reasonable</u> *level of precaution be determined in a case of this type?*

     In a series of provocative essays, Professor Barbara Fried has recently argued that standard cost-benefit analysis represents the "only game in town" when it comes to answering this question. Though cost-benefit analysis's reliance on interpersonal aggregation—its summing of costs and benefits across individuals—has long been faulted for failing to respect the separateness of persons, Fried contends that there is simply no workable alternative for sorting permissible practices of risk imposition from impermissible ones.

     Responding to Fried's challenge, this Article introduces a new interpretation of reasonable precaution—the individualized feasibility principle —that focuses on costs and risks to each affected individual, rather than on costs and risks considered in the aggregate. This principle holds that, when engaging in a significantly risky but socially beneficial activity, an actor is morally obligated to invest in safety precaution until the lesser of the following two points is reached: (i) the point at which further expenditure on safety would threaten the long-term survival of the activity; or (ii) the point at which further expenditure would reduce the well-being of each person responsible for bearing precaution costs by more than it would in-

crease the expected well-being of each person exposed to the risk. Because it does a better job of accommodating widely-held intuitions in important types of cases, the individualized feasibility principle deserves to be considered alongside aggregative standards in the ongoing quest for a comprehensive theory of reasonable precaution.

TABLE OF CONTENTS

Introduction

Negligence law rests on the idea that imposing risks of harm on others is an acceptable and unavoidable part of social life. Most of the socially beneficial activities that make modern life possible—from building skyscrapers to manufacturing prescription drugs to operating power plants—impose nontrivial risks of death or serious bodily harm on workers, consumers, or bystanders, but such risks are widely thought to be morally tolerable provided they are moderated by the exercise of reasonable precaution.

There is no consensus, however, about just what it means to exercise reasonable precaution. United States courts typically rely on the notion of a reasonable person, defining reasonable precautions as those a reasonably careful or prudent person would have taken under the circumstances.[1] However, in the almost seventy years since Judge Learned Hand articulated his famous negligence formula in *United States v. Carroll Towing Co.*,[2] no *substantive* interpretation of reasonable precaution—no interpretation that purports to specify the decision rule that a reasonable person would follow in deciding whether to take a particular precaution—has gained anything close to widespread acceptance.[3]

One idea that *has* found widespread acceptance is Hand's insight that reasonable precaution has something to do with balancing the burdensomeness of taking a particular safety precaution against the probability and severity of the injuries the precaution protects against.[4] It is unreasonable not to take a minimally burdensome precaution if it will significantly reduce

---

[1] *See* RESTATEMENT (THIRD) OF TORTS: LIABILITY FOR PHYSICAL AND EMOTIONAL HARM § 3 cmt. d (Reporter's Note) (2010) ("As far as the jury is concerned, pattern jury instructions frequently explain the negligence standard in terms of the reasonably prudent person."); *see also* Patrick J. Kelly & Laurel A. Wendt, *What Judges Tell Juries about Negligence: A Review of Pattern Jury Instructions*, 77 CHI.-KENT L. REV. 587, 595 (2002) ("In most pattern jury instructions on negligence, negligence is defined by using both the concept of ordinary care and the concept of the conduct of a reasonably careful person or one of her relatives.").

[2] *See* United States v. Carroll Towing Co., 159 F.2d 169, 173 (2d Cir. 1947) ("if the probability [of injury] be called P; the injury L; and the burden [of adequate precautions], B; liability depends upon whether B is less than L multiplied by P: i.e., whether B < PL"); *see also* Moisan v. Loftus, 178 F.2d 148, 149 (2d Cir. 1949); Conway v. O'Brien, 111 F.2d 611, 612 (2d Cir. 1940).

[3] *See* Alan D. Miller & Ronen Perry, *The Reasonable Person*, 87 N.Y.U. L. REV. 323, 327, 328-70 (2012) (describing a range of different substantive interpretations of the reasonable person standard based on differing "normative ethical commitments" that "may be mutually exclusive or inconsistent in fundamental respects").

[4] *See* RESTATEMENT (THIRD) OF TORTS: LIABILITY FOR PHYSICAL AND EMOTIONAL HARM § 3 cmt. d (Reporter's Note) (2010) (observing that risk-benefit balancing approach to negligence has been accepted by a number of leading torts treatises, by many deterrence- and corrective justice-oriented tort scholars, and in judicial opinions "in a large majority of jurisdictions").

the likelihood of someone suffering a serious injury. As the precaution be-comes more burdensome, the idea of not taking the precaution becomes less unreasonable. Yet, beyond this fundamental Handian insight, there is little consensus about exactly where and how to draw the line between reasonable and unreasonable precaution in particular cases.

Perhaps the most influential attempt to do so is based on the notion of aggregate costs. This interpretation of reasonable care rests on the axiom that the goal of risk regulation is, as Guido Calabresi once put it, to "mini-mize the sum of the costs of accidents and the costs of avoiding accidents."[5] If so, then reasonable care should consist in taking all and only those pre-cautions that cost less than the amount they would be expected to save in total accident costs.[6] Put in marginalist terms, reasonable care means spend-ing on safety precaution up to the point at which the next dollar spent yields less than a dollar's reduction in expected accident costs.[7] Call this reasona-ble care as *efficient care*.[8] This interpretation certainly has its virtues. For one thing, it is almost certainly the most straightforward interpretation of Hand's insight. Moreover, it nicely vindicates the intuition that the reason-able level of precaution can never be so burdensome as to make the risky activity not worth engaging in.

But efficient care is dogged by a persistent problem. By taking a "bottom line" approach that focuses exclusively on aggregate costs and benefits, summed across all affected persons, efficient care remains insensi-tive to how costs and benefits are *distributed* among different individuals, and therefore leads to results some find counterintuitive.[9]

For example, suppose a major city with a population of five million people is building a half-mile-long suspension bridge across a river. The

---

[5] Guido Calabresi, The Costs of Accidents: A Legal and Economic Analysis 26 (1970).

[6] *See, e.g.,* Richard A. Posner, *A Theory of Negligence*, 1 J. Legal Stud. 29, 32-33 (1972).

[7] *See, e.g.,* William M. Landes & Richard A. Posner, THE ECONOMIC STRUCTURE OF TORT LAW 87 (1987) (advocating marginalist interpretation of Hand formula); Stephen G. Gilles, *On Determining Negligence: Hand Formula Balancing, the Reasonable Person, and the Jury*, 54 VAND. L. REV. 813, 826 n. 39 (2001) (noting "the issue is the marginal costs and benefits of additional precautions").

[8] The type of efficiency at issue here is Kaldor-Hicks, rather than Pareto. *See, e.g,* Mat-thew D. Adler, WELL-BEING AND FAIR DISTRIBUTION 98-99 (2012) (defining Kaldor-Hicks efficiency).

[9] So-called "distributively weighted" cost-benefit analysis makes some attempt to account for distributional differences, but is still interpersonally aggregative. *See id.* at 109 (2012) ("Distributively weighted CBA takes the form of *summing individual WTP/WTA amounts multiplied by individual weights*.") (emphasis added); *see also* Matthew D. Adler, Cost-Benefit Analysis and Distributional Weights: An Overview (August 20, 2013) (unpublished manuscript) (available at http://www.ssrn.com/abstract =2313388). Even distributively weighted cost-benefit analysis is therefore unlikely to be able to capture moral intuitions founded on interpersonal tradeoffs occurring at the level of the individual.

city is considering whether to invest in special safety netting[10] to protect the 2,000 construction workers on the bridge project (none of whom lives in the city) from fatal falls. Suppose that, although safety netting has been widely used in the construction industry for decades, it has not been universally adopted.[11] If the city invests in the safety netting, one worker can be expected to die in a fall over the course of the five years it will take to complete the bridge. If the city uses only regular scaffolding without netting, two workers can be expected to fall to their deaths during those five years. Using the safety netting will add $10 million to the $1 billion cost of constructing the bridge, which is divided equally among the five million city residents. Assigning a value of $5 million[12] to a human life, the city calculates that the costs of the safety netting ($10 million) substantially exceed the expected accident costs associated with not using it ($5 million).[13] On that account, is it morally permissible for the city to use only regular scaffolding in constructing the bridge?[14]

If reasonable care is efficient care, then the city has no moral obligation to invest in the safety netting. This is true despite that the safety netting cuts each worker's risk of death in half (from 1 in 1,000 to 1 in 2,000), can be expected to save one worker's life, and would impose an additional cost of just $2 on each city resident.

---

[10] Occupational Safety & Health Administration, 29 C.F.R. 1926.502(c) (2013) (designating "safety netting" an acceptable "fall protection system" for employees).

[11] *See, e.g.,* Press Release, Health and Safety Executive (Great Britain), Contractor in Court for Ignoring Safety Risks (January 27, 2014) (available at http://press.hse.gov.uk/ 2014/contractor-in-court-for-ignoring-safety-risks-2/) (noting contractor "had failed to ensure that protective safety measures, such as scaffold edge protection and *safety netting*, were in place to prevent or mitigate a fall from height, leaving the three men at risk of serious or fatal injury") (emphasis added).

[12] Throughout this Article, I use this $5 million figure to represent the monetary value of a human life. It lies at the lower end of the spectrum of recent attempts by federal agencies to set the value of a statistical life (VSL). *See, e.g.,* Eric A. Posner & Cass R. Sunstein, *Dollars and Death*, 72 U. CHI. L. REV. 537, 549-51 (2005) (noting that "most regulatory agencies have now converged on a fairly narrow range for the valuation of life: $5 million to $6.5 million").

[13] Precisely this sort of calculation was at issue in the famous Ford Pinto case. *See* Grimshaw v. Ford Motor Co., 174 Cal. Rptr. 348 (Ct. App. 1981); *see also* Gary T. Schwartz, *The Myth of the Ford Pinto Case*, 43 RUTGERS L. REV. 1013, 1020 (1991) (describing a Ford report produced during discovery in *Grimshaw* which concluded the $11-per-vehicle cost of a safety device designed to prevent fuel tank fires in the Pinto (in 12.5 million vehicles, for a total of $137 million) exceeded the benefit of avoiding the 180 deaths and 180 serious burn injuries expected to occur were the safety device not included (a total of $49.5 million, using $200,000 as the value of a statistical life and $67,000 as the value of injury avoidance)).

[14] I place assumption of risk considerations to the side here. For purposes of the hypothetical, one can assume that the workers mean to willingly assume only the risk that remains *after all reasonable precautions have been taken*.   The workers' assumption of that risk itself plays no role in determining the reasonable level of precaution.

The intuition that the city ought to invest in the safety netting—notwithstanding that doing so is not marginally cost-justified and reduces the aggregate well-being of city residents and bridge workers on net—is, I submit, one that many people would hold.[15] The source of the intuition relates to the manner in which risks and safety costs are distributed: the risks and costs fall on distinct groups of persons, and the risks are concentrated on a group that is a tiny fraction of the size of the vast group of persons among whom the safety costs are spread. As a result, using the netting results in significantly enhanced safety for each worker at a cost that, while significant in the aggregate, imposes a trivial monetary burden on each city resident. Holding that a life-saving precaution like the safety netting need not be taken on account of its aggregate cost leaves efficient care vulnerable to the charge John Rawls famously leveled against utilitarianism: that by focusing exclusively on aggregate costs and benefits, and ignoring the costs and benefits to individuals, utilitarianism fails to "take seriously the distinction between persons."[16]

But is there an alternative interpretation of reasonable precaution that, unlike efficient care, can accommodate common moral intuitions in cases like the suspension bridge? It is all well and good to insist that morality or fairness requires taking more than the efficient level of precaution in some cases, but that position seems empty (or at least unconvincing) unless one can say with some specificity *when* pressing precaution beyond the point of efficiency is required and *how much* further precaution is required. It takes a theory to beat a theory, in other words.[17] So is there a viable alter-

---

[15] *See, e.g.,* Gregory C. Keating, *Pressing Precaution Beyond the Point of Cost-Justification*, 56 VAND. L. REV. 653, 656-57 (2003) (hereinafter "Keating, *Pressing Precaution*") ("[O]ur juries are repulsed by the claim that accidental deaths should not be prevented whenever the costs of prevention exceed the value—economically conceived—of the lives at risk."); Kenneth W. Simons, *Tort Negligence, Cost-Benefit Analysis, and Tradeoffs: A Closer Look at the Controversy*, 41 LOY. L.A. L. REV. 1171, 1180 (2008) ("The lesson that many people take from the Ford Pinto case is that the very act of engaging in cost-benefit analysis displays morally reprehensible callousness."); Cass R. Sunstein, *Moral Heuristics and Risk* in RISK: PHILOSOPHICAL PERSPECTIVES 162 (Tim Lewens ed. 2007) (citing studies showing "ordinary people. . .tend to punish companies that base their decisions on cost-benefit analysis, even if a high valuation is placed on human life"); *cf.* Schwartz, *supra* note __, at 1035-38 ("What seems obvious enough is that there exists a basic belief held by many (indeed most) of the public that it is wrong for a corporation to make decisions that sacrifice the lives of its customers in order to reduce the corporation's costs, to increase its profits.").

[16] *See* John Rawls, A THEORY OF JUSTICE 24 (1999); *see also* Barbara H. Fried, *The Limits of a Nonconsequentialist Approach to Torts*, 18 LEGAL THEORY 231, 231-32 (2012) (hereinafter "Fried, *Limits*") (addressing charge that cost-benefit analysis "fails to respect the separateness of people").

[17] *See, e.g.,* Fried, *Limits*, *supra* note __, at 259-60 ("I think it is fair to say that you cannot beat a bad candidate with no candidate, and in my view nonconsequentialists do not offer a viable candidate.").

native to efficient care and other aggregative interpretations of reasonable precaution?

In a series of provocative and important essays, Professor Barbara Fried has recently argued that no such alternative exists or even could exist.[18] Fried thinks it is impossible to strike an appropriate balance between the competing interests in liberty and security—a balance lying at the heart of the question whether a given risk imposition is morally permissible—without summing up the expected costs and benefits across persons and, in many instances, trading life-and-limb costs to one group of persons against individually trivial monetary benefits to another, much larger group. Non-aggregative approaches to risk regulation "*must* fail," Fried asserts, because "the problem of risk, by its nature can be managed *only* with the sorts of interpersonal tradeoffs. . .in which the numbers count, such that a risk of serious harm can be justified by small benefits to the many."[19]

In this Article, I show that Fried is mistaken, and that a viable, non-aggregative alternative to efficient care does indeed exist. If reasonable precaution is keyed to the costs and risks borne by each affected person, rather than to those borne by aggregates of persons, and to the long-term viability of the socially beneficial activity that poses the risk, rather than to the minimization of costs associated with the activity, a distinct interpretation of reasonable precaution emerges: the ***individualized feasibility principle*** (IFP). A bipartite standard, the IFP holds that, when engaging in a socially beneficial activity that imposes a risk of serious harm on certain individuals, exercising reasonable care means investing in safety precautions until the *lesser* of the following two points is reached : (i) the point at which further investment in safety would burden the well-being of each cost-bearer more than it would enhance the expected well-being of each risk-bearer or (ii) the point at which further investment in safety would jeopardize the long-term viability of the underlying activity.[20] In other words, the IFP requires the risks posed by a socially useful activity to be reduced to the

---

[18] *See id.* at 231-62; Barbara H. Fried, *Can Contractualism Save Us from Aggregation?,* 16 J. ETHICS 1, 39-66 (2012) (hereinafter "Fried, Contractualism"); see also Barbara H. Fried, *What Does Matter? The Case for Killing the Trolley Problem (or Letting It Die)*, 62 PHIL Q. 505, 512-13 (2012) (hereinafter, "Fried, Trolley"). Fried is William W. and Gertrude H. Saunders Professor of Law at Stanford Law School.

[19] Fried, *Contractualism*, *supra* note_, at 39.

[20] I follow Gregory Keating in holding that the principle of economic feasibility plays an important role in fixing the level of reasonable precaution with respect to risks of serious bodily harm imposed by major, socially productive activities. *See* Gregory C. Keating, *A Social Contract Conception of the Tort Law of Accidents*, in PHILOSOPHY AND THE LAW OF TORTS 22, 46-51 (Gerald Postema ed. 2001) (hereinafter "Keating, *Social Contract*"); Keating, *Pressing Precaution*, *supra* note_ , at 684-748; Gregory C. Keating, *Irreparable Injury and Extraordinary Precaution: The Safety and Feasibility Norms in American Accident Law*, 4 THEORETICAL INQ. IN L. 1, 25-88 (2003); Gregory C. Keating, *Pricelessness and Life: An Essay for Guido Calabre si*, 64 MD. L. REV. 159, 180-219 (2005).

maximum extent possible without jeopardizing the long-term survival of the activity, unless doing so would require each cost-bearer to invest more in safety precaution than a risk-bearer would rationally have been willing to invest to protect herself from the risk at issue, were she in the economic situation of the cost-bearer.

To illustrate using the bridge hypothetical, the IFP would require that the safety netting be used. The netting imposes a cost of $2 on each city resident, while reducing each worker's risk of death from 1 in 1,000 to 1 in 2,000. Using $5 million as a (conservative) value of a statistical life, each worker would rationally be willing to pay $2,500 for this reduction in risk. Since a safety gain worth $2,500 to each risk-bearer can be achieved at a cost of just $2 to each cost-bearer, the IFP requires the investment to be made so long as the widespread use of safety netting would not imperil the survival of the bridge construction industry. It would not since, by hypothesis, safety netting has long been widely-used in bridge construction. Thus, under the IFP, the safety netting is a reasonable precaution, notwithstanding that its use is inefficient (since it involves spending $10 million to achieve a $5 million savings in accident costs).[21] The IFP generates this result without summing costs and benefits across persons; instead, under the IFP, the netting's monetary cost to a *single* representative cost-bearer ($2) is compared with the netting's safety benefit to a *single* representative risk-bearer (a risk reduction worth $2,500). For this reason, the IFP can claim to respect the separateness of persons in a way standard cost-benefit analysis does not.

The IFP represents what I believe to be a novel and compelling interpretation of the injunction to exercise reasonable precaution when engaging in a socially beneficial activity that places some persons at risk of serious bodily harm.[22] Because it does a better job of accommodating common intuitions in important types of cases, the IFP deserves to be considered alongside aggregative standards in the ongoing quest for a comprehensive theoretical account of reasonable precaution.

---

[21] An economist might respond that this example shows only that human life has been undervalued, not that aggregative cost-benefit analysis fails to harmonize with moral commonsense. However, an analogous hypothetical can be constructed no matter how high the value assigned to a statistical life. For example, suppose life were assigned a value of *$20 million*, rather than $5 million. The same counterintuitive results follow if we stipulate that the cost of the netting is $40 million and that that cost is to be spread among 20 million city residents. In that case, we get the same result: cost-benefit analysis forbids the netting, even though it will reduce each worker's risk by 50%, will save a life, and will impose a cost of $2 on each city resident.

[22] Though Keating took the first, critical step toward this principle by focusing scholarly attention on the economic feasibility norm and its underlying moral rationale, no commentator has, to my knowledge, articulated the individual risk principle (*see infra*, part II) or proposed conjoining it with the feasibility norm in the manner I do in this Article.

This Article has five parts. In Part I, I introduce the question to be explored and describe Barbara Fried's challenge to any attempt to answer that question without summing costs and benefits across persons. Part II responds directly to Fried's challenge, presenting the individualized feasibility principle as a viable, nonaggregative interpretation of reasonable precaution. In Part III, I explore the theoretical underpinnings of the IFP, drawing on a theory of normative ethics known as *ex ante contractualism*.[23] Part IV explains why the contractualist justification for the IFP does not rest on the sort of interpersonal aggregation that fails to take seriously the distinction between persons. Part V concludes.

I.   The Challenge for Nonaggregative Risk Theory: Is Summing Costs and Benefits Across Persons the Only (Possible) Game in Town?

A.   The Question

Consider once again the suspension bridge hypothetical discussed in the Introduction. Five essential features of this example make it representative of a broad class of risk imposition cases: (1) a socially beneficial activity poses a risk of harm; (2) the risk posed is one of accidentally-inflicted injury; (3) the sort of harm at issue is death or serious physical injury; (4) in terms of its probability, the risk is low, but not insignificant; and (5) the costs of reducing the risk come primarily in the form of wealth, income, or inconvenience. Aside from large-scale construction projects like the building of bridges, dams, highways, stadiums, and skyscrapers, some other examples of activities that share these core features include: public utility projects like the generation and widespread delivery of electric power; machine- or chemical-based factory processes used by firms in manufacturing goods; clinical drug trials; and vaccination initiatives.

---

[23] Ex ante contractualism is a version of contractualist moral theory, which was introduced by T.M. Scanlon in 1982. *See* T.M. Scanlon, *Contractualism and Utilitarianism*, *in* UTILITARIANISM AND BEYOND 103-28 (Amartya Sen & Bernard Williams, eds.) (1982) (hereinafter "Scanlon, *Contractualism*"); T.M. Scanlon, WHAT WE OWE TO EACH OTHER (1998) (especially Chapter 5, pp. 189-247) (hereinafter "Scanlon, WHAT WE OWE"). In his 1982 article, Scanlon credits Rawls with first suggesting the possibility of a contractualist theory of normative ethics. *See* Scanlon, *Contractualism*, *supra* this note, at 103-104; *see also* Rawls, *supra* note __, at 15 ("For it is clear that the contractarian idea can be extended to the choice of more or less an entire ethical system, that is, to a system including principles for all the virtues and not only for justice. . . .Obviously if justice as fairness succeeds reasonably well, a next step would be to study the more general view suggested by the name 'rightness as fairness.'").

In cases of this type, what level of precaution is the actor morally required to exercise?   That is the question I will explore in this Article.[24] I should be clear on the particular kind of answer I am seeking. I seek a substantive decision rule that could be used in specific cases to determine whether the actor was morally obligated to take a particular precaution or to fix the precise amount of money the actor is morally required to invest in safety precautions.[25] Here are two examples of insufficiently specific answers that would fail this test: "The actor is morally required to take all precautions that would have been taken by a reasonably careful person under the circumstances," and "The actor is morally required to take precautions such that the resulting risk imposed is one it is fair to ask others to bear." These answers link the moral obligation to take precaution to a particular concept, i.e., reasonableness or fairness, but they fail to specify how, in a particular case, one would determine whether reasonableness or fairness requires taking some particular precaution.

I should offer two caveats before moving on. First, I will have relatively little to say here about assumption of risk, an issue that obviously affects the morality of risk impositions in particular cases. It is evident that, in appropriate circumstances, a person's free and informed consent to bear a particular risk to their own bodily integrity can render permissible an otherwise impermissible risk imposition.[26] It seems clear, though, that many cases we might be tempted to place under this category fail, upon close analysis, to involve consent made freely or with adequate information. Think, for example, of the unemployed worker who, afraid of being unable to meet her family's basic needs, accepts a factory job knowing that it car-

---

[24] In this Article, I do not purport to advance a *general* interpretation of the reasonable care standard. I mean only to address cases in which a repetitive, coordinated, firm-based activity with substantial social benefits imposes a risk of serious bodily harm on certain individuals (workers, consumers, or bystanders). I do not, for example, mean to address cases in which the private acts of individuals—e.g., riding a bicycle or walking a dog on a public road, making a campfire at a public campsite, mowing a front lawn, etc.—impose nontrivial risks of serious bodily harm on other individuals. *See* Keating, *Social Contract*, *supra* note ___, at 39-42 (distinguishing for purposes of tort liability between the "world of acts," a world of "'isolated, ungeneralized wrongs,'" and the "world of activities," a world "in which certain risks are the regular and routine 'incidents of certain well-known business'") (quoting Oliver Wendell Holmes, *The Path of the Law*, 10 HARV. L. REV. 457, 467 (1897)).

[25] Fried focuses her inquiry on the same sort of substantive decision rule. *See* Fried, *Limits*, *supra* note ., at 232 n.2 ("I set to the side procedural solutions, which do offer a clear alternative to substantive decision rules, consequentialist and nonconsequentialist alike.").

[26] *See, e.g.,* Tim Lewens, *Introduction* in RISK: PHILOSOPHICAL PERSPECTIVES 8-9 (Tim Lewens ed. 2007) ("It is plausible, for example, that risks which a person *chooses* to run can legitimately be much higher than risks which are *imposed* by some other agent.") (emphasis in original).

ries significant risks of serious injury. Does it make sense to say the worker has *freely* assumed the risk in that case? Perhaps not.[27]

Second, my inquiry here is primarily concerned with cases of advertent negligence, i.e., cases in which actors knowingly choose to take (or not to take) a particular precaution or adhere to (not to adhere to) a particular precautionary standard.[28]   The firm-based activities with which I am concerned tend to involve conscious, calculated decisions about whether to impose a particular level of risk or exercise a particular level of precaution. I set to the side cases involving inadvertent failures to take particular precautions, as such cases raise difficult questions about moral wrongdoing and culpability that I do not wish to broach here.[29]

B.   Barbara Fried's Challenge

Fried throws down the gauntlet to nonconsequentialist[30] risk theory in two recent articles that argue for essentially the same core claim, doing so

---

[27] *See* Simons, *supra* note __, at 1181 (setting aside issues of consent and assumption of risk in discussion of morality of risk imposition because "innumerable activities are tolerated in the contemporary world even though it is unrealistic to claim or expect that all those exposed to the risk (including bystanders and even children) subjectively 'consent' in any meaningful sense of the term"); *see also* Robert E. Keeton, Lewis D. Sargentich, & Gregory C. Keating, TORT AND ACCIDENT LAW: CASES AND MATERIALS 1257 (4th ed. 2004) ("[T]he risks and benefits of a consumer product accrue to the same party, the consumer, who can choose to avoid the risk. Employees, on the other hand, rarely have such a choice—they must bear the risk of occupational health hazards, while benefits accrue largely to employers and consumers.") (quoting Federal Respondent's Brief at 55, in Am. Textile Mfrs. Inst. v. Donovan, 452 U.S. 490 (1981)).

[28] On the distinction between advertent and inadvertent negligence, see RESTATEMENT (THIRD) OF TORTS: LIABILITY FOR PHYSICAL AND EMOTIONAL HARM § 3 cmt. k (2010).

[29] *See, e.g.,* Larry Alexander, *Negligence, Crime, and Tort: Comments on Hurd and Simons*, 76 B.U. L. REV. 301, 302 (1996) (stating belief that inadvertent negligence is not morally culpable); *cf.* Dov Waisman*, Negligence, Responsibility, and the Clumsy Samaritan: Is There a Fairness Rationale for the Good Samaritan Immunity?*, 29 GA. ST. U. L. REV. 609, 644-62 (2013) (arguing inadvertent samarital negligence need not involve moral culpability, but typically gives rise to moral responsibility for any resulting injury). *But see* Heidi Hurd, *The Deontology of Negligence*, 76 B.U. L. REV. 249, 270 (1996) ("If it is wrong to do more harm than good in the arena in which deontological maxims do not apply, then it would appear culpable to do an act under circumstances in which the discounted value of the harm that act will cause exceeds the costs of precautions that it would take to prevent that harm.").

[30] The major divide in normative moral theory is between consequentialism and nonconsequentialism. Consequentialists believe that the moral permissibility of an action depends entirely on the consequences of that action (typically, its consequences for the aggregate well-being of all affected individuals). *See, e.g.,* Shelly Kagan, NORMATIVE ETHICS 60 (1998). The most prominent consequentialist view is utilitarianism. *See id.* at 61-62. Nonconsequentialists, while not denying the relevance of consequences, believe that the moral permissibility of an action may be a function of factors other than and in addition to its consequences (such as whether it conforms with a specified norm, whether it respects the rights of all involved, etc.). *See id.* at 70-71. The most prominent nonconsequentialist view is deontology. *See id.* at 72-73.

in slightly different ways.[31] The claim is that nonconsequentialists have failed to come up with a coherent answer to the following question: under what circumstances is it morally permissible to engage in a socially beneficial activity that imposes a nontrivial risk of serious bodily harm on others?[32] In Fried's view, nonconsequentialist theorists have either dodged that question entirely, provided unacceptably vague answers (e.g., "so long as one takes all reasonable precautions"), provided answers that apply only in very unusual circumstances (e.g., cases involving virtually certain consequences to identifiable victims), or provided answers that so closely resemble an interpersonally aggregative approach that they do not warrant the nonconsequentialist label.[33]

When it comes to the problem of specifying when it is morally permissible to engage in socially beneficial conduct that imposes a significant risk of serious harm, consequentialists tend to clash with nonconsequentialists over the question of *interpersonal aggregation*. Most versions of consequentialism are interpersonally aggregative in the sense that the permissibility of a particular action depends on how it will affect the aggregate well-being of all individuals whose well-being it will affect. As Fried puts it, an aggregative theory is one that "rank-order[s] alternative principles for action at least in part based on their aggregate expected benefits (costs), summed across all potentially affected individuals."[34] This kind of approach tends to commit consequentialists to the position that, as Fried puts it, "a risk of serious harm to one person can be justified by small benefits to the many."[35]

Nonconsequentialists, on the other hand, believe that we cannot necessarily justify imposing a significant risk of serious harm on the few for the sake of providing small benefits to the many, even if the aggregate benefit of imposing the risk exceeds its aggregate expected cost. For nonconsequentialists, even if imposing a risk delivers a net aggregate benefit to all affected individuals, it might still be wrong to impose it.[36]

This dispute over interpersonal aggregation lies at the heart of Fried's challenge to nonconsequentialist risk theory. The essence of her

---

[31] While *Limits* emphasizes the failure of deontologically-oriented *tort theory* (in particular, corrective justice theory) to offer a coherent, nonaggregative alternative to cost-benefit analysis for adjudging the permissibility of risk impositions, *Contractualism* focuses on the failure of contractualist *moral theory* to provide such an alternative.

[32] *See* Fried, *Limits*, *supra* note __, at 232 ("[H]ave critics of aggregation offered an analytically coherent *substantive* decision rule for regulating risky conduct that does not itself boil down to some form of aggregation? The short answer is: I do not think so.").

[33] *See id.* at 231-33, 253-59.

[34] Fried, *Contractualism*, *supra* note __, at 39-40.

[35] *Id.* at 39.

[36] This follows directly from the nonconsequentialist tenet that an act's moral permissibility can depend on factors other than its consequences for well-being. *See* Kagan, *supra* note ____, at 70-71.

claim is that it is all well and good for nonconsequentialists to insist that morality in some cases requires taking more than the marginally cost-justified level of precaution, but that that position is empty (or, at the least, unconvincing) unless nonconsequentialists can say with some specificity when pressing precaution beyond the point of cost-justification is required and how much further precaution is required in such cases.[37]

Fried maintains that nonconsequentialists have yet to offer "viable" answers to these questions. "For these purposes," Fried says,

> I mean 'viable' in a very undemanding sense: have nonconsequentialists supplied criteria that, as an operational matter, are capable of differentiating among different forms of risky conduct? I do not reach the further question whether such criteria dominate [cost-benefit calculus] on normative or practical grounds.[38]

Fried's challenge to nonconsequentialist risk theory goes deeper than this, though. She thinks it is no accident that, in her view, there is no extant alternative to interpersonal aggregation when it comes to the difficult case-by-case task of determining whether a particular risk imposition is morally permissible. Fried believes it is *impossible* to accomplish that task without summing costs and benefits across persons:

> Moreover, I will suggest, the effort to come up with nonaggregative principles to handle the problem of risk *must* fail—that the problem of risk, by its nature, can be managed *only* with the sorts of interpersonal tradeoffs that the contractualist enterprise is foundationally committed to rejecting.[39]

Summing costs and benefits across persons is, in Fried's view, the only way to negotiate the inevitable conflict between liberty and security that arises in connection with socially beneficial conduct that imposes a nontrivial risk of harm. In Fried's view, without interpersonal aggregation, a theory forfeits the capacity to explain why morality allows us to engage in the wide range of significantly risky but socially beneficial activities that characterize life in modern, industrialized societies.

Though Fried does at one point expressly consider the possibility that nonconsequentialist moral theory dictates taking a level of precaution that exceeds the level of precaution dictated by efficient care,[40] she fails to

---

[37] *See* Fried, *Limits*, *supra* note __, at 232 ("[H]ave critics of aggregation offered an analytically coherent *substantive* decision rule for regulatin risky conduct that does not itself boil down to some form of aggregation? The short answer is: I do not think so.").

[38] Fried, *Limits*, *supra* note __, at 234.

[39] Fried, *Contractualism*, *supra* note __, at 40 (emphasis in original); *see also* Fried, *Limits*, *supra* note ___, at 238-39.

[40] *See* Fried, *Contractualism*, *supra* note __, at 60-61.

identify the economic feasibility principle, or anything resembling it, as a possible precautionary alternative. For example, she observes,

> What it means to take 'reasonable precautions' is described in many different ways in the nonconsequentialist literature: One should behave in a fashion that is not negligent, that respects the legitimate interests of others to be free from harm, that is not wrongful, unreasonable or unjustifiable, that gives people what they are due; one should adopt safety measures that will substantially lower the risk at relatively low cost. But it is unclear whether these different verbal formulations imply different standards of conduct, and whether any of them differs significantly from the optimal level of precaution dictated by aggregative techniques.[41]

Here Fried alleges that the nonconsequentialist literature on risk regulation has nothing to offer in the way of an alternative precautionary standard other than vague reformulations of the injunction to exercise reasonable precaution.

This allegation is baseless. At least since 2001, when Keating published a paper[42] identifying the economic feasibility principle as the precautionary standard emerging from Kantian social contract theory, the nonconsequentialist literature has indeed featured a determinate precautionary alternative to efficient care. Fried's failure to even once take note of, much less discuss, the possibility of a precautionary standard that is keyed to the survival of the underlying activity is puzzling, particularly in light of the radical nature of her critique, i.e., that there is no viable alternative on the table.[43] But is that oversight fatal to her claim that no precautionary standard can yield plausible, determinate results in ordinary types of cases without relying on interpersonal aggregation?

I believe it is. Building on Keating's foundational work,[44] I describe an alternative precautionary standard—the *individualized feasibility principle*—that avoids reliance on interpersonal aggregation and delivers determinate, morally plausible results in ordinary types of cases, thereby answering Fried's challenge.[45]

---

[41] See id. at 61-62; see also Barbara H. Fried, What Does Matter? The Case for Killing the Trolley Problem (or Letting It Die), 62 PHIL Q. 505, 512-13 (2012).

[42] *See* Keating, *Social Contract*, *supra note* __, at 22-51. In a series of papers that followed, Keating developed his account of the social contract rationale underlying the feasibility and safety norms found in American risk regulation. *See* Keating, *Pressing Precaution*, *supra* note__ , at 684-748; Gregory C. Keating, *Irreparable Injury and Extraordinary Precaution: The Safety and Feasibility Norms in American Accident Law*, 4 THEORETICAL INQ. IN L. 1 (2003); Gregory C. Keating, *Pricelessness and Life: An Essay for Guido Calabresi*, 64 MD. L. REV. 159 (2005).

[43] *See* Fried, *Limits*, *supra* note __, at 232-34.

[44] *See* Keating, works cited *supra* note ___.

[45] Aaron James has responded to Fried's challenge in a general way, arguing that an appropriately specified ex ante version of contractualism need neither forbid interpersonal

II. Answering the Challenge: The Individualized Feasibility Principle

The individualized feasibility principle represents a synthesis of two precautionary standards: the *economic feasibility principle*, which has long been part of American law, and the *individual risk principle*, a novel risk regulation principle I introduce in this Article. In any given case, each of these principles will dictate a distinct level of investment in safety precaution. In some cases, the feasibility principle will dictate a greater investment in safety than will the individual risk principle. In others, the opposite will be true.   In any given case, the IFP directs the actor imposing the risk to comply with whichever of these two principles dictates the *lesser* investment in safety.

A. The Economic Feasibility Principle

The economic feasibility principle holds as follows: When engaging in a socially beneficial activity that imposes a significant risk of serious bodily harm on certain individuals, the actor imposing the risk should take all economically feasible precautions The term "all economically feasible precautions" refers to the set of safety precautions that, in the case of a particular risk, yields the maximum reduction in accident costs (deaths and serious injuries) consistent with the long-term viability of the risky activity in question.[46]

An economic feasibility standard has been part of American law for decades. Most prominently, certain portions of the Occupational Safety and

---

aggregation altogether nor succumb to a perilous slide into morally implausible forms of such aggregation. *See* Aaron James, *Contractualism's (Not So) Slippery Slope*, 18 LEGAL THEORY 263 (2012).   I here take up the task that James expressly declines to tackle: "answer[ing] Fried's doubts about whether, in a vast range of ordinary public policy choices, ideas of 'reasonable cost' or 'adequate opportunity to avoid' can be operationalized other than in aggregative cost-benefit terms." *Id.* at 289; *see also* John Oberdiek, The Morality of Risking: On the Normative Foundations of Risk Regulation 187 n.36 (2003) (unpublished Ph.D. dissertation, University of Pennsylvania) (on file with author and Proquest Learning and Information Company) (offering a nonaggregative theory of the morality of risk imposition, but "postpon[ing] for another occasion" the "less philosophical project" of "[d]elineating concrete standards of permissible risking" that such a theory would generate).

[46] *See* Keating, *Pressing Precaution*, *supra* note __, at 685 (explaining that, under the economic feasibility standard, "[c]ost-justified risks are eliminated, so long as their elimination is compatible with the long-term flourishing of the activity at issue, and significant risks remain if their elimination would threaten the survival of the activity"); *see also* Keating, *Social Contract*, *supra* note __, at 46 ("When risks are significant, and when life and limb are pitted against wealth and income, then, we should reduce risks to the point where they are either no longer unreasonable, or where further reduction would jeopardize the continuation of the activity itself.").

Health Act of 1970[47] implement the feasibility principle: OSHA requires employers to reduce workplace risks posed by toxic materials or harmful physical agents to the extent feasible,[48] i.e., to the maximum extent possible without jeopardizing the long-term survival of the industry.[49] Moreover, a feasibility standard has long been part of American nuisance law.[50]

The economic feasibility principle can account for the intuition that, in many cases, risk-creators are morally obligated to press precaution beyond the point of marginal cost-justifiability. Typically, it will dictate a significantly greater level of precaution than aggregative standards like efficient care.[51] However, the economic feasibility principle stops short of placing implausibly onerous moral constraints on those wishing to engage in socially beneficial but significantly risky activities. The principle avoids what John Broome has termed "moral gridlock"[52] because it allows socially useful activities that are sure to result in accidental deprivations of life and limb to go forward, provided significant risks are reduced to the maximum extent possible without undermining the activity at issue. The economic feasibility standard sets the required level of precaution below (but only *just* below) the point at which engaging in a risky activity becomes so costly or burdensome as to dissuade rational actors from engaging in the activity altogether, thus jeopardizing its continuation.

Two components of the feasibility principle warrant further explanation: the notion of a socially beneficial activity and the notion of an activity's long-term viability.

---

[47] Occupational Safety and Health Act of 1970, 84 Stat. 1590 (1970).

[48] *See* 29 U.S.C. § 655(b)(5) ("The Secretary, in promulgating standards dealing with toxic materials or harmful physical agents under this subsection, 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 dealt with by such standard for the period of his working life.") (emphasis added).

[49] *See* Am. Textile Mfrs. Inst. v. Donovan, 452 U.S. 490, 530 n.55 (1981) (holding cotton dust standard economically feasible because "the industry will be able to maintain long-term profitability and competitiveness); United Steelworkers of Am., AFL-CIO v. Marshall, 647 F.2d 1189, 1265 (D.C. Cir. 1980) ("A standard is feasible if it does not threaten massive dislocation to, or imperil the existence of, the industry.") (internal quotation marks omitted)

[50] *See* Restatement (Second) of Torts § 826(b) (1979).

[51] *See, e.g.,* Keating, *Pressing Precaution*, at 684-85 (noting feasibility principle generally tolerates less risk than the cost-justification standard); Keeton et al., *supra* note __, at 1237-41 (discussing a continuum of precautionary standards in which the feasibility standard is considered less tolerant of risk than the cost-benefit balancing standard). However, if accident costs are largely externalized, it seems at least theoretically possible for an efficient expenditure on precaution not to be feasible (i.e., for the maximum feasible safety expenditure to be sub-efficient).

[52] *See* John Broome, *Trying to Value a Life*, 9 J. PUB. ECON. 91 (1978).

1. Socially Beneficial Activity

What makes a risky activity "socially beneficial"? And why restrict my inquiry here to such activities? As noted above, the type of activities I have in mind are coordinated, repetitive, actuarially large activities typically undertaken and subsidized by firms or government agencies.[53] Examples include large-scale construction projects like the building of bridges, dams, highways, stadiums, and skyscrapers; public utility initiatives like the generation and delivery of electric power, gas, and potable water; machine- or chemical-based factory processes used by firms in manufacturing goods; clinical drug trials; and vaccination initiatives.

A few observations about this class of activities are in order. This category excludes risky social practices comprising the private acts of individuals, e.g., driving a car on a public road, using a conventional gas stove in a home kitchen, walking a large dog on a leash on a public sidewalk. The reason for excluding such practices from my inquiry is that their justification is usually thought to rest on the risk-bearer's reciprocal right to directly engage in the practice at issue, and thereby impose the same risk on others.[54] By contrast, the activities with which I am concerned here typically impose risks on persons who do not themselves engage in the activity at issue and who cannot, therefore, be said to benefit from a reciprocal right to themselves act in a manner that imposes the associated risks on others. It will not do, for example, for a large company to justify the risks its stadium construction project imposes on passersby by saying, "Well, when a passerby builds her stadium, she may impose the identical risk on us." I limit my inquiry here to such non-reciprocal risks.

The category of socially beneficial activities is also meant to exclude commercially profitable activities that are widely criminalized, e.g., the manufacture and distribution of illicit drugs.[55]

---

[53] *See* Keating, *Social Contract, supra* note __, at 40-41 (contrasting the world of *acts*—in which risk impositions are "discrete one-shot events" and the "typical actor is an individual or small firm"—with the world of *activities*—in which "the typical injury arises not out of the diffuse and disorganized acts of unrelated individuals or small firms, but out of the organized activities of firms that are either large themselves, or small parts of relatively well-organized enterprises" and in which safety costs can typically be spread over many individuals); *see also* Holmes, *supra* note __, at 467 ("[T]he torts with which our courts are kept busy today are mainly the incidents of certain well known businesses. They are injuries to person or property by railroads, factories, and the like. The liability for them is estimated, and sooner or later goes into the price paid by the public.").

[54] *See, e.g.,* Keating, *Pressing Precaution*, *supra* note __, at 678 ("The right to impose risks *on* others can justify the imposition of equal risks on us *by* others, because, for example, we may each gain more than we lose from having to bear the risks created by the presence of other cars on the road.").

[55] I follow Fried in restricting the scope of the inquiry in this way. *See* Fried, *Limits*, *supra* note __, at 234-35.

So much for what a socially beneficial activity is *not*. As for a positive definition, I would offer the following necessary (though not sufficient) condition: an activity is socially beneficial only if, in the long term, the activity can be routinely performed in such a way that its aggregate expected benefit to individuals' well-being exceeds its aggregate expected cost to individuals' well-being.[56] Thus, for an activity to be socially beneficial in this "on net" sense, there has to exist a level of precaution such that the gross aggregate benefit the activity delivers exceeds the sum of its aggregate costs (including accident costs and accident prevention costs).[57] Notice that the mere fact that an activity can, in a modern, free-market economy, be undertaken profitably over a relatively long period of time does not guarantee that it satisfies this condition. The production and sale of tobacco products, for example, has been commercially profitable for centuries, but it is questionable whether, given tobacco's harmful health effects, that activity has yielded a net aggregate benefit over the course of its existence.

For purposes of this Article, I need not draw a precise line between those activities that count as socially beneficial in the relevant sense and those that do not. It is sufficient to make clear that, for the most part, I am concerned with non-criminal, commercially profitable activities that, at some level of safety precaution, are capable of yielding a net benefit to aggregate well-being. Following Fried, I take these to be *prima facie* socially productive activities like those described above, activities as to which there is a broad consensus that the activity is worth preserving.[58] The question I explore in this Article is: assuming a risky activity is worth preserving, how safe does morality require that it be made?

2. Long-Term Viability of Risky Activity

The economic feasibility principle says to reduce a given risk to the maximum extent possible without threatening the long-term viability of the

---

[56] It is critical to note that, although the notion of aggregate net benefit provides a useful way of understanding what it means for an activity to be socially beneficial, it does not carry any normative weight in my analysis. The mere fact that an activity is capable of delivering a net benefit to aggregate expected well-being does not guarantee that it is morally permissible to engage in the activity, nor does it entail any claim concerning the level of precaution an actor engaging in the activity is morally required to exercise. Thus, my reliance on the notion of aggregate net benefit to help specify what is meant by a "socially beneficial activity" does not constitute a betrayal of my commitment to providing a nonaggregative interpretation of reasonable precaution.

[57] *See* Fried, *Contractualism*, *supra* note __, at 61 ("If, however much we spend on safety, the sum of lives lost plus safety costs will be greater than the social benefits, it doesn't make sense to go ahead with the project at all.").

[58] *See id.* at 41.

activity giving rise to the risk.[59] What exactly does that mean? Exactly how does one determine how much expenditure on safety precaution a given activity can tolerate without being undermined?

The most concrete answers to these questions come from OSHA, its corresponding regulations, and the cases interpreting them. One oft-quoted opinion serves as a useful starting point:

> A standard is feasible if it does not threaten 'massive dislocation' to, or imperil the existence of, the industry. No matter how initially frightening the projected total or annual costs of compliance appear, a court must examine those costs in relation to the financial health and profitability of the industry and the likely effect of such costs on unit consumer prices. . . . T[]he *practical question is whether the standard threatens the competitive stability of an industry*, or whether any intra-industry or inter-industry discrimination in the standard might wreck such stability or lead to undue concentration.[60]

In a given case, how do OSHA regulators determine whether the imposition of a particular standard threatens the competitive stability of an industry? OSHA approaches this question in the first instance by conducting an industry-by-industry analysis. For each industry, OSHA determines the percentage of the industry's revenues and profits that the costs of complying with the standard represent.[61] On this point, one OSHA report offers the following guidance:

---

[59] Of course, the economically feasible level of precaution must necessarily be technologically feasible, i.e., must be achievable within the limits of available technology. In the context of OSHA health regulations, this means that "the typical firm will be able to develop and install engineering and work practice controls that can meet the PEL [permissible exposure limits] in most of its operations." United Steelworkers of Am., AFL-CIO v. Marshall, 647 F.2d 1189, 1272 (D.C. Cir. 1980). Obviously, requiring a degree of safety that is technologically unattainable is inconsistent with the survival of the activity at issue.

[60] *Id.* at 1265 (quoting Industrial Union Dep't, AFL-CIO v. Hodgson, 499 F.2d 467, 478 (D.C. Cir. 1974)) (emphasis added) (internal citations omitted); *see also* Am. Textile Mfrs. Inst. v. Donovan, 452 U.S. 490, 531 (1981) (describing OSHA's determination that proposed standard for workplace cotton dust exposure was economically feasible because "although some marginal employers may shut down rather than comply, the industry as a whole will not be threatened by the capital requirements of the regulation") (internal quotations and citations omitted); Occupational Exposure to Hexavalent Chromium; Final Rule, 71 Fed. Reg. 10099-10385, Supplementary Information, Section VIII.E, available at: https://www.osha.gov/pls/oshaweb/owadisp.show_document?p_table=FEDERAL_REGISTER&p_id=18599 ("OSHA's obligation is not to determine whether any plants will close, or whether some marginal plants may closer earlier than they otherwise might have, but whether the regulation will eliminate or alter the competitive structure of the industry.").

[61] *See, e.g.*, *id.* (describing results of industry-by-industry analysis of economic impact of proposed permissible exposure limit for hexavalent chromium).

> [W]hile there is no hard and fast rule, in the absence of evidence to the contrary OSHA generally considers a standard economically feasible when the costs of compliance are *less than one percent of revenues*. Common-sense considerations indicate that potential impacts of such a small magnitude are unlikely to eliminate an industry or significantly alter its competitive structure particularly since most industries have at least some ability to raise prices to reflect increased costs. . .There is an enormous variety of year-to-year events that could cause a one percent increase in a business's costs, e.g., increasing fuel costs, unusual one-time expense, changes in costs of materials, increased rents, increased taxes, etc.[62]

In at least one instance, OSHA concluded that where the costs of complying with a particular standard come to less than *both* one percent of an industry's revenues *and* ten percent of its profits, implementation of the standard would not threaten the competitive stability of the industry.[63] The reasoning here seems to be that if an industry is routinely able to absorb revenue and profit fluctuations within certain limits without seeing its competitive stability undermined, then a regulatory standard whose implementation would cause a revenue or profit fluctuation within those same limits would not threaten the industry's competitive stability.[64]

Where a particular industry's compliance costs significantly exceed the threshold levels designated by OSHA (which are typically stated as a percentage of the industry's revenues and profits), OSHA determines on an industry-by-industry basis whether complying with a particular standard will threaten an industry's competitive stability. For example, in analyzing the economic feasibility of proposed permissible exposure limits (PELs) for hexavalent chromium, OSHA expressed concern about how a PEL of 1 [mu] g/m^3 would affect the electroplating industry:

> OSHA is concerned about the economic feasibility of the standard for electroplating at a PEL of 1. At this lower PEL, costs of the standard represent 2.7 percent of revenues and 65 percent of profits. . . . It seems unlikely that a price increase of 2.7 percent would eliminate the industry entirely. OSHA has concluded, however, that the costs associated with such a PEL could alter the competitive structure of the industry. OSHA has concluded this because these costs substantially exceed the average nomi-

---

[62] *Id.* (emphasis added). The sort of "standard" at issue in OSHA health regulations is one that identifies a permissible exposure limit ("PEL") for a particular toxic substance.

[63] *See id.* ("The record does not contain evidence that any of the affected industries for which OSHA found that the costs of complying with the standard will be less than both one percent of prior revenue and ten percent of prior profits will in fact be threatened by the standard. . . .[C]ost changes of less than one percent are routinely passed on and impacts that are less than 10 percent of profits have not been shown to be likely to affect the viability or competitive structure of any of the industries affected by this standard.").

[64] See id.

nal price increases in the industry, and the reasons for these nominal price increases—increases in the cost of labor and energy, for example—will continue. Thus a price increase that would assure continued profitability for the entire industry would require almost tripling the annual nominal price increase. . . .That would represent a significant real price increase that might not be passed forward, particularly by older and less profitable segments of the industry.[65]

The touchstone of OSHA's analysis here seems to be ensuring continued profitability,[66] which leads naturally to a focus on demand elasticity, i.e., the extent to which demand for the industry's product or service will be affected by fluctuations in price. The less elastic the demand, the more an industry can pass on compliance costs to consumers without suffering a reduction in demand and a consequent reduction in profitability. Though, in the above case, OSHA concluded that "demand for electroplating services is relatively inelastic," making possible some degree of cost pass-through, it also concluded that the costs of compliance with the proposed standard (2.7% of revenues) were too large to be entirely passed on to consumers in the form of real price increases. The implication seems to be that compliance with the proposed standard would mean diminishing electroplating firms' profits so much as to effectively make electroplating *an unprofitable enterprise*. That, OSHA implies, is what is meant by threatening the "competitive stability" or altering the "competitive structure" of an industry.

B. The Individual Risk Principle

The individual risk principle represents an individualized, nonaggregative version of the efficient care principle; it is efficient care writ *small*. Whereas efficient care says to minimize the sum of aggregate accident costs and aggregate accident prevention costs, the individual risk principle says to minimize the sum of each individual risk-bearer's expected accident costs and each individual precaution cost-bearer's accident prevention costs. Both principles are marginalist in nature. Whereas efficient care says to invest in safety until the point at which further expenditure would reduce aggregate expected well-being on net, the individual risk principle says to invest in safety precaution until the point at which further expenditure would decrease each cost-bearer's well-being by more than it would in-

---

[65] See id.

[66] *See* Am. Textile Mfrs. Inst. v. Donovan, 452 U.S. 490, 530 n.55 (1981) (holding cotton dust standard economically feasible because "the industry will be able to maintain long-term profitability and competitiveness").

crease each risk-bearer's expected well-being.[67] The individual risk principle therefore contemplates an individualized "single owner"[68] approach to risk regulation, as it asks how much a person who "owned" both a pro rata share of safety costs and pro rata share of expected accident costs would rationally be willing to invest in safety precaution.

The individual risk principle sets a moral ceiling beyond which it is arguably unreasonable to require investment in safety precaution. The basic idea is this: at least where A and B both benefit directly and significantly from the risky activity at issue, one cannot reasonably expect A to accept a setback to her well-being (in the form of an increase in the monetary safety costs she must bear) in order that B can avoid a *smaller* setback to his expected well-being (in the form of an incremental increase in the risk of death or serious bodily harm he must bear). To do so would be to ask A to spend more on safety precautions than B herself would plausibly be willing to spend to protect herself from the risk at issue (were she in A's economic situation).[69]

The individual risk principle is, as its name suggests, keyed to *individual risk*, i.e., the risk borne by each individual risk-bearer.[70] Individual risk is ordinarily expressed as the probability (i.e., a 1 in x chance) that an individual will, over the course of a particular time period (a year, an average lifetime, etc.), suffer a particular harm (death, serious injury, etc.). For example, in the bridge construction hypothetical discussed in the Introduction, each of the 2,000 bridge workers is exposed to a 1 in 2,000 individual risk of a fatal fall over the course of the five-year project if the safety scaffolding is used and a 1 in 1,000 risk if it is not. Notice that individual risk need not (and typically does not) vary with the numbers of persons exposed to the risk. If the city suddenly decided it wanted to double the number of bridge workers, each of the 4,000 workers would still plausibly face a 1 in 2,000 risk of a fatal fall over the life of the project (though the risk might be

---

[67] These two principles will dictate the identical level of precaution only where risk-bearers and cost-bearers are perfectly coextensive and homogeneous in all relevant respects. *See infra* note _____.

[68] *See, e.g.,* Kenneth W. Simons, *Deontology, Negligence, Tort, and Crime*, 76 B.U. L. Rev. 273, 282 (1996) (describing a " 'single-owner' conception" of cost-benefit analysis "which asks what the actor. . .would do if he owned all the resources in question and would therefore internalize all the costs and benefits of the decision").

[69] As discussed in greater detail *infra* in Part III.F, the individual risk principle cap on safety investment is most plausible from a moral point of view in cases where the risk-bearers directly and significantly benefit from the risky activity, e.g., as employees or consumers. In cases where those exposed to the risk at issue do not directly benefit from the risk-imposing activity, it is somewhat more questionable whether the individual risk principle should set a moral ceiling on safety investment.

[70] On the distinction between individual risk and population risk, see, e.g., Matthew D. Adler, *Against "Individual Risk": A Sympathetic Critique of Risk Assessment*, 153 U. PENN. L. REV. 1121, 1126 (2005); *see also* Simons, *supra* note __, at 1218-21.

reduced if the increased number of workers resulted in the bridge being completed in a shorter period of time).

Individual risk is to be contrasted with *population risk*, which represents the number of individuals expected to suffer a particular harm (death, serious injury, etc.) in a specified time period.   Thus, in the bridge hypothetical, using the safety netting makes the population risk fall from two expected deaths to one. Population risk typically *does* vary with the number of persons exposed to the risk.[71] If the number of workers on the bridge doubles and the individual risk to each worker remains the same, the number of expected deaths will double as well.

As discussed in more detail below, because the individual risk principle is keyed to individual risk rather than population risk, the level of precaution it requires in a given case is not *directly* dependent on the number of persons who bear the risk at issue or on the number of persons bearing the cost of reducing the risk. Indeed, the individual risk principle entirely rejects interpersonal aggregation and avoids summing costs and benefits across persons altogether. It does, of course, contemplate a bilateral interpersonal *comparison*: for a given expenditure on safety, the individual risk principle compares the monetary cost of the precaution to each individual cost-bearer against the resulting increase in safety experienced by each individual risk-bearer. But the individual risk principle does not sum precaution costs (or savings) across all affected cost-bearers, nor does it sum accident costs (or savings) across all affected risk-bearers.

Another thing to notice about the individual risk principle is that, unlike the feasibility principle, its application requires determining the number of distinct individuals who shoulder the costs of mitigating a particular risk. (Without doing so, it will be impossible to determine the monetary burden borne by each bearer of precaution costs.) Admittedly, in some circumstances, this may be rather difficult to do. For example, a company engaged in a risky activity may distribute its safety precaution costs among a vast group of consumers, shareholders, or some combination of the two, and determining the exact share of precaution costs each individual shareholder or consumer ends up bearing may be burdensome.   A further complication is that, on account of the diminishing marginal utility of money, the degree to which a specified monetary loss reduces the well-being of a particular individual will plausibly depend on the individual's existing wealth, which will vary significantly from person to person.[72]

---

[71] It is possible to conceive of situations in which this is not the case, however. *See* Reibetanz, *supra* note __, at 302 (discussing "Unexploded Mine" hypothetical in which, as the number of persons exposed to the risk grows, the population risk remains the same (a 100% chance of one death) and the individual risk falls).

[72] In this Article, I largely ignore this complication, but hope to address it in future work. For my purposes here, I assume that all individuals potentially affected by a particular risk

C. The Individualized Feasibility Principle

When the economic feasibility principle is conjoined with the individual risk principle, a bipartite standard emerges: the individualized feasibility principle. As noted above, the IFP directs a risk-creating actor to stop investing in safety precaution only when one of the following two points has been reached: (i) the point at which further expenditure on safety would threaten the long-term survival of the activity (economic feasibility); or (ii) the point at which further expenditure would reduce the well-being of each bearer of precaution costs by more than it would increase the expected well-being of each risk-bearer (individual risk).

According to the IFP, the risks posed by a socially beneficial activity should be reduced to the maximum extent consistent with the activity's long-term survival, unless doing so would require each cost-bearer to invest more in safety precaution than a risk-bearer would rationally have been willing to invest in her own safety, were she in the economic situation of a cost-bearer. Equivalently, the IFP requires each cost-bearer to make any investment in risk-reduction that she would rationally have been willing to make were it her own bodily integrity at risk, up to the point at which further investment would place the long-term survival of the activity in jeopardy.

Thus, the IFP implicitly rejects the notion that the only good comparable in value to the reduction of significant risks of bodily harm is the preservation of the major, productive activities that define modern life.[73] By capping precaution at the point dictated by the individual risk principle,

---

imposition have identical levels of wealth and thus experience the identical loss (or gain) in well-being as the result of a given monetary loss (or gain). Of course, aggregative welfarist theories must confront the same difficulty, as the aggregate reduction to well-being associated with a particular monetary cost will, due to the diminishing marginal utility of money, obviously depend on the existing wealth of each individual responsible for bearing the cost. *See* Matthew D. Adler, Cost-Benefit Analysis and Distributional Weights: An Overview 10-15 (August 20, 2013) (unpublished manuscript) (available at http://www.ssrn.com/abstract =2313388) (explaining how distributionally weighted cost-benefit analysis can take account of varying levels of individual wealth and the diminishing marginal utility of money).

[73] Keating sometimes suggests that one individual's certain monetary loss (or gain) is fundamentally incomparable to—and cannot fairly be traded against—an increase (or reduction) in the low but significant risk of death or serious bodily harm borne by another individual. *See, e.g.,* Keating, *Pressing Precaution*, *supra* note __, at 664-74 (discussing comparability of trivial monetary losses and gains with devastating bodily injuries). As noted in the main text, I disagree with this view. One can coherently subscribe to the comparability of monetary losses (or gains) and changes in one's risk of serious bodily harm without also subscribing to the view that a trivial monetary gain, summed across a vast number of individuals, can justify the imposition of a greater risk of death on a distinct and much smaller group of individuals. Keating sometimes seems to me to conflate comparability and aggregability in this way.

the IFP implies that incremental changes in the monetary costs borne by each individual who subsidizes safety precaution *can* meaningfully be traded against incremental changes in the risks of serious bodily harm borne by each individual risk-bearer. If a given reduction in risk is sufficiently costly for each cost-bearer and sufficiently small for each risk-bearer, the IFP holds it morally indefensible to require the reduction, even if doing so would not threaten the survival of the underlying activity. However, by simultaneously capping precaution at the economic feasibility point, the IFP accommodates the intuition that, in the context of socially beneficial activity, the reasonable level of precaution can never be one which is so stringent as to make the activity unviable. The IFP caps investment in safety short of the level required by the individual risk principle if this is necessary to preserve the underlying socially beneficial activity.

1. Does the IFP Lead to Morally Implausible Results?

One implication of the IFP is that the level of precaution morally required with respect to a particular risk is not directly dependent on the number of persons exposed to that risk. Unlike the efficient care principle, neither the economic feasibility principle nor the individual risk principle is directly sensitive to the size of the population exposed to the risk.[74] The underlying thought is that the degree of protection to which each risk-bearer is morally entitled should not directly depend on how many other persons also happen to be exposed to the risk.[75] Is this defensible? Should reasonable precaution be keyed to both the individual risk borne by each risk-bearer *and* the number of risk-bearers?[76]

In many cases, it seems somewhat implausible that the necessary level of precaution should be a direct function of the number of persons exposed to the risk. Should a bus driver use greater care when the bus is full of passengers than when it is empty save for one or two?[77] Should a high-rise apartment building be constructed to keep its occupants safer than occupants of a single-family dwelling, simply because more people are at

---

[74] *See* Adler, *Against Individual Risk*, *supra* note __, at 1240-41 (noting many regulatory standards in the United States are insensitive to population size and keyed to individual risk only).

[75] *See* Oberdiek, Morality of Risking, *supra* note __, at 144 ("[N]o difference can be justified in the relative care taken in constructing buildings that serve differing numbers of people. Since each person stands the same probability of suffering a like harm—since each person in either kind of structure runs the same risk—each person is entitled. . .to the same level of precaution or the same standard of care.").

[76] *See* Oberdiek, Morality of Risking, *supra* note __, at 144 ("[U]nder the Hand test, greater care must be taken in the construction of apartment high-rises than single-family homes since so many more could suffer significant harm if, say, the apartment building's foundation faltered.").

[77] *See* Simons, *supra* note __, at 1174.

risk in the high rise?[78] A number of commentators have answered questions like these in the negative.[79]

On the other hand, a standard's lack of direct sensitivity to population size might imply that whether 10 workers or 10,000 workers are exposed to a particular workplace risk, each worker's risk must be reduced by the identical amount.[80] If a 1 in 1,000 risk of death falls on each individual worker, the risk can be expected to result in 10 deaths in the latter case, but only a 1% chance of a single death in the former. To some commentators, it has seemed implausible to require no greater reduction of each worker's risk in the 10,000-worker case than in the 10-worker case.[81]

Notice, however, that under the IFP, the requisite level of precaution may *indirectly* depend on the size of the population exposed to the risk. For example, as practiced by OSHA, the economic feasibility principle specifies a particular permissible exposure limit for each worker that must be satisfied by all firms within a particular industry.[82] It seems likely that the more at-risk employees a firm has, the more it will need to invest to satisfy this uniform standard. For example, the total safety investment necessary to reduce each worker's individual risk of death from 1 in 1,000 to 1 in 750 would almost certainly be greater in the case of a 10,000-worker factory than in the case of a 10-worker factory.[83] Thus, under the economic feasibility principle, the greater the number of persons at risk, the more each cost-bearer will likely be required to invest in safety precautions. It seems to me that the intuition driving Adler and other commentators critical of precautionary standards keyed to individual risk is that, when a risk imposition can be expected to cause 10 deaths, morality compels us to *do more* than when a risk imposition merely creates a 1% chance of a single death, even if both risk impositions are identical with respect to the individual risk borne by each risk-bearer. Insofar as it tends to require a greater total (and

---

[78] See John Oberdiek, The Ethics in Risk Regulation: Toward a Contractualist Re-Orientation, 36 RUTGERS L.J. 199, 203-204 (2004).

[79] *See* Simons, *supra* note ___, at 1174; Oberdiek, *Ethics in Risk Regulation*, *supra* note _, at 203-204; *cf.* Scanlon, WHAT WE OWE, *supra* note ___, at 236.

[80] For example, in setting permissible exposure limits (PELs) for hazardous substances under the economic feasibility standard, OSHA designates a PEL that applies regardless of the number of workers at risk in a given firm, industry, or group of industries. *See, e.g.,* Am. Textile Mfrs. Inst., Inc. v. Donovan, 452 U.S. 490, 501-504 (1981) (describing OSHA's setting PEL for cotton dust exposure without regard for population size). Of course, achieving the identical reduction in each worker's individual risk may well require a different per-worker safety expenditure depending on the number of workers exposed to the risk.

[81] *See, e.g.,* Adler, *Against Individual Risk*, *supra* note __, at 1241 (calling it "morally arbitrary" to require identical reduction in individual risk without regard to number of risk-bearers).

[82] *See supra* note 80.

[83] The per-worker safety investment may be greater or smaller in the 10,000-worker case, depending on the circumstances.

pro rata) safety expenditure when more people are at risk, the economic feasibility principle accommodates this intuition.

Under the individual risk principle, the relationship between the requisite level of precaution and the number of persons at risk is more complicated. Whereas the economic feasibility principle requires the identical reduction in individual risk without regard to the number of risk-bearers, under the individual risk principle, the requisite reduction in individual risk may itself be affected by an increase in either the number of cost-bearers (since this will necessarily reduce each cost-bearer's share of precaution costs) or the number of risk-bearers (since this will likely raise the cost of reducing each risk-bearer's individual risk of harm). The individual risk principle is keyed to the marginal tradeoff between each cost-bearer's share of safety precaution costs and each risk-bearer's individual risk. As the number of persons exposed to the risk grows, the marginal risk-reducing efficacy of each dollar spent on safety precaution will plausibly change. For example, the first $100,000 spent on safety precautions at a 10-worker factory would almost certainly reduce each worker's individual risk by more than the first $100,000 spent on safety precautions at a 10,000-worker factory. Further, the point at which further expenditure on safety precaution would produce no appreciable reduction in each worker's risk would likely be reached sooner in the case of a 10-worker factor than in that of a 10,000-worker factory. For these reasons, and holding the number of cost-bearers constant, an increase in the number of risk-bearers may dictate a greater, smaller, or equal pro rata investment in safety precautions, depending on the circumstances.

To the extent the IFP's lack of direct sensitivity to the size of the risk-bearing population seems morally implausible—and I suspect it may to Fried and other aggregationists—it is important to bear in mind that aggregative interpretations of reasonable precaution are not without their own counterintuitive implications. As the suspension bridge hypothetical discussed in the Introduction makes clear, summing costs and benefits across persons can lead to results that conflict dramatically with widely-held moral convictions. Indeed, no extant theory of reasonable precaution that I am aware of has the virtue of delivering results that accord perfectly with moral intuitions in all situations. If it is implausible to regard the imposition of a 1 in 1,000 risk of death on 10 people as requiring the same regulatory response as the imposition of the same risk on 10,000 people, it seems no less implausible to regard the imposition of a 50% risk of death on two people and the imposition of a 1 in 1 million risk of death on 1 million people as morally indistinguishable on the grounds that both can be expected to result in the death of a single person. Even granting that the IFP may generate implausible results in certain situations, these results seem no *more* implausible to me than those associated with aggregative precautionary standards like efficient care.

III. A Contractualist Justification for the IFP

I have so far argued that the individualized feasibility principle represents a viable alternative to efficient care. Like efficient care, it is a substantive interpretation of reasonable precaution that is capable of delivering precise guidance concerning the requisite level of precaution in specific cases. Unlike efficient care, the IFP is capable of accommodating the intuition that, where risks are concentrated on a group of persons that is a small fraction of the size of the group among whom precaution costs are spread, acting reasonably may require investing more in safety precaution than is marginally cost-justified.

In this Part, I explore the theoretical underpinnings of the IFP. If the efficient care standard is straightforwardly grounded in utilitarianism's directive to maximize aggregate well-being,[84] what is the theoretical justification for the IFP? Given that the IFP typically dictates a level of precaution that does *not* optimize aggregate well-being, its theoretical grounding is not surprisingly found generally in nonconsequentialism, and specifically in a relatively new nonconsequentialist theory of normative ethics known as *ex ante contractualism*. In this Part, I argue that if the morality of risk imposition is determined from the standpoint of ex ante contractualism, the risks posed by socially beneficial activity should be mitigated to the extent required by the individualized feasibility principle.

A. Contractualism's Core Ideas

Contractualism is a theory of normative ethics introduced by the philosopher T.M. Scanlon in 1982.[85] It is not a general theory of morality, but rather a theory of *interpersonal* morality or, to use Scanlon's famous phrase, of "what we owe to each other."[86] From its inception, contractualism has purported to represent an alternative to utilitarianism and, in partic-

---

[84] *See* Oberdiek, Morality of Risking, *supra* note ____, at 143("The Hand test is essentially a consequentialist handmaiden, based on justification to the world at large and envisaging interpersonal aggregation of burdens and benefits.").

[85] *See* T.M. Scanlon, *Contractualism and Utilitarianism*, *in* UTILITARIANISM AND BEYOND 103-28 (Amartya Sen & Bernard Williams, eds.) (1982) (hereinafter "Scanlon, *Contractualism*"); T.M. Scanlon, WHAT WE OWE TO EACH OTHER (1998) (especially Chapter 5, pp. 189-247) (hereinafter "Scanlon, WHAT WE OWE").

[86] Thus, contractualism does not directly address questions of political morality, i.e., the moral strictures applicable to the actions and policies of coercive institutions, nor does it address our moral obligations to animals or future persons. *Cf.* Arthur Ripstein, EQUALITY, RESPONSIBILITY, AND THE LAW 5 (1999) (noting that political morality, "the morality governing the exercise of force, has its own standards of responsibility that may well be out of place in other contexts.").

ular, to embody an argument against the core utilitarian tenet that all moral questions reduce to questions about the consequences of acts (or rules) for the aggregate well-being of all affected individuals.[87]

Contractualism offers the following formula for determining the rightness or wrongness of a particular act:

> An act is wrong if its performance under the circumstances would be disallowed by any set of principles for the general regulation of behavior that no one could reasonably reject as a basis for informed, unforced, general agreement.[88]

What exactly does that mean? And in what sense is this formula supposed to represent an alternative to utilitarianism?

The contractualist formula involves three core ideas: (i) the idea that interpersonal morality presupposes the requirement of justifiability to *each* affected person considered as an individual, rather than the requirement of justifiability to *all* affected persons considered in the aggregate;[89] (ii) the idea that the moral status of a particular act (its rightness or wrongness) is a function of the moral validity of the *general principle* licensing the act;[90] and (iii) tying together the first two tenets, the idea that a principle is justifiable to each person if and only if it would command the free assent of all persons, i.e., if and only if no person could *reasonably reject* it as a principle for the general regulation of behavior. Thus, under contractualism, an act is morally permissible if and only if no one could reasonably reject the general principle permitting the act.

When, according to contractualism, can a principle be reasonably rejected? The concept underlying the notion of reasonable rejectability is the *minimax criterion*.[91] Among a set of candidate principles, the nonrejectable principle P is the one of which the following is true: the strongest complaint any person could make against P, were P generally accepted, is weaker than the strongest complaint that could be made against every other alternative principle. As Scanlon puts it, "[S]omeone can reasonably reject a principle if there is some alternative to which no other person has a com-

---

[87] Scanlon, *Contractualism*, *supra* note __, at 103.

[88] Scanlon, WHAT WE OWE, *supra* note __, at 153.

[89] *Id.* at 153, n. 8 ("What is basic to contractualism as I understand it is the idea of justifiability to each person (on grounds that he or she could not reasonably reject).").

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

[91] I follow Sophia Reibetanz in using the minimax criterion to explain the notion of reasonable rejectability. *See* Sophia Reibetanz, *Contractualism and Aggregation,* 108 ETHICS 296, 300 (1998) (describing "minimax complaint model" of reasonably rejectability).

plaint that is as strong."[92] The principle no one could reasonably reject is the principle that, among a set of candidate principles, *minimizes* the strength of the complaint that could be lodged by the *maximally* burdened person.[93] Supposing I am the person that would be most burdened by general acceptance of a particular principle, I still cannot *reasonably* reject that principle if every alternative principle would, if generally accepted, impose a greater burden on someone else.

Contractualism thus contemplates a *rejectability inquiry* the goal of which is to identify the principle satisfying the minimax criterion. This is a fundamentally comparative inquiry that takes into account not just the extent to which general acceptance of each candidate principle burdens each affected person in an absolute sense, but also the differential each person experiences in the burdens they would bear under the respective principles.[94] That is, supposing A is the most burdened party under Principle P-1 and B the most burdened party under Principle P-2, we ask not only if A's burden under Principle P-1 is weightier than B's burden under Principle P-2, but also if A's gain in moving from P-1 to P-2 is more significant than B's gain in moving from P-2 to P-1. The question to be asked, in other words, is whether it would be unreasonable for A to refuse to accept the burden she must bear under P-1 in order that B can enjoy the benefit she must relinquish under P-2.[95]

To make this all a bit more concrete, consider how contractualism would determine whether it is morally permissible to lie to someone threat-

---

[92] Scanlon, WHAT WE OWE, *supra* note __, at 229.

[93] Applying the minimax rule to complaints is equivalent to applying the more familiar "maximin" rule to outcomes for individual well-being. *Cf.* Rawls, *supra* note _____, at 133 ("The maximin rule tells us to rank alternatives by their worst possible outcome: we are to adopt the alternative the worst outcome of which is superior to the worst outcomes of the others.").

[94] *See* Scanlon, *Contractualism*, *supra* note __, at 113 ("Whether it would be unreasonable for me to reject a certain principle, given the aim of finding principles which no one with this aim could reasonably reject, depends not only on how much actions allowed by that principle might hurt me in absolute terms but also on how that potential loss compares with other potential losses to others under this principle and alternatives to it."); Reibetanz, *supra* note __, at 299 (in the rejectability inquiry, "we take a person's *complaint* about a principle to be a function both of her absolute level of well-being under acceptance of that principle and of the burden which acceptance of that principle imposes upon her (that is, the amount by which she would be worse off under acceptance of that principle than under acceptance of some alternative)") (emphasis in original).

[95] *See* Scanlon, *Contractualism* at 123. For example, suppose that under P-1, A's well-being is 150 and B's 101, and that under P-2, A's well-being is 102 and B's 103. The mere fact that the loser under P-2 (A) is slightly better off than the loser under P-1 (B) does not necessarily mean that P-2 is the nonrejectable principle satisfying the minimax criterion. This is because A might plausibly be taken to have a stronger complaint with P-2 being chosen over P-1 than B would have with P-1 being chosen over P-2. Why? Because it would arguably be unreasonable for B to refuse to accept a two-unit (roughly two percent) reduction in well-being in order that A can enjoy a 49-unit (roughly 50 percent) increase in well-being.

ening to wrongfully harm another person. To work with a specific example, suppose that, while he is robbing me at gunpoint late one night, the robber demands to know where I live. In response, I tell him that I do not live in the immediate area, when in fact I live a few houses away. I do so because I wish to protect my sleeping family from danger and my household possessions from theft. Is my lie morally permissible?

Per contractualism, my act of lying under these circumstances is morally permissible only if it would be allowed by a general behavioral principle that no one could reasonably reject. Rejecting as implausible a principle granting categorical permission to lie to others, consider two alternative principles: (i) Qualified Permission To Lie: It is permissible to lie to another person when the person is threatening to wrongfully harm someone and when lying reduces the likelihood that the person will succeed in wrongfully harming them; (ii) Categorical Prohibition On Lying: It is never permissible to lie to another person. The question is which of these principles satisfies the minimax criterion, i.e., which principle imposes a lighter burden on the person it burdens most heavily?

The persons most burdened by general acceptance of the Qualified Permission principle would, it seems, be those on the receiving end of the sanctioned lies. Those threatening to wrongfully harm others might complain that, notwithstanding their own wrongful conduct, being lied to shows them disrespect or diminishes their well-being in some respect. On the other hand, the persons most burdened by general acceptance of the Categorical Prohibition principle would, it seems, be persons forced to divulge truthful information to those bent on using such information to wrongfully perpetrate serious harm, as well as persons placed at increased risk of being wrongfully harmed as the result of such compelled disclosures.

The question then becomes, who bears the greater burden: those would-be wrongful actors on the receiving end of the lies permitted by the Qualified Permission or those placed at increased risk of suffering wrongful harm as a result of the truthful disclosures compelled by the Categorical Prohibition? It seems obvious that the latter group bears the far greater burden. The German citizen forced to truthfully disclose to a group of Nazi officers that he is hiding a Jewish family in his cellar, and the Jewish family likely to be wrongfully harmed as a result of the disclosure, no doubt have a stronger complaint with the Categorical Prohibition than the robber in the above example has with the Qualified Permission. In other words, it seems clear that persons like the German citizen or the Jewish family could reasonably refuse to accept the burden they must bear under the Categorical Prohibition so that persons like the robber or the Nazi officers can enjoy the benefits they would have to give up under the Qualified Permission. Thus, the Qualified Permission To Lie is not reasonably rejectable. Because that principle licenses my lying to the robber under the circumstances, that act is morally permissible.

One of the most distinctive features of contractualism is its pluralism concerning the types of reasons that carry weight in moral reasoning. On Scanlon's view, a person can reasonably reject a candidate principle not only in virtue of its consequences for the person's well-being, but also in virtue of its consistency (or lack thereof) with the person's rights, entitlements, or, as the self-defense example illustrates, fairness interests.[96] Thus, contractualism is properly considered a nonconsequentialist moral theory, since it holds that the moral permissibility of an act may properly depend on factors other than the consequences of the act (or its licensing principle) for individuals' well-being.

Another critically important feature of contractualism is what has been termed its *individualist restriction*, i.e., "the insistence that the justifiability of a moral principle depends only on various *individuals'* reasons for objecting to that principle and alternatives to it."[97] In other words, according to the individualist restriction, the strength of a complaint lodged against a particular principle can never be a function of the *sum* of different individuals' gain (or loss) in well-being under that principle as compared with some alternative principle. Contractualism instead contemplates a series of "pairwise comparisons"[98] in which *one* representative individual's burden under a particular principle is compared to *one other* representative individual's burden under an alternative principle. I adhere to the individualist restriction here because doing so ensures that the precautionary decision rule that emerges from my contractualist analysis will be genuinely nonaggregative, thereby respecting the separateness of persons in a way interpersonally aggregative standards do not.

B.   A Dilemma for Contractualism?

Fried seems to believe that, of all forms of nonconsequentialism, contractualism holds the greatest promise of providing a viable, nonaggregative standard for regulating risks. However, despite the potential Fried thinks contractualism holds, she still believes it will necessarily fail to deliver a workable standard in the final analysis.[99]

---

[96] *See* Scanlon, WHAT WE OWE at 229 (discussing the ways in which his version of contractualism eschews the welfarism of the so-called Complaint Model).

[97] *See id.* at 229 (emphasis in original).

[98] *See* Reibetanz, *supra* note __, at 300 (under contractualism, "we determine whether an individual could reasonably reject some principle by making a series of pairwise comparisons: each person's complaint about that principle is compared separately with the complaint of every other person").

[99] Fried, *Limits*, *supra* note __, at 259 ("Contractualist approaches to the problem of harm are a more complicated case. I do think they can provide a morally meaningful, nonconsequentialist justification for picking one scheme of risk regulation over another, and for choosing one that deviates from standard, unweighted aggregation. . . But at the end of the

Her reasoning for this conclusion takes the form of a dilemma. The dilemma emerges from consideration of the two alternative epistemic points of view—the *ex ante* and the *ex post*—from which candidate principles could be evaluated in contractualism's rejectability inquiry. These two points of view represent alternative assumptions about the type and extent of information available concerning how each affected individual fares under the candidate principles.

Under ex ante contractualism, candidate principles are evaluated based on their *expected* outcomes for each affected individual's well-being.[100] This would mean that the rejectability inquiry would be based on probabilistic data about each individual's expected fate under each precautionary principle, i.e., her likelihood of suffering death or serious injury from the risks the principle tolerates, the size of her expected benefit from the risky activities the principle sanctions (or prohibits), the amount of the monetary costs she would be expected to bear under the principle, etc.

Under ex post contractualism, on the other hand, candidate principles are evaluated based on their *actual* outcomes for each individual's well-being.[101] In this case, the rejectability inquiry would be based on data about each individual's actual fate under each precautionary principle, including, most saliently, whether she ends up being killed, seriously injured, or left unharmed by the risks the principle tolerates.

With the ex ante/ex post distinction in mind, Fried's dilemma can be stated as follows. If possible precautionary standards are evaluated from an ex post point of view, contractualism will implausibly deem impermissible most forms of socially beneficial, though significantly risky, conduct. If a precautionary standard tolerates a risk that results in the death of even a single person, that standard could, it seems, be reasonably rejected, since no person will have a stronger complaint than one based on the loss of life.[102] As James puts it, "Complaints of death will always carry the day."[103] Thus, if the ex post point of view governs, any risky activity that results in the death of at least one person will be impermissible according to contractual-

---

day, I believe all of the plausible schemes from which to choose will necessarily be aggregative in one form or another.").

[100] *See* James, *supra* note __, at 267 ("Let us call *ex ante contractualism* the. . .view that only *expected* outcomes count as grounds for complaint or objection (including expected outcomes of a principle's general adoption), mounted on behalf of each potentially affected party, from some specified epistemic position.") (emphasis in original).

[101] *See id.* ("Let us call *ex post contractualism* the view that we should evaluate what decision is reasonably acceptable only in light of its *actual outcomes*, as they actually unfold over time.") (emphasis in original).

[102] *See* Fried, *Contractualism*, at 9-10 (noting that, under ex post contractualism, "any principle that authorizes actions that risk gravely harming at least one person will be rejected by the hypothetical representative person who, by happenstance or by being permitted to peek ahead, learns she will be the unlucky one").

[103] James, *supra* note __, at 272.

ism, since there will be at least one person who could reasonably reject the principle that allows the activity to go forward. But this would bar all sorts of risky activities—large-scale construction projects, etc.—that, assuming reasonable precautions are taken, seem intuitively acceptable, notwithstanding that they are certain to result in the accidental loss of life and limb.[104]

If, on the other hand, the contractualist inquiry assesses precautionary standards from an ex ante point of view, it will, in Fried's view, end up settling on a standard that is indistinguishable from the sort of interpersonally aggregative decision rule used in CBA:

> If representative persons are imagined to choose general principles from an ex ante POV, each will prefer the principles that optimize her expected subjective value, given her ex ante preferences, circumstances, etc. Assuming a plausible range of risk aversion and altruism, a normal distribution of preferences about outcomes *and evenly distributed upside and downside risks*, the principles that optimize a given individual's expected position will be roughly equivalent to the principles that optimize aggregate wellbeing. (To put it another way, under conditions of uncertainty, optimal individual choice tends to converge with optimal social choice.) Thus, we end up with an aggregate solution by a different route.[105]

According to Fried, if contractualism evaluates precautionary principles from the *ex ante* standpoint, it will settle on roughly the same decision rule as that favored by aggregative welfarists: take all and only those precautions that are marginally cost-justified, summing costs and benefits across persons. This implies that, if contractualists adopt the *ex ante* point of view, they forfeit the capacity to explain why pressing precaution beyond the point of cost-justification can be morally obligatory.

Thus, Fried thinks that contractualism either yields the same decision rule favored by aggregative welfarists or imposes an implausible ban on most risky, socially beneficial activity. "The basic problem facing con-

---

[104] *See* James, *supra* note __, at 268-72 (arguing that ex post contractualism results in "moral gridlock," i.e., a ban on virtually all socially beneficial but risky activities).

[105] Fried, *Contractualism* at 43-44 (emphasis added). Fried's analysis here is importantly limited by the assumption described in the language I have italicized. If "upside and downside risks" are truly evenly distributed, then Fried is correct that a representative individual could reasonably reject any precautionary standard that failed to optimize aggregate wellbeing. The problem is that the risks posed by a socially beneficial activity rarely fall evenly on the individuals they fall on and, even more critically, there is often incomplete overlap between the set of individuals who bear the "downside" risks of the activity, the set of individuals who stand to benefit from the activity (those who bear the activity's "upside" risks), and the set of individuals who bear the costs of reducing the downside risks. Thus, Fried's analysis suffers from a problem that has dogged consequentialist moral theory virtually since its inception: a failure to take into account morally relevant variances in the *distribution* of the effect of a particular act or principle on individuals' wellbeing.

tractualists," she opines, "is that adopting an ex ante POV proves too little and adopting an ex post POV proves too much."[106]

C. Ex Ante Contractualism and Generic Reasons

Fried can generate an ostensible dilemma only by relying on an inadequately nuanced description of the epistemic position from which complaints would be lodged under ex ante contractualism. When ex ante contractualism is made sensitive to asymmetries in the distribution of risks, costs, and benefits among different individuals, ex ante contractualism can indeed yield a substantive decision rule for safety precaution that diverges from efficient care and other aggregative standards.[107]

In her description of ex ante contractualism, Fried builds in assumptions that effectively efface the real-world distributional differences of which contractualism is designed to take account. Although she seems to recognize that the subjective preferences and circumstances of the complainants in the rejectability inquiry will vary, by stipulating an "even distribution of upside and downside risks," she assumes away the single most important manner in which circumstances *will* and *do* in fact vary in the context of risk regulation. In the real world, particular risks posed by particular activities rarely fall evenly on the people they fall on. Even more critically, some risks are borne by beneficiaries of the activity at issue (such as workers or consumers), whereas other risks are borne chiefly by non-beneficiaries (such as bystanders). Some risks are borne primarily by persons who also bear a portion of the costs of reducing those same risks (many consumption risks), whereas, as the suspension bridge hypothetical illustrates, other risks are borne chiefly by persons who do not bear any appreciable share of those costs (most workplace risks). Such morally relevant differences are negated in Fried's treatment of ex ante contractualism.

To be fair, Fried does briefly discuss a case in which an activity's risks and benefits are distributed unevenly among the affected individuals.[108] She considers a case in which the efficient course of action would be to site a toxic waste dump in an area where land values are lowest (Poor-

---

[106] *See* James, *supra* note __, at 265 ("[I]f contractualism allows 'ex post' objections, by considering actual outcomes, it becomes difficult to justify the risks created by most public policy, leaving contractualism at odds with moral commonsense in much the way utilitarianism is. But if contractualism instead takes a full 'ex ante' form, by considering only *expected* outcomes, it becomes unclear how it is supposed to recommend something other than aggregative cost-benefit decisionmaking.")

[107] *See id.* at 266 ("Fried's dilemma nevertheless fails on its second horn. Suitably elaborated, an 'ex ante contractualism' meaningfully constrains aggregative cost-benefit decision making.").

[108] *See* Fried, *Contractualism, supra* note __, at 57-58.

ville).[109] Because of their proximity to the dump, Poorville residents bear a disproportionately high health risk.[110] Is this decision morally defensible? Fried considers three ways in which it might be made so: provide cash compensation to Poorville residents; give risk-bearers a veto over whether the risky activity should proceed at all; justify the efficient course of action to Poorville residents by "bundling" risks together, such that the risks posed by the dump are offset by the benefit provided by other, similar risks.[111] Notably absent from Fried's consideration is the possibility of siting the dump in the location where it poses the lowest risk to human health. Assuming all possible dump locations pose the identical risk to human health (which seems rather unlikely as a practical matter), is it morally defensible to site the dump in Poorville on account of its having the smallest possible effect on property values? From a contractualist perspective, it seems not. It seems the only fair principle to follow in such circumstances is (again) one Fried does not consider: choose the site randomly and compensate the residents living close to the site in an amount sufficient to allow each of them to relocate, if they wish. This may lead to an inefficient result (e.g., if a site in a rich neighborhood is the one randomly selected for the dump), but it seems to me to be the only course of action that no one could reasonably reject. A principle allowing siting decisions to be made on the basis of property values, if generally accepted and applied repeatedly over time, would likely result in a disproportionate share of toxic dumps and other hazardous operations being sited in poor neighborhoods, placing a resident of such a neighborhood at a disproportionately high risk of harm. This, it seems to me, would give someone who lives in a poor neighborhood a basis for reasonably rejecting such a principle in favor of a principle dictating a random or proportionate siting policy.

Let us return to the question of how one might use ex ante contractualism to generate a distinct interpretation of reasonable precaution. If we are to improve upon Fried's analysis, the question that must be confronted is how to use ex ante contractualism in a way that takes account of the morally relevant differences in how costs, risks, and benefits can be distributed among distinct groups of persons.

I believe the beginning of an answer lies in Scanlon's notion of *generic reasons*. According to Scanlon, contractualism's rejectability inquiry is to be conducted neither from the perspective of a hypothetical "average" person nor from the perspective of specific individuals. Scanlon instead contemplates an intermediate perspective, in which complaints are based on "reasons that we can see that people have in virtue of their situation, characterized in general terms, and such things as their aims and capabilities and

---

[109] *See id.* at 57
[110] See id.
[111] *See id.* at 58.

the conditions in which they are placed."[112] Scanlon terms these "generic reasons."

Generic reasons include reasons based on desires or goals that could plausibly be attributed to *all* persons. "We commonly take it," Scanlon observes, "that people have strong reasons to want to avoid bodily injury, to be able to rely on assurances they are given, and to have control over what happens to their own bodies."[113] However, as the definition quoted in the previous paragraph suggests, a rejectability inquiry based on generic reasons is also meant to take account of *differences* in people's situations and of the way these differences would presumably affect their willingness to endorse candidate principles. Scanlon explicitly acknowledges this: "Not everyone is affected by a given principle in the same way, and generic reasons are not limited to reasons that the majority of people have."[114]

This suggests that, in evaluating principles governing the level of precaution to be exercised when engaging in a socially beneficial activity that poses a significant risk of serious harm, we need to consider the most salient respects in which people can be differently situated with respect to a particular risk posed by a particular activity. There are three key respects in which people can be so differently situated: (1) in the extent to which they benefit from the activity; (2) in the extent to which they bear the risk at issue; and (3) in the extent to which they bear the cost of safety precautions that reduce that risk. This implies that we need to evaluate risks from at least three different perspectives: those who benefit from the activity at issue, those who bear the activity's risks, and those who bear the costs of safety precautions that reduce those risks.

Only in rare circumstances will these three groups overlap perfectly.[115] For example, as the suspension bridge hypothetical illustrates, most

---

[112] Scanlon, WHAT WE OWE, at 204.

[113] Id.

[114] *Id.* at 204-205. Scanlon goes on to observe, "If even a small number of people would be adversely affected by a general permission for agents to act in a certain way, then this gives rise to a potential reason for rejecting that principle." *Id.* at 205.

[115] When they do—i.e., when the identical group of persons realizes the benefits of a risky activity, bears the activity's risks, and bears the costs of mitigating those risks, and when benefits, risks, and costs are evenly distributed throughout that group—then Fried is correct that ex ante contractualism would dictate taking the same level of precaution as that dictated by the efficient care principle. If all participants in the rejectability inquiry are, by hypothesis, identically situated with respect to the risk at issue and if all relevant risks, benefits, and costs are evenly distributed, then the only principle that is justifiable to *each* participant is the principle that maximizes the average well-being of *all* participants. Every other principle is rejectable because, compared with the efficient care principle, it leaves each affected person worse off. James appears to be in agreement on this point. *See* James, *supra* note __, at 278 (acknowledging that an "aggregative rule of social choice" would be "justified" under ex ante contractualism in cases where "everyone potentially affected has roughly the same ex ante prospects of benefit" and "*no one can otherwise mount a personal objection*") (italics in original).

workplace risks are borne by persons (workers) who do not bear the costs of reducing those risks through safety precautions. Workers generally do not enjoy the cost savings that result from leaving certain workplace safety precautions left untaken. Those savings tend to accrue to consumers or shareholders or some combination of the two.[116] The same could be said of socially productive activities that pose environmental risks to bystanders. The person exposed to a heightened cancer risk as a result of living next to a cement factory that emits toxic fumes into the air bears that risk without also bearing the costs of safety precautions that reduce it (or realizing the savings associated with leaving such precautions left untaken). Thus, in many cases, there will be little overlap between those exposed to the risk at issue and those who bear the costs of reducing the risk through the exercise of safety precautions. It seems that, contrary to Fried's claim, upside and downside risks are quite often *not* evenly distributed.

Thus, following James and Keating,[117] I here deploy an ex ante contractualism that evaluates candidate principles based on the generic reasons that risk-bearers, cost-bearers, and beneficiaries would have to object to those principles in light of their expected outcomes.

D. Well-being, Willingness-to-Accept, and Willingness-to-Pay

Under ex ante contractualism, candidate principles are evaluated based on their expected outcomes for individuals' well-being. I need to say something about the particular conception of well-being I will rely on in evaluating candidate precautionary standards.

The ex ante contractualism I deploy here presupposes the coherence of interpersonal comparisons of levels of well-being and further undertakes to quantify gains and losses to individuals' respective levels of well-being in terms of willingness-to-accept (WTA) and willingness-to-pay (WTP).[118] WTA and WTP together embody a preference-based, monetary conception of well-being, one that is typically used in cost-benefit analysis (CBA). WTA refers to the minimum amount of money a person would be willing to

---

[116] *See* Keeton et al., *supra* note ___, at 1257 ("[T]he risks and benefits of a consumer product accrue to the same party, the consumer, who can choose to avoid the risk. Employees, on the other hand, rarely have such a choice—they must bear the risk of occupational health hazards, while benefits accrue largely to employers and consumers.") (quoting Federal Respondent's Brief in Am. Textile Mfrs. Inst. v. Donovan, 452 U.S. 490 (1981), at 55).

[117] *See* James, *supra* note __, at 274 (endorsing ex ante contractualism over ex post contractualism or a hybrid ex ante-ex post contractualism); Keating, *Pressing Precaution*, at 674-84 (implicitly relying on an ex ante version of social contract theory).

[118] The coherence of interpersonal comparisons of levels of well-being is taken for granted by most social welfare theorists. *See, e.g.,* Adler, WELL-BEING, *supra* note ___, at xv-xvi, 185-201 (offering and defending an account of well-being that allows for interpersonal comparability).

accept to put up with something undesirable, such as a specified increase in their risk of death or serious bodily injury. WTP refers to the maximum amount of money a person would be willing to pay for something desirable, such as a specified reduction in their risk of death or serious bodily injury.[119]

Three types of changes in individuals' well-being will be of particular concern to me in this Article. First, as is generally accepted, a person who bears a risk of death or serious bodily injury experiences a reduction in well-being (compared with not being exposed to the risk at all), even if the risk never materializes. Typically, this sort of reduction is quantified in terms of a WTA in the amount of $1/r \times V$, where $1/r$ represents the probability of the risk materializing in injury to that person and V represents the WTA associated with sustaining the injury itself. Thus, if life is valued at $5 million, the WTA associated with, for example, bearing a 1 in 1,000 risk of death is $5,000 and the WTA associated with bearing a 1 in 10,000 risk of death is $500.[120] The underlying idea is that a rational person would be willing to accept a sufficiently large amount of money to be exposed to an additional risk of death that, while not insignificant, is still quite low. Second, a person who bears a monetary cost (e.g., the cost of safety precautions) experiences a reduction in well-being. Here, quantifying the reduction in well-being in terms of WTA is straightforward: the WTA associated with bearing a monetary cost is simply the amount of the cost itself. Finally, a person who derives a benefit from a particular industrial activity—or from a coordinated system of such activities—experiences a gain in well-

---

[119] Suppose X is better off under outcome O than under outcome O'. A preference-based account of well-being maintains that this just means that X *prefers* O to O'. We might then ask *how much* better off X is under O than under O', i.e., how *strongly* X prefers O to O'. One way to answer this question would be to specify the *minimum amount of money X would be willing to accept* for a move from O to O'—this is X's WTA. Another way to answer the question would be to specify the *maximum amount of money X would be willing to pay* for a move from O' to O—this is X's WTP. This is precisely how CBA quantifies changes in individuals' well-being. *See, e.g.,* Adler, WELL-BEING, *supra* note ___, at 6.

[120] *See, e.g.*, Eric A. Posner & Cass R. Sunstein, DOLLARS AND DEATH, 72 U. Chi. L. Rev. 537, 560 (2005); Thomas Kniesner, W. Kip Viscusi, & James P. Ziliak, *Willingness to Accept Equals Willingness to Pay for Labor Market Estimates of the Value of Statistical Life,* (Vanderbilt Univ. Law School Law & Econ. Working Paper No. 13-06, 2012) at 1-3; W. Kip Viscusi, *The Value of Life: Estimates with Risks by Occupation and Industry*, 42 J. ECON. INQUIRY 29 (2004)*;* W. Kip Viscusi, *Estimating the Value of a Statistical Life Using CFOI Data* (Vanderbilt Univ. Law School Law & Econ. Working Paper No. 13-17, 2013). Therefore, a reduction in a person's risk of death from 1 in 1,000 to 1 in 10,000 would be "worth" $4,500 to her in the following sense: the amount of money she would be willing to *accept* to bear a 1 in 10,000 risk of death is $4,500 less than the amount she would be willing to accept to bear a 1 in 1,000 risk of death. Alternatively, she would be willing to *pay* $4,500 for the reduction. *See* Kniesner et al., *supra* this note, at 1 (arguing "that there is no significant divergence between willingness to accept and willingness to pay in the important case of the decisions workers make concerning exposure to fatal injury risk").

being that can be quantified in terms of WTP. By the same token, when a particular socially beneficial activity is discontinued, the reduction in well-being each beneficiary consequently experiences can be quantified in terms of WTA.[121]

E.   Justifying the Economic Feasibility Cap on Precaution

Under the IFP, investment in safety precaution is capped at the economic feasibility point, i.e., the point at which further reduction of the risk would threaten to undermine the survival of the socially beneficial activity giving rise to the risk. What is the ex ante contractualist rationale for capping safety expenditures at this point? In contractualist terms: why is it that no one could reasonably reject a precautionary standard that caps safety expenditures at the economic feasibility point?

Contractualism's emphasis on considering the consequences of a precautionary standard's *general* acceptance, as well as its acceptance in the case of a particular risk or activity, is key here.[122] It means that, in the rejectability inquiry, each affected individual can lodge complaints based not only on how acceptance of a given precautionary standard would affect her in the case of a particular risk posed by a particular activity, but also on how she would be affected, over the course of her lifetime, by the standard's application to *all* risks to which it applies and to *all* activities that pose such risks.[123]

General acceptance of a precautionary standard lacking a cap at the point of maximal economic feasibility[124] would plausibly result in the dis-

---

[121] As noted above, in calculating WTA and WTP in this Article, I assume for sake of simplicity that all affected individuals have identical levels of existing wealth, and so assume away the complications that arise due to the diminishing marginal utility of money. Thus, all affected individuals are assumed to experience the identical reduction in well-being as the result of bearing a specified monetary cost. Further, all individuals are assumed to have the identical WTA for a specified increase in their risk of death and the identical WTP for a specified reduction in that risk. *See supra* note ___. In future work, I hope to address this important issue in greater detail.

[122] *See* Scanlon, WHAT WE OWE, *supra* note ___, at 202-203 ("[W]hen we are considering the acceptability or rejectability of a particular principle, we must take into account not only the consequences of particular actions, but also the consequences of general performance or nonperformance of such actions and of the other implications (for both agents and others) of having agents be licensed and directed to think in the way that that principle requires.").

[123] Scanlon calls this feature of contractualism "intrapersonal aggregation," i.e., "aggregation *within* each person's life, summing up all the ways in which a principle demanding a certain level of care would constrain that life, rather than aggregation *across* lives, adding up the costs or benefits to different individuals." *Id.* at 237; *see also* Oberdiek, Morality of Risking, *supra* note ___, at 146-51 (discussing the notion of intrapersonal aggregation).

[124] Examples of precautionary standards that would often dictate a greater-than-feasible expenditure on precaution include a "safety" standard (which directs risks to be reduced to the point at which they are no longer significant regardless of cost or impact on the industry

continuation of many, if not most, of the life-improving, liberty-enhancing activities that define modern life.   Many, if not most, of the socially beneficial activities undertaken in modern, industrialized societies pose risks that either cannot, given present technology, be eliminated or would be so costly to eliminate that the underlying activity would no longer generate a net social benefit (thus jeopardizing the activity's long-term survival).[125]

Were the majority, or even a significant proportion, of risky, socially beneficial activities discontinued, we would live in a very different world, one resembling the pre-industrial, more-or-less agrarian society prevalent in the Western world prior to the Industrial Revolution. Major industrial activities such as large scale construction projects, the production of gasoline and natural gas, the manufacture of automobiles and other gas- or electric-powered road vehicles, and the manufacture of prescription drugs would cease to exist, as they pose risks that cannot feasibly be eliminated or reduced to a level that is no longer significant.   Such a state of affairs would give rise to a profound complaint from a person who enjoys the manifold benefits that flow from the complex of socially beneficial, but risky, activities that make possible modern life.   General acceptance of a pure safety standard, a technological feasibility standard, or any other precautionary principle lacking a cap at the economic feasibility point would deprive each beneficiary of the plausibly enormous benefit she realizes from the wide class of activities whose risks cannot feasibly be eliminated or reduced to the point of insignificance.

The key question from a contractualist perspective is whether this complaint is stronger than a risk-bearer's ex ante complaint with a principle that caps precaution at the point of economic feasibility.   Who, in other words, bears the heavier burden: the person deprived of the benefits of the (plausibly large) class of activities whose risks cannot feasibly be eliminated or reduced to the point at which they are no longer significant, or the person exposed to a significant risk of death or serious injury from one such activity?

---

and a "technological feasibility" standard (which directs risks to be reduced as much as possible given available technology, regardless of cost or the impact on the industry). *See* Keeton et al., *supra* note ___, at 1237-39, 1252-53.

[125] *See, e.g.,* Fried, *Limits, supra* note ____, at 260 ("Because virtually everything we do (or, acting as the state, permit others to do) carries some irreducible risk of serious harm to others, virtually everything we do (or permit others to do) entails interpersonal tradeoffs."); Fried, *Contractualism*, *supra* note ____, at 61 ("But in most arenas of life. . .the point at which further investments in safety will cease to produce any positive return in safety is far beyond the projected benefits of the project. The result of interpreting ''reasonable'' in this fashion will thus be moral gridlock by a different route. No building project will show a net expected benefit, given the enormous amount one would be required to spend on safety precautions, and hence no building project will go forward.").

While it seems clear that each of us would rather be deprived of the manifold benefits of industrialized society than suffer death or serious bodily harm,[126] the question becomes a much closer one when the certain deprivation of the manifold benefits of modern industrialized society is the price to be paid for avoiding a mere *risk* to one's physical integrity. As Keating points out, it seems hard to deny that preserving the life-improving, liberty-enhancing benefits to be derived from the many activities in modern, industrialized societies posing significant risks of serious bodily harm is at least *comparable* in value to protecting one's bodily integrity from significant risk.[127] Where the risk of death or serious injury is low, but still nontrivial (say, somewhere between 1 in 1,000 and 1 in 50,000), the overwhelming odds are, supposing that I am exposed to such a risk, that my physical integrity will remain intact. In light of that probability, and given the profound benefits to be realized from the many activities whose risks cannot feasibly be eliminated, it seems likely that even someone who faces a significant risk of death from a particular socially beneficial activity would prefer the feasibility principle, which tolerates that risk but preserves all socially beneficial activities posing similar risks, to a principle which eliminates the risk at the price of shutting down the vast majority of such activities.

There is a further point to consider. Among the benefits provided by the productive activities that pose significant risks is a dramatic reduction in the risks to bodily integrity posed by non-anthropogenic causes like disease, natural disasters, and climatic variations. This means that the discontinuation of the majority of such activities could plausibly be supposed to result in a *profound* increase in the risks to bodily integrity faced by the beneficiaries of such activities.[128] If the price for continuing such activities were the creation of *huge* risks to bodily integrity—something on the order of a 1 in 10 risk of death—then the risk-bearer's complaint might well carry the day. But the risks of death or serious injury posed by most socially ben-

---

[126] If the choice is between living healthy and able-bodied in an agrarian, pre-industrial society and living as a quadriplegic in modern, industrialized society, most of us would, I imagine, choose the former. Of course, even pre-industrial society was not without socially beneficial activities that generated significant risks of harm (e.g., large-scale construction, shipbuilding, etc.). One thing that seems to define modern, industrial society is that the risks posed by such activities became comparable in magnitude to the risks posed by the composite of natural causes (like disease and natural disasters) and anthropogenic activities that are not clearly socially beneficial (e.g.., war, slavery, etc.).

[127] *See* Keating, *Pressing Precaution*, at 727 ("Shutting down most of the major productive activities in our economy *would* be a harm comparable to bearing a significant risk of devastating injury.").

[128] Just as one illustration, the world average life expectancy at birth in 1900 was 31 years, and under 50 years in even the richest countries. By the mid-20th century, the average rose to 48 years. By 2005, this figure had risen to 65.6 years, and over 80 years in some countries. *See* World Health Organization Global Health Histories Presentation, "Health, history, and hard choices: Funding dilemmas in a fast-changing world," (August 2006) (available at: http://www.who.int/global health_histories/seminars/presentation07.pdf).

eficial activities can feasibly be reduced to a level that, while not insignificant, is still quite low (less than 1 in 100).[129] If a significant but low risk of serious bodily harm to A can be eliminated only at the price of dramatically increasing the risk of serious bodily harm to B, then it would seem that B has a stronger complaint with a principle that eliminates the risk to A then A has with a principle that tolerates that risk, but reduces it to the maximum extent feasible. That is, it seems unreasonable to ask beneficiaries of activities posing significant risks to accept an increase in their risk of death and serious injury to levels that are perilously high in absolute terms in order that bearers of such risks can enjoy a reduction in their risk of death and serious injury from a level that is significant, but still low in absolute terms, to a level that is insignificant. On this basis, it seems any beneficiary of the wide class of socially beneficial activities posing ineliminable risks of serious harm would have a stronger complaint with a precautionary principle lacking a cap at the economic feasibility point than a risk-bearer would have with a precautionary principle that incorporated such a cap.

In the case of workplace or consumption risks, this conclusion seems to follow with particular force because the risk-bearing worker or consumer benefits—often in a profound way—*from the very activity that places her at risk*. For example, the oil refinery worker who bears a significant risk of death from exposure to benzene, a toxic chemical used in the refining process, would be greatly disadvantaged by a cessation of oil refining activity, since that activity provides her with gainful employment.[130] A safety standard would bring oil refining to a halt unless the risks it poses to workers could be reduced to an insignificant level.[131] Under the feasibility principle, on the other hand, the worker is guaranteed a continued livelihood from oil refining; her benzene risk is reduced to the maximum extent possible without violating that guarantee.[132]

---

[129] *See, e.g.,* Keeton, et al.*, supra* note ___, at 1234-35 ("Typical occupational risk of death in occupations of average risk are 2.7 per 1,000 for all manufacturing and 1.62 per 1,000 for all service employment. Typical lifetime occupational risks of death in occupations of relatively low risk are 0.48 per 1,000 in electric equipment and 0.07 per 1,000 in retail clothing.") (quoting 52 Fed. Reg. 34460 (1987)).

[130] *See generally* Indus. Union Dep't v. Am. Petroleum Inst., 448 U.S. 607 (1980) (discussing OSHA standards for permissible workplace exposure to benzene).

[131] *See id.* at 641 ("[W]e think it is clear that the statute was not designed to require employers to provide absolutely risk-free workplaces whenever it is technologically feasible to do so, so long as the cost is not great enough to destroy an entire industry. Rather, both the language and structure of the Act, as well as its legislative history, indicate that it was intended to require the elimination, as far as feasible, of significant risks of harm.").

[132] *See* Keating, *Pressing Precaution*, *supra* note ___, at 727 ("The argument against shutting down most of society's major productive activities is an argument of fairness—the workers employed by those activities would be harmed in the long run by the elimination of those activities, even though these activities exact a significant toll on the lives and health of those very workers.").

In light of the above considerations, it seems that any beneficiary of the class of socially beneficial activities that pose significant risks of death or serious injury—even a beneficiary who herself is exposed to one or more such risks—would have a stronger objection to any principle lacking a cap at the point of economic feasibility than the bearer of a significant risk would have to a principle incorporating such a cap. In other words, it would be unreasonable to ask someone to forego the manifold liberty-enhancing, life-improving, life-extending benefits of socially productive activities that pose significant risks in order that someone else should be able avoid bearing a significant, but still low, risk of death or serious injury. Thus, I conclude that no one could reasonably reject a precautionary standard incorporating a cap at the point of economic feasibility.

F.   Justifying the Individual Risk Principle Cap on Precaution

What is the ex ante contractualist justification for capping safety precaution at the point dictated by the individual risk principle, i.e., the point at which further investment in safety would reduce the well-being of each cost-bearer more than it would increase the expected well-being of each risk-bearer?

As discussed in Part II, the basic intuition behind capping precaution at this point is that not to do so would mean requiring each cost-bearer to invest in safety precaution past the point at which a risk-bearer would rationally cease investing in her own safety, were she asked to shoulder a single cost-bearer's share of precaution costs. Under ex ante contractualism, could a cost-bearer reasonably reject a principle sanctioning such a tradeoff?

To make this more concrete, consider a variation of the suspension bridge hypothetical discussed in the Introduction. As above, assume the safety netting is an economically feasible safety precaution, i.e., that its use throughout the bridge construction industry would not threaten the industry's competitive stability or long-term viability. However, suppose it is not a city of 5 million that is building the bridge, but rather a town of 25,000. Suppose further that the netting will reduce each of the 2,000 workers' risk of dying in a fall by just 5% (rather than by 50%, as in the original hypothetical). At $10 million, the netting imposes a cost of $400 on each town resident, while lowering each worker's risk of a fatal fall from 1 in 1,000 to (roughly) 1 in 1,053.[133]   Assigning life a value of $5 million, each worker would rationally be willing to pay no more than $250 for such a reduction in her own risk of death.

---

[133] To be precise, if the netting is used, each worker's risk of death falls from 20 in 20,000 to 19 in 20,000.

The individual risk principle would prohibit investment in the safety netting under these circumstances, since the netting would impose a monetary cost of $400 on each cost-bearer for the sake of providing each risk-bearer with a reduction in risk worth only $250.[134]   In other words, the individual risk principle forbids investment in the safety netting because it would set back each city resident's well-being (through a monetary cost) more than it would increase each bridge worker's expected well-being (through increased safety). If no rational worker would be willing to spend $400 to achieve a 5% reduction in her own risk of death, it would arguably be unreasonable to expect a city resident to make that safety investment.

Could a principle that disallowed safety expenditures of this kind—that is, a precautionary standard that capped safety precaution at the point dictated by the individual risk principle—be reasonably rejected under ex ante contractualism? The person most burdened by general acceptance of such a cap would be the *risk-bearer* exposed to higher risks of death or serious bodily harm than she would be exposed to under an alternative principle that did not cap safety investment at this point (e.g., a pure economic feasibility principle). On the other hand, the person most burdened by general acceptance of a precautionary standard that did *not* cap precaution at this point would be the *cost-bearer* forced to shoulder greater safety costs than she would have to shoulder were precaution costs capped at the point dictated by the individual risk principle.

Who has the stronger complaint? It is evident that the cost-bearer does. A representative risk-bearer's complaint with the individual risk cap would have to be based on her exposure to the sort of risk the cap tolerates, i.e., a risk that sets back her expected well-being by less than it would set back each cost-bearer's well-being to eliminate. By contrast, a representative cost-bearer's complaint with a precautionary standard lacking the individual risk cap would be based on her having to shoulder the sort of monetary cost the cap precludes, i.e., a cost that sets back her well-being by more than the risk it eliminates would have set back each risk-bearer's expected well-being. The risk-bearer is complaining about a burden whose elimination would impose a greater burden on the cost-bearer; the cost-bearer, on the other hand, is complaining about a burden whose elimination would

---

[134] It is worth noting that efficient care would also reject the investment in the netting under these circumstances, but for very different reasons. With its focus on aggregate costs, efficient care rejects the netting in this scenario because its total cost exceeds the resulting savings in total accident costs, i.e., installing the netting involves investing $10 million in order to reduce expected accident costs by $500,000. By contrast, under the IFP, the netting is rejected because the burden its use imposes on each individual cost-bearer (a monetary cost of $400) exceeds the burden its nonuse imposes on each individual risk-bearer (a slightly increased risk of death, one which each worker would rationally be willing to accept for $250).

impose a lesser burden on the risk bearer. The risk-bearer cannot *reasonably* refuse to bear that burden.   It would be unreasonable for a risk-bearer to refuse to accept the setback to her expected well-being (i.e., the greater risk) associated with capping precaution at the point dictated by the individual risk principle in order that each cost-bearer can avoid the necessarily *greater* setback to well-being (i.e., the increased monetary cost) associated with pressing precaution beyond that point. It follows that no one could reasonably reject a precautionary standard that capped safety expenditures at the point dictated by the individual risk principle.

The foregoing analysis is premised on the idea that, in determining the reasonable level of precaution, it is appropriate to make straightforward interpersonal tradeoffs in well-being between a person exposed to a risk of serious harm and a person who bears a share of the costs of reducing that risk through safety precautions. Specifically, the cap on safety expenditures imposed by the individual risk principle prohibits taking a specified safety precaution if its monetary cost to each cost-bearer exceeds the amount of money that each risk-bearer would rationally be willing to pay for the resulting reduction in risk.

This sort of tradeoff between one person's monetary burden and another person's risk of serious bodily harm seems appropriate when the risk at issue is imposed as a necessary incident of an activity that provides *both* persons—cost-bearer and risk-bearer—with a direct and significant benefit, as is generally true in the case of workplace risks and consumption risks. For example, in the case of a workplace risk posed by a firm's socially beneficial activity, the worker exposed to the risk and the firm shareholder who bears a share of the costs of reducing that risk are cooperatively engaged in an enterprise that stands to benefit them both, as it provides the worker with a livelihood and the shareholder with a share of profits. The worker's contribution is, among other things, to tolerate the activity's ineliminable risks; the shareholder's contribution is, among other things, to shoulder the costs of safety precautions that reduce those risks. Both contributions are necessary for the activity to move forward, and both involve a reduction in well-being. In light of these facts, and given that both worker and shareholder stand to benefit directly and significantly from the risky activity, fairness arguably requires that their contributions to the activity (i.e., the setback to well-being each must tolerate in order for the activity to move forward) be marginally equalized in the manner required by the individual risk principle. True, there is a sense in which the shareholder, as an owner of the firm that engages in the risky activity, is imposing the risk on the worker for her own benefit, but the risk imposition redounds to the worker's benefit as well, as it is the price of the activity's existence, and the activity benefits the worker significantly on net.

What, though, about a case in which the person exposed to the risk does *not* directly benefit from the risky activity?[135] Can the cap on safety precaution imposed by the individual risk principle defensibly be applied in such a case?[136] In other words, can the same sort of straightforward inter-personal tradeoff be made between one person's monetary burden and another's risk of serious bodily harm when the risk-bearer is a bystander who derives no direct benefit from the activity, e.g., someone who happens to live next to a cement factory that emits carcinogenic fumes into the air, but who neither works at the factory nor consumes its products?[137]

This is a difficult normative question that requires a lengthier treatment than I can provide here. It may be that, when it comes to risks imposed on persons who derive no direct benefit from the risk-creating activity, capping safety expenditures at the point dictated by the individual risk principle is not, in fact, defensible from a contractualist point of view.[138] Even if that were true, though, it would not follow that imposing such risks is indefensible altogether under ex ante contractualism. A principle that categorically forbade the imposition of risks on non-beneficiaries would be reasonably rejectable under ex ante contractualism, since such a principle would foreclose any socially beneficial activity that posed ineliminable risks of harm to persons who derive no direct benefit from that activity. And this would plausibly have a stultifying effect on socially beneficial activity generally, since few socially beneficial activities, it seems, could claim to impose ineliminable risks of serious harm *exclusively* on persons who directly benefit from the activity. If that is true, then such a principle could be reasonably rejected for much the same reasons as a precautionary standard that failed to cap safety expenditures at the point of economic feasibility, i.e., its general acceptance would force most, if not all, socially productive activity to grind to a halt.[139] At the most, then, ex ante contractualism would require

---

[135] *See, e.g.,* Sven Ove Hansson, *Risk and Ethics: Three Approaches* in RISK: PHILOSOPHICAL PERSPECTIVES 21 (Tim Lewens ed. 2007) (identifying a key question in the analysis of the ethics of risk as, "To what extent do the risk exposed benefit from the risk exposure?").

[136] I thank Aaron James for impressing upon me the importance of addressing this question, which I plan to explore in greater detail in future work.

[137] *See, e.g.,* Richard W. Wright, *The Standards of Care in Negligence Law* in PHILOSOPHICAL FOUNDATIONS OF TORT LAW 261-68 (David G. Owen ed. 1995) (advocating more stringent standard of care for situations in which the defendant put the plaintiff at risk to benefit the defendant or some third party than for situations in which the defendant put the plaintiff at risk at least partially to benefit the plaintiff).

[138] *See, e.g.,* Simons, *supra* note ___, at 1212 & n.96 ("[W]here the potential victims do not benefit from the risky activity, many nonconsequentialists would object that a simple benefit/risk or cost/benefit analysis is inadequate to justify the risk, insofar as one party benefits at the other's expense. ... The important point is that, everything else being equal, non-beneficiaries are entitled to greater protection from risks of harm than are beneficiaries.").

[139] *See supra*, Part III.E.

that the risks a socially beneficial activity imposes on non-beneficiaries like the cement factory neighbor be reduced to the extent economically feasible.

G.   Justifying the IFP as a Floor for Safety Precaution

In the previous two sections, I have argued that ex ante contractualism sets the IFP as a moral ceiling on investment in safety precautions that mitigate the risks of serious harm posed by socially beneficial activities. What I have not yet done is explain why ex ante contractualism sets the IFP as a moral *floor* for safety investment in such cases. Why, from a contractualist point of view, must precaution be pressed all the way to the point dictated by the IFP, particularly if, as will often be the case, it will be inefficient to do so?

The reason for this is straightforward. Until the level of safety precaution dictated by the IFP is reached, further investment in safety (i) would enhance each risk-bearer's expected well-being through increased safety by more than it would set back each cost-bearer's well-being through a monetary cost and (ii) would *not* pose a threat to the long-term viability of the activity that gives rise to the risk. It would be unreasonable for a cost-bearer to refuse to make such an investment, which would necessarily burden her by less than *not* making the investment would burden each risk-bearer. It would be unreasonable, in other words, for a cost-bearer to refuse to accept the setback to her expected well-being (i.e., the greater monetary cost) associated with pressing precaution to the point dictated by the IFP in order that each risk-bearer can avoid the necessarily *greater* setback to well-being (i.e., the greater risk) associated with capping precaution below that point. The basic idea is this: when I benefit directly from an activity that poses a risk of death or serious injury to you, it is unreasonable for me to invest any less in safety precautions than I would rationally have been willing to invest were it my *own* bodily integrity that was at stake.

For example, returning to the original suspension bridge hypothetical described in the Introduction, it would be unreasonable for a city resident to refuse to pay $2 so that each bridge worker can enjoy a risk reduction worth $2,500 (recall that if the safety netting is used, each worker's risk of death falls from 1 in 1,000 to 1 in 2,000). After all, were her own bodily integrity at issue, a rational city resident would gladly pay $2 for such a reduction in the risk of death she faces. It would be unreasonable for a city resident to refuse to make this safety investment simply because it is someone else's bodily integrity at stake rather than her own. It follows that, according to ex ante contractualism, no one could reasonably reject a precautionary standard that required investing in safety precaution up to the point dictated by the IFP.

IV. Is The IFP Truly Nonaggregative?

The contractualist justification I have just offered for the individualized feasibility principle does not rest on the sort of interpersonal aggregation that fails to take seriously the separateness of persons. Fried conjectures that, even if contractualism can generate a plausible alternative to efficient care as a decision rule for regulating risky conduct, any plausible alternative will "necessarily be aggregative in one form or another."[140] According to Fried, an aggregative argument is one that "rank-order[s] alternative principles for action at least in part based on their aggregate expected benefits (costs), summed across all potentially affected individuals."[141]

The contractualist argument I present for the IFP simply does not do this. One indicator of this is the fact that, except in rare circumstances, the IFP that emerges from the contractualist inquiry dictates an *inefficient* expenditure on safety precaution (typically an inefficiently *high* expenditure). If generally accepted, the IFP would be expected to result in a lower level of average and aggregate well-being than a welfare-maximizing standard like efficient care.

Of course, as Fried would no doubt observe, the mere fact that the contractualist inquiry settles on a non-optimific precautionary standard does not, of itself, guarantee that aggregative reasoning plays no role in contractualism's endorsement of that standard over the alternatives. We need to look more closely at the IFP to ensure that it does not rely on interpersonal aggregation in some more subtle way. One might suspect, for example, that some sort of aggregative reasoning has to play a role in justifying the economic feasibility principle, which caps precaution costs at the maximum amount the socially beneficial activity will bear. Isn't the point of capping precaution at the feasibility point to preserve the benefits each risky activity delivers, and aren't those benefits worth preserving only in the aggregate?

In other words, doesn't the feasibility principle effectively trade small life-improving benefits to the many against significant increases in risks of serious bodily injury to the few? If precaution is pressed only to the last point at which an activity's aggregate benefits exceed its aggregate costs, aren't we still using the incremental gains in well-being the activity delivers to a large number of individuals to justify the imposition of fatal or devastating injuries on the few? Take an activity like the production of televisions. If the health risks posed to workers in television factories are to be mitigated only to the extent possible without rendering the activity of television production a net loss (and thereby jeopardizing its continuation), it seems the feasibility principle sanctions the imposition of such risks for the

---

[140] Fried, *Limits*, *supra* note __, at 259.
[141] Fried, *Contractualism*, *supra* note __, at 39-40.

sake of preserving the (seemingly modest) entertainment and convenience benefits the typical television consumer enjoys, aggregated across the vast number of consumers who enjoy those benefits.

However, to show that the contractualist justification for the feasibility standard is aggregative in Fried's sense, it is not sufficient to show merely that aggregate (or average) expected well-being is greater in a world governed by the feasibility principle than in a world governed by a more stringent precautionary standard (e.g., a safety standard or a technological feasibility standard). Even if it could be shown that aggregate (or average) expected well-being is *necessarily* greater under a standard that tolerates significant risks of serious bodily injury only to the extent necessary to preserve socially beneficial activities than under a standard directing us to eliminate such risks even at the cost of precluding such activities, this fact need not be the *normative* basis for preferring the former standard to the latter.[142] What would have to be shown is that, by contractualism's lights, at least part of *the reason why* precautionary standards like safety or technological feasibility are reasonably rejectable is that they result in lower aggregate (or average) well-being than a precautionary principle incorporating the feasibility standard.

It seems clear to me that this is not the case. When we compared the generic reasons differently-situated individuals would have to object to the IFP to the generic reasons such individuals would have to object to other precautionary principles, we *never* appealed to the respective principles' effects on aggregate well-being. Indeed, we never appealed to the principle's expected effects on the well-being of more than a single (representative) individual. What we appealed to was the effect that each candidate principle would, if generally accepted and adhered to, be expected to have on the well-being of *each differently-situated individual* over the course of that one individual's lifetime.

This sort of calculus is, of course, *comparative* in the sense that it compares the change in well-being we would expect one individual to experience with the change in well-being we would expect another, differently-situated individual to experience. But it is not *aggregative* because the strength of an individual's complaint with a given principle is never allowed to be a function of the *sum* of multiple individuals' gain or loss in well-being under that principle. Under contractualism, we are comparing one

---

[142] *Cf.* Oberdiek, *Ethics in Risk Regulation*, *supra* note __, at 201 ("It is true, of course, that highways both are but-for causes of many deaths and make the lives of many other people more convenient. Yet, importantly, it does not follow that highways and their accompanying deaths are justifiable in virtue of making the lives of many other people more convenient. That highways have a property, in other words, of 'massive loss to some, minor benefit to many others,' entails nothing about that property's moral relevance in the justification of highways.").

representative individual's complaint under a given principle against one other representative individual's complaint under an alternative principle.[143]

Now, it must be admitted that, under the IFP, the numbers do count in one limited respect. Under the IFP, the number of persons who benefit from a risky activity and the number of persons who bear the costs of reducing the activity's risks can be *determinants of* the amount of the safety expenditure the IFP dictates with respect to a risk posed by that particular activity. For example, the total safety expenditure required by the individual risk principle is based, in part, on the amount each cost-bearer will bear in precaution costs, which is obviously a function of the number of individuals among whom total precaution costs are divided. However, it does not follow that the *justification for* that expenditure involves summing costs and benefits across multiple individuals and allotting normative weight to the candidate principles' respective expected yields in aggregate well-being. That is, under the IFP, the numbers may have calculative significance in setting the requisite expenditure on precaution, but, as noted above, they have no justificatory significance with respect to that expenditure. Under contractualism, the reason that no one can reasonably reject the IFP has nothing to do with its effect on aggregate (or average) well-being. The IFP is not reasonably rejectable because no *one individual's* generic reasons for objecting to it are as strong as the respective generic reasons other individuals would have for objecting to each of the salient alternatives.

V. Conclusion

In this Article, I have introduced and defended an alternative interpretation of reasonable precaution, one that, in my view, deserves to be considered alongside efficient care and other aggregative standards. I hope to have shown that the individualized feasibility principle is capable of serving as a viable substantive decision rule for regulating risky activities, one that can reliably guide decisions about safety investment in specific cases. I hope further to have shown that the deliverances of the IFP will typically differ, often significantly, from those of aggregative standards. In particular, where an activity's risks are concentrated on a group of persons that is a small fraction of the size of the group among whom safety precaution costs are spread, the IFP can accommodate the intuition that acting reasonably may require pressing precaution beyond the point that would maximize aggregate or average well-being. Finally, I hope to have shown that, although the IFP may occasionally yield results some find morally implausible, these

---

[143] *See* Reibetanz, *supra* note __, at 300 (under contractualism, "we determine whether an individual could reasonably reject some principle by making a series of pairwise comparisons: each person's complaint about that principle is compared separately with the complaint of every other person.").

results are no more implausible than those associated with aggregative standards.

If I have succeeded in meeting these goals, Fried's challenge will have been decisively answered. Of course, the ultimate question is not whether the IFP is a viable interpretation of reasonable precaution, but whether it is, all things considered, a *better* one. But this question can only be answered if the most promising alternatives—aggregative and nonaggregative—are brought forward and evaluated against each other in the clear light of day.[144] Aggregationists such as Fried may well reject the IFP in the end, but the debate ought to proceed in recognition that viable alternatives exist on both sides, rather than in the mistaken assumption that nonconsequentialism simply has nothing to offer in the way of a workable decision rule for regulating the risks posed by socially beneficial activities.

Whatever contractualism's flaws as a theory of normative ethics, its insistence that valid moral principles be justifiable to each individual represents what Aaron James has called an "attractive expression of the 'separateness of persons' from a moral point of view."[145] My goal has been to show that the IFP (or something closely resembling it) is the principle that emerges when one approaches the question of how to regulate risky, socially beneficial activities from the standpoint of ex ante contractualism. If that is true, then unlike aggregative standards of care, the IFP can claim to answer directly to the idea that, as Rawls put it, each of us "has an inviolability founded on justice or . . . natural right which even the welfare of everyone else cannot override."[146]

---

[144] *See* Fried, *Limits*, *supra* note ___, at 262 ("If there is a viable alternative, it needs to be put on the table so that it can become part of the public debate over risk regulation and be assessed relative to cost-benefit analysis and other versions of aggregation."). As noted above, the economic feasibility principle had been suggested by Keating as a nonaggregative alternative as early as 2001. *See supra*, Part I.B. The individualized feasibility principle represents a second nonaggregative alternative.

[145] James, *supra* note ___ , at 264.

[146] Rawls, *supra* note ___, at 24-25.