

CONTRACTUALISM'S (NOT SO) SLIPPERY SLOPE

Aaron James*

Department of Philosophy, University of California, Irvine

Familiar questions about whether or how far to impose risks of harm for social benefit present a fundamental dilemma for contractualist moral theories. If 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 fully “*ex ante*” form by considering only *expected* outcomes, it becomes unclear how it recommends something other than aggregative cost-benefit decision-making. Focusing on T.M. Scanlon’s version, this paper develops this basic choice of interpretation and recommends the *ex ante* version. The paper explains how contractualism is inconsistent with John Harsanyi–style utilitarianism and how contractualism supplies a principled framework for walking a careful line between the “bad aggregation” characteristic of utilitarianism and the “good aggregation” that is both unavoidable and fully appropriate in public life.

What should contractualist moral theories in the style of John Rawls or T.M. Scanlon say about the justifiable imposition of risk for social benefit? The answer is surprisingly obscure, and yet it is of crucial importance if contractualism is finally to offer a commonsensical alternative to utilitarianism.

In public life, we often decide whether or to what extent to adopt policies that have the following features: (1) a large number of people stand to benefit; (2) a relative few (often as yet unidentified) are subject to serious risks of harm (whether moderate risk of grave harm, such as death, or high risks of lesser injury); (3) the harm that would befall any of the few, if or when the risks go awry, is dramatically more significant than the potential benefits to any of the many; and yet (4) the policies can be expected to promote the greatest total benefit, aggregating across different lives (i.e., the total social gains will exceed the total social costs, perhaps because total costs are minimized).

*I am especially indebted to Barbara Fried for both written comments and numerous conversations at the Center for Advanced Studies in the Behavioral Sciences, Stanford University, over the 2009–2010 academic year, and to discussions with T.M. Scanlon during that same period. I also owe a special debt to A.J. Julius. For comments or discussion, I thank Michael Cholbi, Kory DeClark, Margaret Gilbert, Greg Keating, Rahul Kumar, John Linarelli, Sharon Lloyd, Steven Munzer, Amanda Trefethen, and Leif Wenar.

Questions of this general kind come in familiar variety. Should we build a bridge, school, or opera house despite the fact that workers will be subject to serious risks of injury or death in the construction process, and, if so, at what cost in workplace precautions? Should we allow clinical trials of a drug that promises to save millions of lives despite serious risks that some participants will have adverse reactions, and is it sufficient if participants enroll voluntarily? Should we raise driving speed limits (e.g., from 65 to 75 mph) despite the fact that this would predictably increase accident rates, if this would boost overall productivity as drivers save time on the road? Should developing countries retain relaxed health and safety standards in order to broaden employment and promote economic growth, even though “sweatshop” laborers desperate for work are more likely to become seriously ill or physically injured? Should we further remove international barriers to trade in goods, services, and capital in order to augment the wealth of nations, despite increased risks of worker displacement, wage stagnation, and financial contagion? And so on.

Utilitarianism tells us in all such cases to maximize the total expected benefits, aggregating across different lives. The risks of serious harm to the few may require efforts at prevention or compensation, but *only* to the extent they optimize interpersonally aggregated welfare.

According to contractualism, any policy for social benefit is subject to a further requirement: it must be justifiable to each and to all on reasonably acceptable grounds. If indeed any of the few can reasonably object to the risks imposed upon them, a policy is morally impermissible at any aggregated opportunity cost.

Much therefore depends on what grounds qualify as “reasonably acceptable.” In the version defended by Rawls and Scanlon, only the interests or personal claims of different individuals are relevant, to the exclusion of direct interpersonal aggregation of gains or costs.¹ The serious risks to the relative few thus cannot be defended simply by summing up the less significant benefits—of a new bridge, or vaccine trials, lax labor laws, or free trade—to what may be millions of people. We only compare the potential benefits to each potential beneficiary, taking them one by one without regard for their numbers. When those benefits are dramatically less significant than the risks to the few, any of the few will be in a position to reasonably object to the imposition.

So elaborated, contractualism offers an attractive expression of the “separateness of persons” from a moral point of view. It is also open to at least two kinds of objection. The first objection is that the resulting ban on interpersonal aggregation is too restrictive. Thus Derek Parfit offers a barrage of counterexamples to Scanlon’s restriction to “personal reasons,”² suggesting

1. JOHN RAWLS, *A THEORY OF JUSTICE* (1971); T.M. SCANLON, *WHAT WE OWE TO EACH OTHER* (1998).

2. SCANLON, *supra* note 1, at 229.

that contractualism should also allow us to weigh in the *impersonal* value of an aggregated social benefit.³ I offer a partial reply to this objection below. As shown below, Scanlon's individualistic contractualism has underappreciated resources for permitting aggregative policies in situations of uncertainty and risk. Our initial focus, however, is a second and perhaps more basic problem that arises when we assume Scanlon's interpretation of the "separateness of persons."

This different problem arises when we consider what exactly the objection of someone subjected to risk for social benefit would be. Is his or her potential objection the *risk* of injury, whatever the actual outcome, or is it the injurious *outcome* itself? As Elizabeth Ashford and Barbara Fried argue in different ways, this choice between these *ex ante* and *ex post* readings is highly consequential.⁴ As shown below, if 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 fully *ex ante* form by considering only *expected* outcomes, it becomes unclear how it is supposed to recommend something other than aggregative cost-benefit decision-making. Either way, contractualism faces a perilous slide into something a lot like the aggregative utilitarianism to which it was supposed to offer a commonsensical alternative.

This dilemma exposes an important way that contractualism stands underdeveloped. Because Rawls's principles of justice are specifically for the "basic structure" of society, they are indeterminate about policy choice *within* the broad constraints of those principles; they apply to activities such as vaccine trials or bridge or opera house construction only indirectly, if at all.⁵ Scanlon's version is supposed to apply directly to policies that impose risk for general benefit, but as shown below, it is unsettled between *ex ante* and *ex post* perspectives in a way that leaves its bearing on public policy unclear. Utilitarianism's most attractive feature, from Jeremy Bentham to John Stuart Mill to John Harsanyi, has always been its applicability in central questions of public life. To the extent that contractualism comes too close to utilitarian-style justification without offering a distinctive alternative, it can seem to fail where success matters most: it cannot claim to rival, let alone replace, utilitarianism (or similar purely aggregative cost-benefit calculations) as a basis for public choice.

3. DEREK PARFIT, *2 ON WHAT MATTERS* (2011), ch. 21.

4. Elizabeth Ashford, *The Demandingness of Scanlon's Contractualism*, 113 *ETHICS* 273–302 (2003), focuses on the *ex post* interpretation. Barbara Fried, *Can Contractualism Save Us from Aggregation*, 16 *J. ETHICS* 39–66 (2012), available at <http://www.springerlink.com/content/f8v117544h876204/>; and Barbara Fried, *Is There a Coherent Alternative to a Cost/Benefit Calculus for Regulating Risky Conduct?* (unpublished manuscript), emphasizes the choice between *ex ante* and *ex post* perspectives. ALLAN GIBBARD, *RECONCILING OUR AIMS* (Barry Stroud ed., 2008), presents an *ex ante*-style challenge to explain why we should reject a Scanlon-Harsanyi marriage, which I consider in detail below.

5. The issue is perhaps part of the "legislative stage" of Rawls's "four stage sequence," though it amounts to a significant further step of application. RAWLS, *supra* note 1, at 195–201.

I admit that “*ex post* contractualism” does slide into something too much like utilitarianism. The suggested dilemma nevertheless fails on its second horn: suitably elaborated, an “*ex ante* contractualism” meaningfully constrains aggregative cost-benefit decision-making. It supplies a principled framework for walking a careful line between the “bad aggregation” characteristic of utilitarianism and the “good aggregation” that is both unavoidable and fully appropriate in public life. Although I can hardly offer a full account of this kind, let alone apply it in a way that settles complex policy disputes, I will clear the ground for the overall approach by highlighting resources, both familiar and underappreciated, for its development.

I. EX ANTE OR EX POST?

Following Scanlon, let us call *contractualism* the view that the rightness (or justness, or moral justifiability) of a policy decision (or its grounding regulative principle) depends only on what each of several different affected parties could reasonably reject, from their respective standpoints, on the basis of personal reasons that could be given on their own behalf. A judgment that someone’s objection to a principle is “reasonable” is a judgment that that the objection is *sufficient* to defeat any other person’s objection to a proposed alternative principle.

Rawls models what is reasonably acceptable to one, in this sense, by what one’s representative would self-interestedly agree to from behind a veil of ignorance. Scanlon instead allows full information and moralizes motivation: reasoning is to be guided by the aim of finding generally acceptable principles. Given that aim, then, even when we know or presume certain information about a situation, it is a further question whether it should count as morally relevant or whether it should be disregarded for purposes of moral reasoning about what different people in their various situations are owed.

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.⁶ Because a *regulative principle* of conduct is at issue, we also count any actual consequences that result because policies of the kind in question are generally allowed.

Let us call *ex ante contractualism* the opposing 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.⁷

6. This is suggested in T.M. Scanlon, *Contractualism and Utilitarianism*, in *UTILITARIANISM AND BEYOND* 122–123 (Bernard Williams & Amartya Sen eds., 1982). Parfit’s contractualism is especially clear about addressing actual rather than expected outcomes; PARFIT, *supra* note 3, at 1:162. See also Derek Parfit, *Justifiability to Each Person*, 16 *RATIO* 368–390 (2003).

7. This is suggested in Gregory C. Keating, *Irreparable Injury and Extraordinary Precaution: The Safety and Feasibility Norms in American Accident Law*, 4 *THEORETICAL INQUIRIES LAW* (2003);

On the *ex ante* version, the relevant kind of epistemic position can be specified in different ways.⁸ In general, however, we do not consider expected outcomes from the epistemological point of view of the affected parties in light of their various actual reasons. We consider only the information available to the *agents* to whom the relevant principles are addressed as expectations of conduct. For unless we work within the perspective available to the agents in question, they will not be in a position to follow and apply those principles in practice.⁹ Scanlon suggests something like this in treating relevant grounds for objection as arising from “standpoints” identified in terms of their various associated “generic” reasons. Such reasons pick out “general characteristics” of a person’s situation not in its full actuality but as it might be *seen*, given “commonly available information about what people have reason to want.”¹⁰

When we assess how I have governed myself, then, what is relevant is not what actually happens to others, or even what others in *their* positions can reasonably expect to happen, but what I can expect or could have expected to happen, given the type of information commonly available in my situation, appropriately characterized. To take an institutional case, in considering a driving speed limit increase, say, the *ex ante* version does not count injuries or deaths that actually materialize for namable individuals and datable outcomes except insofar as they provide evidence, updated over time, about what risks were or are being created by the ongoing institution or activity for different social positions under general descriptions (e.g., “pedestrian,” “driver,” “reckless driver”). We specify the relevant risks in light of actuarial calculation and other information available to the relevant regulatory authority rather than in light of the epistemic situation of, say, the motorist sizing up the risks of driving to the store.¹¹

Gregory C. Keating, *Pressing Precaution beyond the Point of Cost-Justification*, 56 VAND. L. REV. (2003); and Gregory C. Keating, *Pricelessness and Life: An Essay from Guido Calabresi*, 65 MD. L. REV. (2005). Keating puts contractualist argument in terms of expected values but without an explicit limitation to them. Rahul Kumar, *Risks and Wrongs* (unpublished paper), defends a fully *ex ante* version of contractualism.

8. Rawls’s theory is *ex ante* in several respects. First, a decision made from behind the veil of ignorance is *ex ante* in the hypothetical sense that it is assumed to be made “in advance” of knowing one’s actual situation. Second, even looking beyond the veil, outcomes are to be evaluated as just or unjust only in terms of the justifiability of the institutional procedures that produce them. Third, Rawls’s theory is epistemically constrained to the exclusion of an esoteric morality. As Rawls puts the idea, “Conceptions of justice must be justified for the conditions of our life *as we know it* or not at all”; RAWLS, *supra* note 1, at 398, italics added.

9. The *ex ante* contractualist must also firmly deny the possibility of “moral luck.” T.M. SCANLON, *MORAL DIMENSIONS: PERMISSIBILITY, MEANING, BLAME* (2008), 148–151, develops an account of how the phenomenon might be explained away. This may be seen to flow from ideas of self-governance and mutual recognition, as discussed in Aaron James, *The Significance of Distribution*, in *REASONS AND RECOGNITION: ESSAYS IN HONOR OF T.M. SCANLON* (2011).

10. SCANLON, *WHAT WE OWE*, *supra* note 1, at 204.

11. Principles of conduct are thus agent-relative in the sense that some agents but not others may have the assumed informational frame of reference. The risk that I might kill someone on a trip to the store (when I comply with all established regulations) might be too small for me to take account of. But that same risk imposed over thousands of trips to the store might

On the *ex post* version, by contrast, we evaluate a policy not given the high likelihood that many more people will die per se but in light of the actual deaths that result, seen from some period of time after the change or from an imagined perspective that takes the whole of history into view.¹² The idea is not simply that we update our sense of what different parties can reasonably claim as we observe emergent outcomes as a policy works over time; the *ex ante* version can happily admit this, so long as the outcomes are seen as changing our *epistemic* position. On the *ex post* version, the real claim of each party potentially affected is given only by the actual outcome that has resulted or will result, quite aside from the risks and uncertainties associated with its coming about.¹³

Between these opposing views there is a *hybrid* version that admits and balances both expected and actual outcomes.¹⁴ As with the *ex post* version, we evaluate principles in light of their actual outcomes for different parties. As with the *ex ante* version, we add the imposition of risk itself as a separate consideration (e.g., I *nearly* hit you with my car while checking my cell phone and so wrong you, even if no material harm is done). What I owe you, then, depends on how we should balance different considerations both of how I could reasonably expect my conduct to affect you and of how you are in fact affected.

II. MORAL GRIDLOCK

Despite Scanlon's appeal to "generic" reasons and "standpoints," some passages in *What We Owe to Each Other* seem to endorse a fully *ex post* (rather than mere hybrid) position. Scanlon writes:

in considering whether a principle could reasonably be rejected we should consider the weightiness of the burdens it involves, for those on whom they fall, and the importance of the benefits it offers, for those who enjoy them, leaving aside the likelihood of one's actually falling in either of these two classes.¹⁵

be relevant for how regulations themselves are set or adjusted when the larger pattern of risk imposition can be known by a regulatory authority.

12. An *ex post* version can rely on the available evidence in any particular decision, but the independent facts of the case will settle when the right thing is done.

13. The *ex ante* contractualist can, however, identify who is ultimately wronged with who is ultimately injured. The objection is then that the agent acted while *in a position to know that someone in the same (appropriately generically described) situation would likely get hurt*. Note also that other forms of *ex post* decision are not at issue, as when a court is deciding whether someone should be compensated for some harm already done. On the *ex ante* version, the question here is still about what compensatory efforts could be reasonably expected going forward, given the available evidence about what might address the prior wrong done. For the *ex post* variant, what is at issue the harmful outcome and restoration of the victim in and of itself.

14. Sophia Reibetanz, *Contractualism and Aggregation*, 108 ETHICS 296–311 (1998); and James Lenman, *Contractualism and Risk Imposition*, 7 POL. PHIL. & ECON. (2008).

15. SCANLON, *WHAT WE OWE*, *supra* note 1, at 208.

Scanlon admits that it is “intuitively obvious that the likelihood that a form of behavior will lead to harm is an important factor in determining its permissibility.” He then suggests that this does not necessarily require that we “take this probability into account . . . as a factor that, in one way or another, diminishes the complaint of a person who suffers this harm.”¹⁶ Instead:

The probability that a form of conduct will cause harm can be relevant not as a factor diminishing the “complaint” of the affected parties (discounting the harm by the likelihood of their suffering it) but rather as an indicator of the care that the agent has to take to avoid causing harm.¹⁷

It is not, then, the probability of harm but the harm itself that counts.

If we accept an *ex post* contractualism, however, it becomes difficult to explain why any number of socially beneficial and seemingly reasonable activities could be justifiable in practice.¹⁸ Most or many socially beneficial activities will in fact cause at least one death somewhere at some point in the subsequent history of the world, perhaps by an unforeseeable chain of events. But on the *ex post* variant, whenever an activity—commercial air travel, say—brings any risk of serious injury—such as a slight risk of premature death on the ground by crashing plane—we have to consider not a given person’s slight risk of serious harm but rather the serious harm itself. And as long as the harm is indeed serious—and what is personally more serious than premature death?—it will in effect veto any and all activities whose prohibition would not be subject by anyone else to comparably serious complaint. We have what John Broome calls “moral gridlock.”¹⁹

To elaborate, a single death or serious injury will amount to a veto in part because of contractualism’s ban on interpersonal aggregation; we cannot justify the death or injury of the one by the aggregated social gain of air travel and the like. The situation thus becomes analogous to Scanlon’s transmitter case. If an imagined technician—called “Jones”—stands to suffer serious injury unless the World Cup match is taken off the air, we are simply to *disregard* as morally irrelevant the fact that each of millions or billions of viewers will thereby suffer a worsened afternoon. Even if the total of those inconveniences would easily be greater than the loss to Jones, there is no one who enjoys that total gain and so no one who is in a position to defeat Jones’s strong objection to being left to suffer in the transmission machinery.

To be sure, a general prohibition on public goods creation would be much worse for any given person than a mere spoiled afternoon. As Scanlon puts the potential complaint, it would be “too confining” if we were not allowed

16. *Id.* at 209.

17. *Id.*

18. Ashford, *supra* note 4; and Fried, *Can Contractualism*, *supra* note 4.

19. John Broome, *Trying to Value a Life*, 9 J. PUB. ECON. 91–100 (1978); which is cited in Fried, *Can Contractualism*, *supra* note 4.

to do some things that put others at risk.²⁰ Those subjected to risks of injury can therefore ask at most for limitations on certain *types* of risk imposition. But the question then is how to tailor expectations to specific types of conduct, and the *ex post* reading leaves it unclear how any such tailoring could permit life or liberty as we know it. As long as I am left to lead a half-decent life, it is hard to see how even a severe deprivation of my liberty would defeat someone else's complaint of premature death. I might prefer that *I* be given liberty or given death, but "give me liberty or give *him* death" is not a very plausible moral argument.

Scanlon's confinement argument might be strengthened in several ways, but none are especially close to commonsense moral convictions of the sort contractualists usually try to respect. Almost any activity creates a slight risk of serious injury or death (e.g., I start my well-maintained car in the usual way, which *might* trigger a freak explosion that takes out a crowded parking garage). Thus one might argue that unless *pretty liberal* risk creation is allowed, we preclude the very activity that makes a mutually beneficial society possible. Still, there need only be a single case in which someone's premature death will not be justifiable to him or her by the benefits of advanced society, and it seems there will usually be at least one such objector. To take a standard example, the Amish farmer who eschews much of the benefit of general societal life but unwittingly winds up living below an airport flight path will be left worse off, on balance, when he prematurely dies in a plane crash. In that case, it seems that any optional activities, including commercial air travel, would have to be dramatically scaled back. One could perhaps argue that there is no realistic way of ensuring that people lead even "half-decent" lives in a modern society without permitting considerable risk creation and that those who suffer premature death would die even sooner in a violent state of nature or if we somehow reverted to a premodern age. But this argument seems at best a stretch.²¹

Taking a different tack, one might appeal to the voluntary assumption of risk. Perhaps the Amish farmer fairly assumes his risk of death by airplane accident as long as he chooses not to move out of the airport flight path. Perhaps he has an adequate opportunity to avoid his exposure only when his society offers support in this (e.g., housing subsidies or low-interest loans), but if he is given this opportunity and yet chooses to stay put, he takes his life into his own hands. Still, while this might apply in some cases, it is a

20. SCANLON, WHAT WE OWE, *supra* note 1, at 209.

21. I assume the overall argument of this paragraph would not be undercut by normalizing relevant objections. That is, it might seem that any principles for the regulation of conduct can be only so specific about which risks an addressed agent is supposed to track and manage, in which case potential objections to harm cannot be narrowly tailored to specific cases, especially when the risks of harm are exceedingly low. This has considerable force on the *ex ante* version of contractualism but not on the *ex post* version. On the *ex post* version, the potential objection to harm is still the harm that actually happens, while the potential demands of self-governance (concerning gathering information, forgoing valuable opportunities, and so on) are merely countervailing objections. They will not carry the day when premature death is in question.

tall order to suppose that it applies in every case in which a social recluse or other limited beneficiary of society dies by a freak airplane (or train or car) accident. In those cases, an objection to death without sufficient compensatory benefits will retain its full force and defeat the noncomparable objections of anyone else. And again, only one such person is needed for the risks to count as unjustifiable to everyone.

More promising is a "tiebreaking" argument, which cites not the supposed compensatory benefit to the Amish farmer himself but the deaths that would result unless optional activities such as air travel are allowed. Interpersonal aggregation is permissible, Scanlon argues, when it involves similar harms (letting one die to save five from death), or when the harms are similar enough to count as "comparable" (e.g., letting one die to save five from permanent paralysis). And even then, numbers do not count per se. What is relevant is only that unless we take any additional person as "breaking the tie," each of the extra persons on the side of the greater number could reasonably object that his or her presence is not being registered.²² For if one were permitted or required not to save the larger number, the moral situation would be as if there were the same number of people on either side (e.g., one to one instead of one to five). Since extra persons are present, this alone should be sufficient to make it permissible and perhaps required to save the larger group.

To apply this kind of argument, then, it may be said that just as almost every activity creates a slight risk of someone's death somewhere at some point, the prohibition of almost every activity equally creates a slight risk that someone will die when they otherwise would have lived. If each death by falling planes is matched by a death when air travel is banned (some far-flung person then dies without receiving a lifesaving organ), the comparably severe objections are a wash, and the lesser burdens of confined liberty carry the day.

It is not entirely clear what this would mean in practice. The relevant facts will often be difficult if not impossible to ascertain, leaving it at best unclear whether either air travel or our modern way of life comes out as justified or not. Presumably, then, we are simply to do our best to minimize deaths. This is not quite utilitarianism, since we still cannot factor in noncomparable gains (these provide tiebreaking considerations only when comparably more serious considerations of death are a wash). Yet if this is the *only* relevant moral basis for policy choice, *ex post* contractualism seems reduced to something very much like utilitarian justification, at least in one crucial respect: it is equally foreign from the point of view of the commonsense moral convictions that contractualism seeks to capture and explain.²³

22. *Id.* at 232; following FRANCES M. KAMM, 2 MORALITY, MORTALITY (1993), at 116–117.

23. Ashford, *supra* note 4, takes this to support utilitarianism, which is then not "too demanding," at least not by comparison to contractualism.

To see why that might be, consider that commonsense arguably often supports a more direct *ex ante* argument, which addresses risks of death or serious injury in terms of one's own expected prospects within a generally beneficial system of cooperation. Most of us find it acceptable to exempt ambulances from normal traffic rules.²⁴ We find this acceptable despite the fact that we all thereby face increased risks of injury or death by ambulance accident, because each of us stands a good chance of needing expedited passage to a hospital at some point. The acceptability argument need not cite the fact that overall deaths are minimized when ambulances are exempted from normal traffic rules or that overall welfare is improved but only each person's own *ex ante* advantage. We each might insist that further precautionary measures are taken as specific circumstances become known—that the driver who is allowed to run red lights is not equally permitted to mow down the hapless pedestrian in front him when he could instead swerve or slow down.²⁵ If he is distracted at the wheel, he will have wrongly failed to take due care. But if he does all he can to avoid the accident once the situation of heightened risk materializes, even the killing may be justified, because the more general practice, perhaps in conjunction with situational precautionary expectations, is justified on grounds of its *ex ante* benefits to all.²⁶ Although the practice will not feel especially sensible when we turn out to be the hapless soul who frozenly watches as he is struck, many are prepared to regard such outcomes, even *ex post*, as tough bad luck rather than a failure of society to take due care (e.g., in deathbed reflections, one says to oneself “well, these things do happen”).

Likewise, and more generally, we may see good prospects of longer-term gain from commercial air travel, despite the risks of falling planes; from wide vaccine distribution, despite the risk of an adverse reaction; from bridge or school or opera house construction, despite a risk of construction accident; or from a wide net of practices that cumulatively make for a generally prosperous society despite any number of risks. Even in the odd case in which we are subject to risks without direct corresponding benefits, many will savor the blessings of liberty; we nevertheless stand to gain from participation in a larger scheme of conduct that affords expansive liberty to act in ways that impose risks upon others for our own personal gain.

It is hard to see how this style of argument could be sustained, however, if we consider only actual outcomes as according to *ex post* contractualism. Complaints of death will always carry the day.

24. The example is due to KAMM, *supra* note 22, at 303.

25. As in “Ambulance II,” from FRANCES M. KAMM, *INTRICATE ETHICS* (2007), at 273.

26. The justified background exemption from traffic rules means the driver could not have been expected to avoid creating the dangerous situation. He *can* justify having set in motion the process that may turn out to have a fatal moment of carelessness; in that case, it is *only* the moment of carelessness that is unjustified.

This suggests that contractualism should at least take a hybrid form; it should at least admit considerations of expected benefit and harm as relevant to justification. But because the hybrid view also admits the relevance of actual outcomes, it remains unclear how the above commonsensical arguments will go through. As long as someone's relevant ground for complaint is not his or her risk of death but death itself, it will not clearly make a large difference if we tally in his or her own expected gains. When expected gains to others are the concern, the complaint of premature death will defeat any noncomparable benefits, expected or received, which could be adduced as countervailing considerations.²⁷ If the argument is instead supposed to rest wholly on comparable complaints (of death or very serious injury), we are back to a utilitarian-style justification that justifies policies only by comparing death rates.

Ex ante contractualism, by comparison, readily accounts for such examples. We then count not benefit or harm but *only* expected benefit or expected harm. So the Amish farmer's relevant complaint is not his death (if he dies) but the *slight* risk of death to which he is exposed. And if his risk is only slight, it becomes hard to see how he would have a reasonable objection to our modern way of life—especially if his risk of injury has been substantially reduced, if he has been given an adequate opportunity to limit his exposure, and/or if he does invariably benefit to some extent from being in or around a postindustrial age (e.g., by having a longer expected life span). We can likewise accept the suggested argument for ambulance exemptions as well as its extension across a broad range of cases in which risks are imposed for the good of all.

It is important that *ex ante* contractualism will equally ignore the actuarial certainty of *someone's* death, lest it wind up in the same boat as the *ex post* version. Suppose, for example, that placing additives in the groundwater can be expected to benefit everyone modestly (e.g., dental benefits from fluoride placed in the water supply). While any given person's risk of injury is slight, suppose we know with virtual certainty that *someone or other* will be killed by an adverse reaction, although we cannot know in advance, in actionable terms, who that person will be. Would fluoridation be permissible? It may not seem so. Is not someone being unreasonably asked to accept modest dental benefits in exchange for his or her death? And if so, is not an *ex post* or at least a hybrid view appropriate after all?²⁸

27. Lenman, *supra* note 14, at 116, suggests that reasonable *ex post* objections will take into account considerations of what the agent could have known and the ensuing strains of commitment, so as to blunt the contrast between *ex post* and *ex ante* objections. As suggested in note 21 *supra*, such burdens of governance would still only be countervailing considerations, and it is hard to see how they would be decisive in the face of powerful *ex post* objections to premature death. They gain greater force within a fully *ex ante* framework that compares risks of death rather than the bare fact of death itself.

28. Versions of this question are pressed in different ways by Reibetanz, *supra* note 14; and Marc Fleurbaey, *Assessing Risky Social Situations*, 118 J. POL. ECON. 649–680 (2010).

The *ex ante* contractualist might hold his or her ground as follows. When risks fall only to the expected beneficiaries, the case is arguably the same as a single-person fluoridation case: you could justifiably fluoridate a single person's water for the sake of dental benefits if his or her risks of serious injury or death were very small. In the same way, we are fine with speed limits that allow small or even significant risks of our being seriously injured for the sake of saving time in a trip to the store. We might then say the same thing in the multiperson version: if each and every expected beneficiary of the fluoridation policy faces only a small risk of death with good prospects of benefit, the high likelihood that someone or other will die will not change the moral situation.²⁹

To be sure, the multipersonal case *would* be relevantly different if special vulnerabilities to adverse reactions could be attributed to certain parties identified under some actionable description (e.g., "persons with such-and-such treatable disposition to adverse reaction," rather than merely "person who is killed"). Then special precautions might be required. They also might be needed for those who cannot be expected to benefit from the policy. If someone's risk of death is high enough, the policy might even be unjustifiable despite the benefits for others, short of some further special justification, such as comparable harm to others or voluntary risk assumption. If, on the other hand, the risk of death is low, the case becomes that of the Amish farmer, in which we do not have to forgo modest but likely benefits in order to prevent slight risks of serious injury or death. That leaves a range of harder cases in which significant risks must be compared with other significant risks or prospects of gain (of which more below). For now, the point is merely that we can begin to see how the bare fact that someone or other will almost certainly die might not be relevant to morally justified action *per se*. Because it is a general reality of almost all public policy choices, it is of no particular relevance to one choice rather than another. It becomes of particular relevance only when the available information affords some way of making a difference to what deaths ensue.

I tentatively conclude, then, that *ex ante* contractualism is more likely to capture commonsense morality than either *ex post* or hybrid contractualisms.

III. EX ANTE CONTRACTUALIST UTILITARIANISM?

Although contractualism is our main focus, I should emphasize the generality of our problem. It arises in a more severe form, for example, for welfare consequentialism. Welfarist views that require the compensation of losers (e.g., as required by Pareto efficiency) have to reckon with the fact that we cannot, alas, compensate people after they have died. One standard move is therefore to shift to a purely *ex ante* welfarism: only an *expected* welfare

29. Johann Frick, Contractualism and the Ethics of Risk (unpublished manuscript), nicely develops this thought.

improvement for all is required, even as some will, of course, ultimately be made and left worse off. But this move can seem arbitrary in a way that it is not arbitrary for the *ex ante* contractualist.

For one thing, insofar as expected welfare is decided by so many gambles, it has no automatic moral relevance; we need some further reason why what we owe someone is changed by what risks he or she would take. Nor can one simply invoke the value of a person's welfare as such, which is at bottom an *ex post* notion. Even if voluntary gambles can in some cases serve as a proxy for a person's good, the underlying concern is with how well or badly a person's life *actually goes*, quite aside from what would fulfill our best welfare expectations (which can turn out to be seriously wrong).

Contractualist views, by contrast, begin from a different underlying moral concern: not welfare or outcomes per se, but what treatment is justifiable to each and to all, where justifiable treatment depends entirely on what is within our regulative powers. On a natural reading, we judge what others have done to us based not on worldly outcomes per se but on what they knew or could have known would happen and the risks they took with our lives. When that is not the ultimate concern, mere bad outcomes are simply bad luck (or at least a fundamentally different, perhaps *ex post* kind of moral concern). The *ex ante* perspective thus has a natural moral rationale.³⁰

This brings us to the second horn of our basic dilemma: that *ex ante* contractualism courts a different kind of slide into something too much like utilitarianism. To see this, consider how the case of ambulance exemptions might support an aggregative rule of social choice in light of John Harsanyi's famous equiprobability theorem.³¹ According to the theorem, an aggregative rule of social choice follows under (roughly) the following conditions: (1) each is as likely as any other to enjoy the benefits and/or suffer the costs (e.g., to need an ambulance and/or wind up in an accident); (2) each is concerned to do as well for him- or herself, prospectively, as he or she can (by making premature death as unlikely as possible); and (3) each is willing to tolerate a measure of risk to that end (each of us can live with one's modest risks of death or injury in a large population, given our good chances of needing speedy rescue at some point). Ambulance exemptions count as justified, then, because when we tally up the exemption's total expected benefits (e.g., in lives saved), and subtract out the total expected costs (in deaths, cars wrecked, and productive time lost), we come up with a greater resulting overall value than we would for the situation in which ambulances follow stricter rules.

While this is only a theoretical possibility so far, I suggest above that it is quite commonsensical to take the ambulance case to be representative of a

30. For development of this theme in light of the idea of "moral recognition," see James, *supra* note 9.

31. John Harsanyi, *Cardinal Utility in Welfare Economics and in the Theory of Risk Taking*, 61 J. POL. ECON. 434–435 (1953). See JOHN BROOME, *WEIGHING GOODS* (1991), at 51–58, for discussion and criticism.

wide class of public policy choices. An aggregative rule of social choice thus seems to provide a correspondingly wide basis for policy choice, and in a way that is not only supported by common sense but also fully consistent with the “separateness of persons.” And to the extent such a rule is widely applicable, it may be said that *ex ante* contractualism bears the burden of explaining why and when, if at all, it should be rejected. To the extent it lacks the resources to do so in an informative, non-ad hoc way, an aggregative rule of social choice may seem the natural default.³²

Harsanyi’s famous challenge to Rawls can be understood in these terms.³³ While Harsanyi’s equiprobability is itself merely a formal construct, it poses the question: if we can tweak Rawls’s favored construction and produce utilitarianism, why draw the veil in Rawls’s way? Why assume parties behind the veil decide under *uncertainty*, without assuming they have an equal chance of being anyone, while being strongly averse to risk? Rawls has answers for his particular subject matter, the basic structure of society.³⁴ But as noted earlier, this special case at best indirectly addresses the more general policy contexts of our concern. And it will not necessarily help to invoke Scanlon’s more widely applicable version. As Allan Gibbard emphasizes, the challenge equally applies.³⁵ Assuming that no one can reasonably reject the rule of social choice that would be selected from behind Harsanyi’s veil, utilitarianism becomes what we owe to each other. The question then is: Why assume otherwise?

Scanlon does defend his full information characterization.³⁶ Even if his arguments are successful, however, this is modest advance by itself. As Gibbard argues, the challenge can also or instead be posed in terms of Harsanyi’s “second theorem,” which makes no use of a veil of ignorance. It requires (roughly) only the following three assumptions: (1) each person’s preferences are coherent (they satisfy the axioms of expected utility theory); (2) social preferences are also coherent; and (3) a Pareto condition is satisfied (roughly, if one person prefers a state of affairs, everyone else is at least

32. Fried, *Can Contractualism*, *supra* note 4, suggests this at several points, though in conjunction with her several challenges to proposed limitations on aggregation.

33. Harsanyi, *supra* note 31.

34. For discussion, see JOHN RAWLS, *JUSTICE AS FAIRNESS: A RESTATEMENT* (Erin Kelly ed., 2001), at 106–109. When a veiled decision concerns one’s whole life prospects, and there are grave prospects of profound disadvantage under pervasively influential basic institutions (e.g., being enslaved in a utilitarian society), what is socially acceptable from a moral point of view is not adequately represented by a rational gamble under conditions of risk when probabilities are known (to be equal, in this case) and one is willing to take one’s chances (with a gambler’s modest aversion to risk). The more appropriate model is what one would choose under conditions of fundamental uncertainty (the parties simply do not know who they are), with a correspondingly heightened aversion to winding up in a permanently disadvantaged state. Rational choosers then follow the maximin rule of choice and thus reject an aggregative principle of social choice in favor more egalitarian principles. Rawls’s claim is that only the more cautious representation will explain to those who wind up in the worst position why, save the demands of impartiality, they are being given their full due.

35. GIBBARD, *supra* note 4.

36. Scanlon, *Contractualism*, *supra* note 6, at 120–128.

indifferent to it).³⁷ Assuming these assumptions can be couched within Scanlon's contractualism, the result is a utilitarian principle that is, again, consistent with the "separateness of persons."³⁸ We need only take the strength of relevant personal complaints to be solely a function of each person's good, seen either as expected welfare,³⁹ as what John Broome calls each person's "personal good," or as what Gibbard calls each person's "goal scale."⁴⁰

Should we accept the proposed Scanlon-Harsanyi marriage? If we must speak now or forever hold our peace, our main concern is as follows: the two are not right for each other. We take Harsanyi's two theorems in turn below.

As regards Harsanyi's equiprobability result, our reply is that there is simply no specifically moral reason to think that it applies even in the first instance, and certainly not with full generality from a contractualist point of view. And until some such moral grounding is provided, there is no challenge to explain why it should be rejected.⁴¹

An initial general problem is that Harsanyi does not clearly work in the right information space. What gambles self-interest optimizers would be willing to take in their respective circumstances will not necessarily be morally relevant. As noted above, when we seek to justify a principle for conduct, what will be relevant is what the addressed *agent* of the conduct could be expected to know and account for, given his or her epistemic situation, about the interests of others (e.g., what a government can know about accident rates in setting speed limits). Because there will often be only so much an agent can take into account in a particular action, the relevant informational framings often will not involve gambles based on specific calculable probabilities but rather nonspecific considerations such as "If I/we do that, people could really get hurt," or "If I/we do that, those people would very likely die."

And even when specific probabilities of injury are known or measurable, and not mostly uncertain, it will be a further question how generic interests are properly described. The appropriate characterization, as guided by the

37. John Harsanyi, *Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility*, 63 J. POL. ECON. 309–321 (1955).

38. GIBBARD, *supra* note 4, at 50–52, 59–60.

39. See the "Complaint Model," which SCANLON, *WHAT WE OWE*, *supra* note 1, at 229, discusses and rejects, as well as the discussion in Reibetanz, *supra* note 14. SCANLON, *WHAT WE OWE*, *supra* note 1, ch. 3, rejects the assumed welfarism, admitting as morally relevant only aspects of welfare, along with other kinds of complaints.

40. BROOME, *WEIGHING GOODS*, *supra* note 31, chs. 8 and 10; GIBBARD, *supra* note 4, at 42–43, 64–66; and Broome's formulation at BROOME, *WEIGHING GOODS*, *supra* note 31, at 114–115. It is worth emphasizing that in any case, a workable basis for interpersonal comparisons may have little relation to traditional utilitarian concern with personal welfare.

41. BROOME, *WEIGHING GOODS*, *supra* note 31, at 58, summarizes his extended development of the point as follows: Harsanyi's equiprobability argument "does not really yield philosophically interesting information about the nature of morality. Instead, it makes a philosophically unsupported assumption about the nature of morality." For a similar criticism, see BRIAN BARRY, *THEORIES OF JUSTICE: A TREATISE ON SOCIAL JUSTICE* (1989), at 334–335.

aim of finding principles that could be justified to all, need not line up with what would optimize anyone's rational self-interest or their *ex ante* personal good. Any such characterization must be defended in light of the situation at hand. In the fluoridation and ambulance cases, for example, it may well be relevant that each stands on balance to benefit, *ex ante*, from a policy. But the claim that this is so is a claim about what is relevant to what such people are owed, as described, again, given the information available from the agent's point of view.⁴²

While this is not to deny that a Harsanyi framing might still be appropriate in some range of cases, it does mean that Harsanyi's first theorem is more limited than it initially appears: it implies only that an aggregative "rule" of social choice is *formally coextensive* with an applicable principle for some relevant range of cases. That by itself says nothing about why an imposition of risk would be justified. When it is justified, because everyone potentially affected has roughly the same *ex ante* prospects of benefit (whether from ambulance exemptions, fluoridation, organ or vaccine distribution, or the blessings of liberty), this will be so only because *no one can otherwise mount a reasonable personal objection to this*. In principle, any number of further relevant grounds for objection might enter. Quite aside from marginal cases such as the Amish farmer, people often have special vulnerabilities even when they have "roughly" similar *ex ante* prospects, and so will often have reasonable grounds either for expecting special precautions or for a real choice in the risks to which they are imposed. Indeed, we can respect common sense and yet notice ways it overestimates the extent to which people have the same or very similar interests. The basic moral question is not whether we *can* squint and see everyone as basically in the same *ex ante* situation in *some* respects, but whether there are also relevant differences even among generically described positions (e.g., "working mother" versus "worker"), given our general accumulated knowledge of social life. To the extent people's situations should be differentiated, the Harsanyi result will not apply. It will apply, as a default or otherwise, *only* when such special considerations are also being adequately addressed.

Turning to Harsanyi's second theorem, the proposed suitor is an even worse match for Scanlon's contractualism. The second theorem is flatly inconsistent with the separateness of persons as represented in the very structure of Scanlon's theory. It assumes the axioms of expected utility theory, including what is often called the "sure-thing principle."⁴³ But the sure-thing principle is false as long as we assign significance to the *distribution*

42. Thus, as Scanlon explains, in considering relevant grounds for mutual aid, we ignore the fact that the fortunate Joneses are unlikely to at some point themselves require assistance. We instead consider their relevant interests under more generic descriptions such as "the importance of being able to get aid *should one need it*" and "the degree of inconvenience involved in giving it, *should one be called upon to do so*." SCANLON, WHAT WE OWE, *supra* note 1, at 207.

43. See BROOME, WEIGHING GOODS, *supra* note 31, at 94–95.

of risk, as we must in justification to each person taken separately. In that case, a conception of rationality that required indifference to risk distribution would be false, or at best morally irrelevant from a contractualist point of view.⁴⁴

To elaborate, consider the counterexample to the sure-thing principle proposed by Peter Diamond.⁴⁵ Two people in a hospital require a kidney transplant to live. There is only one kidney available, and the hospital administrator must choose between giving the kidney to one of the patients and distributing it at random, say, by a lottery that gives each an equal chance. The sure-thing principle implies that we should, as a matter of rationality, be wholly indifferent to the distribution of risk. We should have no reason to prefer the lottery, which more equally distributes death prospects. Diamond suggests, however, that we might well rationally prefer the lottery because it gives each "a fair shake." In that case, the sure-thing principle is false.⁴⁶ Nor is Diamond's counterexample a marginal case. As Broome explains:

Diamond's stylized example represents a widespread practical problem. For a given quantity of risk, is it better to have it more rather than less equally distributed? Radiation leaking from the nuclear power stations will kill a number of people. Should nuclear policy be designed so that the risk of death is evenly distributed across the population, rather than concentrated on a smaller group of people? Suppose one hundred people will die; is it better to have ten million people exposed to a .00001 chance of dying, than ten thousand exposed to a .01 chance?⁴⁷

Here, as with the kidney example, there is nothing irrational in being concerned with how risks are distributed.

Indeed, the very structure of *ex ante* contractualism implies that any relevant notion of rationality will be sensitive to risk distribution. We need only assume that each of the transplant patients and everyone exposed to nuclear radiation have a morally relevant interest in being less subject to the risks in question. Any policy choice will have to be justifiable to each person affected in light of such interests, and so, to the extent the risks could be reduced for any one party, it follows that, barring special justification, the failure to do so must be justified in terms of expected benefits to someone else. In that sense, the distribution of risks among different parties is an essential moral concern. The axioms of expected utility theory will not then be satisfied, and Harsanyi's second theorem *cannot* be situated within Scanlon's contractualism.

44. I am grateful to A.J. Julius for drawing my attention to this line of objection.

45. Peter Diamond, *Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility*: Comment, 75 J. POL. ECON. 765–766 (1967).

46. For detailed explication, see BROOME, WEIGHING GOODS, *supra* note 31, at 111.

47. *Id.* at 111–112.

Thus, much as with his equiprobability result, Harsanyi's second theorem only *seems* to generate a burden of justification. In fact it has no moral resonance, presumptive or otherwise, from a contractualist perspective. So there is nothing objectionable in assuming that it does not apply. In this clean and dry formal space, there is no slope to slip on.

The point here is general. Any *ex ante* utilitarianism or purely aggregative cost-benefit procedure will be insensitive to the distribution of risk per se. Concerns of risk distribution must be added, if at all, as supplementary considerations by a thumb placed on the cost-benefit scales. *Ex ante* contractualism, on the other hand, renders risk distribution relevant by virtue of the very structure of justification (at least given very modest assumptions about people's interests). To underline the point, consider the following two cases from James Lenman.

In a first case we compare the following two policies:

In pursuit of a certain public good, G, a certain government kills twenty of its 20 million citizens.

In pursuit of a precisely similar public good, G, a certain government imposes on each of its 20 million citizens a one-in-one-million risk of being killed.

Without further details, these policies are equivalent from an aggregative cost-benefit perspective. But they surely are morally different (even if G cannot be obtained in a less costly way, especially if there is no extenuating emergency situation).⁴⁸ The *ex ante* contractualist can explain why as follows: the twenty who face government-instituted killing face a high probability of death, and therefore each has a significant complaint which is lacking for each of the 20 million who are all exposed, under the alternative policy, to only a slight risk of termination (they could have been among the chosen few). Assuming there are no further alternatives, the claim of each of the twenty to have their risk dramatically reduced will outweigh the claim of the others against having their risk slightly increased. Indeed, if the subjects are each similarly situated, the risk must be shared equally.

In Lenman's second case, we are again considering imposing some risk upon 20 million people, but instead we suppose that in the absence of any precautions, each faces a one-in-500,000 risk of death.⁴⁹ We have either of two precautionary measures:

Reduce the risk of each of the 20 million to one in 1 million.

Reduce the risk to 19 million of them to one in 19 million, while increasing the risk to the remaining 1 million to one in 100,000.

48. Lenman, *supra* note 14, at 100.

49. *Id.* at 107.

The first precautionary policy has a higher aggregative cost-benefit score (it reduces expected mortality from forty to twenty, whereas the second only reduces it to eleven). But the two policies are morally quite different. According to *ex ante* contractualism, any such redistribution of risk would require a justification to those subjected to increased chances of harm. A justification might well be provided: perhaps everyone is given an adequate opportunity to avoid the increased risk, but 1 million people choose to assume it. Yet the mere fact of an aggregated gain would not be sufficient. And if no such special justification is forthcoming, the move away from equality of risk imposition would be unjustifiable to each of those subject to greater risk.

IV. PERMISSIBLE AGGREGATION

Although nothing here precludes aggregation per se, given appropriate grounds, it does suggest that there is no especially slippery slope into unwanted forms of it. Aggregative decision-making not only *requires justification*; the most natural grounds for its justification are ruled out: we still cannot cite *impersonal* reasons concerned with how the social world can be expected to go over time from no particular person's point of view. A defense of aggregation must, then, take either of two forms: one must defend the moral relevance of purely impersonal values as a proper goal of public policy (even at what will be, for some, great personal expense) or argue that the appropriate specific circumstances arise in light of the personal prospects of those potentially affected. Either way, the right to make an aggregative decision is not the moral default.

That is the beginnings of a reply to our initial dilemma. Even if we reject *ex post* and hybrid contractualism, there is no *general* reason why *ex ante* contractualism faces an unwanted slide into aggregative justification of the kind it was originally designed to resist. Still, this is not yet to address hard questions of how particular risks, in particular policy contexts, are to be identified and weighed. While Fried presses our dilemma in part to raise questions about contractualism generally, she also means to highlight a vast range of garden-variety public-policy choices in which, she argues, contractualism has no alternative to aggregative forms of decision. Fried's argument turns in large part on her skepticism about particular ways of limiting aggregation rather than the general framework of contractualism itself. I am not trying to answer her skepticism on this score. My present point is more general: contractualism can happily welcome aggregation—provided it is justified for the right sorts of reasons—while remaining fundamentally distinct from utilitarian and other welfarist views.

This is significant because it may go a long way toward answering the important objection, noted above, that contractualism is unduly limited in that regard. As I now explain, the *ex ante* version permits aggregative decision

in several quite different ways for different fundamentally nonaggregative reasons while maintaining its ban on unrestricted aggregative reasoning more generally.⁵⁰ Consider several of these different ways in turn.

A. General Expected Benefit

Above I suggest that situations in which all stand, *ex ante*, to benefit—as in the fluoridation or ambulance cases—might support an aggregative decision rule on contractualist grounds. I emphasize above that the case for this depends on addressing asymmetrical risks (or other grounds for complaint) but without settling how those risks are to be identified, especially in light of background risks or prospects of benefit. As I also suggest in discussing the Amish farmer, it remains open to the *ex ante* contractualist to defend significant *intrapersonal* aggregation, that is, to justify specific risks as consistent with “due care” because compensated by expected benefits in a larger practice of risk imposition. A man’s health might be put at risk because a bridge construction project is degrading his water supply, for example, but if the principle that governs the risk imposition equally covers activities that can be expected to benefit him (e.g., school and opera house construction), then he might be seen as a net expected beneficiary of the general permission to impose a class of risks throughout society.

The question, then, is how such “bundling” might be constrained. Even if one is better off in society than in a state of nature, that will not show that *any* imposition of risk satisfies the standard of due care. Still, one might therefore worry: Is there not an interpersonal slippery slope that allows almost *any* imposition of risk, given some supposedly compensatory expected benefits of social life identified under some very abstract description?

Ex ante contractualism lacks a general answer to this question. It mainly leaves the matter to case-specific judgment. It can also plausibly invoke certain general burdens of intrapersonal justification. For one thing, if “bundling” is to be justified, we have to expect that those affected will see benefits *over the course of their lives* rather than in the economist’s “long run” (when we are all dead, as Keynes famously quipped). Moreover, credibility is crucial. Serious risks will not be justified by benefits that are unlikely, even if they would be large if received. Likewise, as the size of the potential gain decreases, its likelihood should increase. So, for example, given that the removal of barriers to international trade promises only modest standard-of-living increases for low-skilled workers, trade liberalization arguably must be accompanied by social insurance and redistributive schemes that virtually ensure those promised gains are received (in contrast, say, with recent U.S.

50. Among “aggregative decisions” we might include any aggregative cost-benefit calculation, including policies that require them of officials (perhaps in particular specified forms), whether or not the decision’s licensing principle is itself aggregative in either its content or underlying justification.

economic history, in which international trade, technological change, and weak social safety nets have led to decades of stagnating wages).⁵¹

Still, there will presumably be many cases in which risks are not symmetrically imposed, even overall. In that case, *ex ante* contractualism can seem to face its own problem of “moral gridlock.” Many projects aimed at general benefit, such as stadium or opera house construction, arguably create risks that cannot be justified, either by “bundling” or by *ex post* measures of compensation, even as they offer only mild entertainment benefits spread across large numbers of people. Will not those who are (on balance) subject to greater risks be in a position to reasonably reject the activities being permitted, and might this render impermissible many or even most activities we take for granted?⁵²

Above I express worry that *ex post* and hybrid contractualism sharply break from common sense by in effect rendering such decisions a matter of “tiebreaking” aggregation of deaths or serious injuries. My suggestion is that *ex ante* contractualism fares far better in this regard because it admits arguments from general *ex ante* benefit and leaves considerable room for bundling of intrapersonal risks without being hobbled by *ex post* outcomes. This may go a long way toward answering the concern that too little of what we take for granted will be allowed, but without foreclosing plausible limits on how much risk can be imposed, even at a significant aggregated opportunity cost. We find just this general kind of balance, for example, in central areas of U.S. accident law.⁵³ In some cases, a “safety” standard prohibits any “significant” risk, but without banning the imposition of risk all together (“insignificant” risks are allowed, even as they may in fact cause injury). In other cases, a still weaker “feasibility” standard applies, which allows “significant” risk for the long-term survival of the activity but still requires levels of precaution that may come at a huge cost in forgone economic gains. In this and other cases, there remain in a range of different policy contexts difficult questions about which risks matter and how. But in conjunction with the other possible justifications below, this at least suggests how there might be no general problem of gridlock.

B. Comparability

If my point is that “tiebreaking” interpersonal aggregation need not be the sole basis for policy, that is not, of course, to deny the prevalence of cases in which it is indeed the appropriate concern. In those cases, the *ex ante* approach has a further theoretical leg up: it can expansively characterize the “comparability” of harms.

51. For discussion, see AARON JAMES, *FAIRNESS IN PRACTICE: A SOCIAL CONTRACT FOR A GLOBAL ECONOMY* (2012), ch. 7.

52. Fried presses this objection in Fried, *Can Contractualism*, *supra* note 4, sec. 3.2.

53. Keating, *Irreparable Injury*, *supra* note 7.

To see this, recall Scanlon's transmitter case, which assumes more or less certain outcomes. For *ex ante* contractualism, certainty is merely a limiting (and perhaps rare) case along a probabilistic continuum.⁵⁴ When similar or comparable harms have roughly the same chances, and the case is otherwise the same, we reach the same result. If there is a 50 percent chance that Jones will suffer potentially fatal electric shocks while being lodged in the World Cup transmission machinery, and there is a 50 percent chance that millions or billions of viewers will suffer a ruined afternoon, it will still count as impermissible to leave Jones to the whims of electrical fate. This will be so even when it can be expected to come at a huge aggregated cost-benefit loss.

That will not be so, however, when the risks of harm count as "comparable." Scanlon suggests that death is comparable to total paralysis but not comparable, say, to having one's legs broken. So while one could save a larger group faced with total paralysis instead of a smaller group faced with death, one would be forbidden from similarly saving a larger group of people facing broken legs. But suppose now that a smaller group faces only a 30 percent chance of death while the larger group faces a 90 percent chance of total paralysis. Then the case for saving the would-be paralytic is strengthened: in a choice between a larger and a smaller group, any of the extra people in the larger group could appeal to that greater prospect of harm as essential if their presence is to be adequately counted.

But the same might be true when we are instead choosing between death and a harm of a much lesser magnitude, such as broken legs. One might be required to save a single person who has a 98 percent chance of suffering two broken legs instead of a single person faced with a 5 percent chance of death. Accordingly, one might be permitted to save a large group of people from high chances of broken legs instead of saving a smaller group faced with small risks of death.

How much aggregation this licenses will depend on the case. The argumentative resource does at least suggest that contractualism is less limited than Parfit and others assume.

C. Risk Assumption

I mention above a different class of cases in which special vulnerabilities to injury may be justified: they may be justified not by intrapersonal aggregation or by appeal to comparable risks or prospects for others but as voluntarily assumed. That is, we can promote the interpersonally aggregated

54. Reibetanz, *supra* note 14, at 301–302, explains how Scanlon's limits on aggregation can be rephrased entirely in terms of expected benefit and harm: taking regulative principles in light of their expected consequences over time, one's prospects are assessed in light of one's overall chances, over time, of falling into groups of larger or smaller numbers of persons who could require assistance.

greater good, even when the benefits are modest but widely spread, in the name of the person's own decision to assume the resulting risks (or his or her failure to avoid them), given an adequate range of alternatives. Bridge construction projects could then go ahead for the greater good, despite our knowledge that workers are highly likely to die, so long as the workers voluntarily take the jobs, being fully aware of what they are getting into, with adequate alternative lines of work.⁵⁵

The question is then when employment opportunities are "adequate." The law and economics school will say that opportunities are adequate when and only when they are aggregatively "cost-justified" (i.e., when a further dollar spent on precautionary measures would yield less than a dollar's worth of expected accident costs). For the contractualist, however, this position cannot be maintained simply by appeal to the impersonal value of greater efficiency or higher average standards of living, or by any other form of interpersonal aggregation not justified on some personal ground that contractualism allows. A worker who takes a job that poses serious risks of death, where this is his only way to feed his family given his sharply limited labor opportunities, may be said to not to have "voluntary" assumed the risk in a sense that legitimates outcomes such as his death, whatever the cost-benefit argument is in its favor. If subjecting him to dangerous work is to be ultimately justified, it must be justified in other terms.

Other arguments might of course be available. There may be a risk-sensitive "tiebreaking" argument: a worker might be conscripted into war or other dangerous work in order to reduce the comparably serious risks to others (where the potential harm may be the same, or simply less serious and far more likely). In other cases, such as with "sweatshops" in developing countries, special risks of injury may be justifiable for nonaggregative reasons *despite* the manifest inadequacy of alternative forms of employment (i.e., subsistence farming on increasingly crowded plots of land). It may be enough that each sweatshop laborer him- or herself can be expected to do better by making higher and increasing wages despite the workplace health and safety risks. The fact that workers line up for dangerous sweatshop jobs would not *itself* legitimate the risks imposed, though it might give evidence that the jobs are indeed better than any other the workers are aware of—that is, that conditions of *ex ante* benefit hold.⁵⁶

55. Scanlon's transmission towers case might thus be different from the case of Jones the World Cup technician, because Jones is not assumed voluntarily to take on risks of getting stuck in the transmission machinery. Either the technical problems are assumed to be a freak accident and are not simply known to come with the job, or Jones is assumed not to have adequate alternative employment options. SCANLON, *WHAT WE OWE*, *supra* note 1, at 236.

56. This is not to justify any outcomes that result; the continuous reduction of workplace risks through sustained safety standard improvement with international support can still be required, because such measures are reasonably affordable. This is an instance of the point above that the permissibility of using an aggregative decision rule on Harsanyian grounds depends not only on general *ex ante* benefit but also on the assumption that no one can mount a reasonable complaint against the risks he or she is subjected to. For detailed discussion of the case, see JAMES, *FAIRNESS*, *supra* note 51, ch. 10.

D. Institutional Mandate

It is often the case that legitimate public institutions are entrusted with scarce resources for a specified social purpose, and an aggregative decision would advance this aim. When funds have been set aside for the construction and support of public hospitals, for instance, the intended social purpose might be best served by building hospitals in populated areas, where they serve the largest number of people, rather than in areas where a smaller number of perhaps more needy people live. In such cases, there is a specifically moral argument for “doing the most good” in the sense specified by the intended social purpose.⁵⁷ Anything less is a wasted opportunity, a misuse of public resources for their mandated goal.⁵⁸

This is not, however, to *assume* the appropriateness of aggregative choice outside the relevant agency. The legitimacy of the established form of political authority may not itself turn on aggregative considerations and it may or may not “do the most good” by a broader standard. This is not necessarily to endorse bureaucratic compartmentalization. To the extent the relevant institution ignores important claims (as an urban hospital ignores the medical needs of people in less populated areas), the justifiability of its specified mandate and its single-minded advancement will depend on whether those other important claims are met through other institutional arrangements. And when the other arrangements are inadequate, it may be permissible or required for an institution to reach beyond its mandate and take up other causes (e.g., as when public universities branch into early educational outreach given the failures of public secondary schools).

Even so, the question of whether a whole range of institutions adequately meets a relevant range of claims will not *thereby* depend on a comparably general aggregative standard. A whole set of institutions might have a specific personal foundation (e.g., in Rawls’s principles for a society’s basic structure or their application to basic health care). Even if a range of institutions would be better or more coherently governed if governed together by a similar aggregative standard, there is no general presumption in favor of governance that is “better overall” or “more coherent,” at least not without appeal either to impersonal values that are here forbidden or to personal values that have yet to be specified.

E. Moral Neutrality

Finally, I should mention a more general way that contractualism permits aggregation more straightforwardly. While it blocks impersonal consideration

57. One might thus admit the relevance of numbers in Parfit’s “Case Four,” in PARFIT, ON WHAT MATTERS, *supra* note 3, at 197; and “Case Six,” *id.* at 200–201. In the specified institutional contexts, we can fail to give one person a much greater benefit in order to give equal or comparable but much smaller benefits to a much larger number of people (and without abandoning Scanlon’s restriction to personal reasons). Though Parfit intends such cases also to apply outside of our specified institutional setting.

58. SCANLON, WHAT WE OWE, *supra* note 1, at 237 n.42, mentions this basis for aggregation.

of how the world goes over time from no particular person's perspective, it does so *only* when matters of right and wrong, justice and injustice, are in question. In some cases there may be no such question. For instance, suppose that among people with roughly the same expected life spans, we can use scarce resources either to extend one person's life dramatically (e.g., by an additional forty years) or to extend a very large number of people's lives by a significant but much smaller amount of time (five more years).⁵⁹ Is this an issue of right and wrong? Not if people in either group lack relevant grounds for complaint. In that case, either option would in effect be morally permissible. Why might people lack relevant grounds for complaint? It seems clear that the bare fact of death, as such, does not present a valid complaint of "what we owe to each other." The "complaint" that one will eventually die is at most a gripe in register of "cosmic fairness," since death *per se* is, of course, a fact of life beyond all human control.

The same point may apply, however, to life's rough duration, even when it falls to some extent within someone's power. Here we might enter a judgment of relevance: I may lack a reasonable objection to not being given forty additional years, at a significant opportunity cost to many others, not because anyone else enjoys a *stronger* claim to the necessary resources but because, at least in the situation in question, this is not the kind of thing to which I can lay claim in the name of what we owe to each other.⁶⁰ Life has to end at some point, and though I have some rough interest in having that point come later rather than sooner, this is not clearly morally relevant if my death is not in some sense premature.⁶¹ In this kind of case, contractualism can admit that the issue should be settled by aggregative impersonal reasons.⁶²

It is important that this question of whether or when ideas of right and wrong apply is itself a *moral question*. This point is often neglected, for

59. This is "Case 6" in PARFIT, ON WHAT MATTERS, *supra* note 3, at 200.

60. "Threshold deontology" offers one version of this claim. I mean to suggest a more general line of argument, which may or may not assume that claims are always sensitive to some aggregative opportunity cost threshold.

61. The point also applies to Parfit's "Case 7" (*id.* at 203), in which the options are either equally permissible or underdescribed. Parfit's "Case 5" (*id.* at 199) may be different because of implicit reference to *normalcy* of life span in a sense that is broadly subject to societal influence. So, e.g., among people with life-shortening medical conditions, we can use scarce medical resources either to give one person a normal life span (forty more years, instead of an early death at thirty years of age) or marginally to extend a larger number of normal life spans (from seventy to seventy-five total years). Here one might lay a claim of justice to a measure of normalization. But the claim need not be decisive. Aggregative concerns of Institutional Mandate may still apply.

62. This may require rejecting what Parfit calls Scanlon's "Greater Burden Claim" (*id.* at 192). Given that claim, reasonable rejection of a principle depends in part on its "burden" as compared to an alternative principle's being in force, where this *always includes any benefits one would forgo*. The suggestion in the text is that we instead treat forgone benefits as relevant only in appropriate contexts, as supported by context-specific generic reasons claims. Relaxing the Greater Burden Claim as a sweeping constraint relieves some of the pressure Parfit places upon Scanlon's restriction to personal reasons. For this and other grounds for aggregation, Parfit may therefore be mistaken to suggest (*id.* at 200) that "It is clearly Scanlon's Individual Restriction which is making Scanlon's Formula go astray." On the contrary, the Individual Restriction is not clearly the culprit.

instance, by the law and economics defense of “cost-justified” precaution in accident law, which urges that the alternatives to a no-fault scheme (such as New Zealand–style social insurance) leave everyone less well off than they could be.⁶³ Such a scheme might well be morally permissible for aggregative cost-benefit reasons *if* there is no issue of justice or fairness at stake in the choice between no-fault and fault-based systems. But the case for that would have to be made, and the *mere* fact that some or all might be better off than otherwise would not suffice for that moral argument. Accident law may not be the appropriate place to economize (again, as U.S. law assumes). Even when there is a distinct *fairness* argument for a no-fault system,⁶⁴ this is not to say that fairness will not come with a price in economic gains forgone. The fact that some or all would be “worse off” in the sense of suffering a mere opportunity cost might make no moral difference *per se*.

V. THE CONTRACTUALIST’S CONJECTURE

Lest aggregation seem *too easy* to justify, we might summarize the *ex ante* contractualist limits on justification as the following disjunctive moral constraint—a checklist, if you will, for permissible aggregative decision-making. According to what we might call the *contractualist’s conjecture*:

An aggregative decision is morally permissible in public affairs if at least one of the following conditions is met and impermissible if none (among those which apply) are satisfied.⁶⁵

Moral Neutrality: the policy alternatives at issue are all morally permissible, in which case an aggregative decision is not an issue of right or wrong, justice or injustice.

Tiebreaking: the decision favors a larger over a smaller number of people faced with similar or otherwise comparable prospects of benefit or harm.

Institutional Mandate: a legitimate public institution is entrusted with scarce resources for a specified social purpose, and an aggregative decision would advance this aim.

General Expected Benefit: goods are created among those who have similar chances of overall personal benefit, and no one potentially affected can reasonably object to the imposition of risk.

63. LOUIS KAPLOW & STEVEN SHAVELL, *FAIRNESS VERSUS WELFARE* (2002), emphasize this point. From a contractualist point of view, they mistakenly assume that considerations of greater benefit are decisive without offering necessary moral arguments.

64. Jeremy Waldron, *Moments of Carelessness and Massive Loss*, in *PHILOSOPHICAL FOUNDATIONS OF TORT LAW* (David G. Owen ed., 1995), offers a Rawlsian defense of a no-fault scheme.

65. I assume there are alternatives, such as following established administrative practice (as distinct from relying on cost-benefit constructions) with sensitivity to political pressure or special interests. Officials could also simply muddle along, deciding by moral intuition or assumed principle. There is also, of course, a range of cost-benefit approaches for any particular case, only some of which may be justifiable.

Due Precaution: the decision in question seeks social benefit but takes precautions that reduce risks of harm at a reasonable cost to those potentially adversely affected.

Risk Assumption: the risks created by the decision are voluntarily assumed by those exposed to them, who have been given an adequate range of alternatives.

Other Personal Grounds: the decision can be justified by other "personal" grounds without direct appeal to impersonal values.⁶⁶

This framework of reasoning must still, of course, be applied in real policy contexts. The contractualist's *conjecture*, at any rate, is that this set of ideas yields a plausible understanding of both the limits of aggregation and its legitimating rationales: given further case-specific analyses, it allows us to walk the fine line between aggregation of the right and the wrong kinds.

Much then depends on how the general framework of reasoning applies, especially when we move away from stylized cases to actual policy contexts. We have not answered 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. Our point, again, is more general: friendliness to aggregative decision-making is not necessarily a problem provided it has the appropriate grounds, and indeed, this considerably aids a contractualist project that has been dogged by concerns of blocking aggregation of rightful kinds. Many contexts may well turn out to present stably the appropriate conditions for aggregative choice, even as a kind of policy default. Yet the default will have no further generality beyond those particular settings and in any case will always stand open to scrutiny from a fundamentally nonaggregative moral perspective. Where rival views converge in practical recommendation, the *ex ante* contractualist can claim the better account of why aggregation should be justified. Seen as a characterization of morality in public policy choice, an *ex ante* contractualism can bid fair to being meaningfully distinct from and perhaps even superior to utilitarianism and other welfarist consequentialist views simply because of the general way aggregative decision is constrained.

VI. HOW SHOULD WE ASSESS STAKES?

We conclude by addressing a methodological objection. *Ex ante* contractualism as described so far affords no general account of how we are to assess

66. SCANLON, WHAT WE OWE, *supra* note 1, at 221, allows that personal claims, e.g., against environmental degradation might be bolstered by the impersonal value of environmental preservation. Thus the fact that the overall cost of aggressive climate change mitigation is relatively small (according to BJÖRN LOMBORG, THE SKEPTICAL ENVIRONMENTALIST (2001), at 323, it amounts to one in fifty years of economic growth) might strengthen personal objections to anything less than aggressive measures. Personal objection might also arise from the personal interest we each have in the nature of the society in which we live, including, e.g., a personal interest in societal goods such as national security or national wealth. For the case of the national "gains of international trade," see JAMES, FAIRNESS, *supra* note 51, ch. 6.

the stakes of different parties involved in a given decision. It may therefore seem objectionably arbitrary or ad hoc. Is not some general rule of judgment for assessing stakes required? Otherwise, even if there is no slope into utilitarianism, are we not left grasping at slippery intuitions more or less in the dark?

Scanlon himself worries that an *ex ante* approach would generally have to discount harms according to how likely they are to occur, and that any general rule for doing so would be open to ready counterexample. As he explains, if we are to discount harms by their improbability, it seems a coercive medical experimentation lottery could be justifiable as long as the pool is large and chances of harm to any one person are very small. This is not a counterexample, however, if *ex ante* contractualism assumes no general discounting rule for assessing stakes and instead leaves the matter to context-specific judgment.⁶⁷ We could then explain the case as follows. The imagined lottery might be fine if people freely enroll to aid the advancement of medicine, just as many astronauts might jump at the chance of a one-way trip into the further reaches of outer space. To the extent the risks of medical experimentation are not voluntarily assumed, however, it would not matter how low the probability of being selected for forced experimentation is.

Scanlon reaches that conclusion by banning probabilities. We can reach the same result by *disqualifying any specific probability assignment as substantively irrelevant while assuming a credible possibility of harm*, at least in that kind of situation. The grounds for this might be our especially powerful reasons to control our bodies, which Scanlon explicates, in a similar connection, in this way: “it is not unreasonable to refuse to regard one’s own life and body as ‘on call,’ to be sacrificed when it is needed to save others who are at risk.”⁶⁸ For that reason, the *specific* chances of harm are not to be weighed up in comparison to expected benefits but are rather “silenced” from a moral point of view.⁶⁹ The lottery is, then, no more permissible than a case in which a government holds an execution lottery for the sake of greater security, even if each subject has only a tiny risk of secret abduction

67. So, e.g., in the unexploded mine case discussed in Reibetanz, *supra* note 14, at 302, we can treat the case as one of easy rescue: one knows that each worker (in a group of one hundred) has a substantial chance of death, while one is merely sure to become ill by a rescue effort. Because the claims of each worker to assistance are not discounted (according to the specific probability that he or she among the one hundred will be the one injured), the substantial chance of each worker’s death thus defeats one’s appeal to one’s own high risk of mild illness. This suggestion is similar to the objection mentioned in *id.* at 303 n.12, and I believe it is not undermined by the suggested reply, which appeals to an *ex post* reading of contractualism.

68. SCANLON, WHAT WE OWE, *supra* note 1, at 209.

69. One could still imagine a further case in which only a forced medical experimentation lottery would reliably spur medical advances that would save millions from death. The issue here is still not one’s probability of harm but how much one’s interest in bodily self-determination weighs against the possibility of saving lives—that is, whether or to what extent that interest has a silencing role. Claiming that it does have a silencing role would not require relaxing Scanlon’s restriction to personal reasons, as PARFIT, ON WHAT MATTERS, *supra* note 3, at 209–210, recommends in light of a similar example.

and termination and even if, by chance, no one is ever in fact killed.⁷⁰ Objections to arbitrary treatment might render the unlikelihood of harm irrelevant, much as would claims to bodily control.

Of course, judgments of moral relevance will usually only shape the field of considerations to be evaluated; most policy contexts will still require substantive assessments both of what the different relevant risks and prospects are for different parties and how they compare. The task for *ex ante* contractualism is to offer that kind of substantive analysis for a range of realistic policy contexts.

One may object that *ex ante* contractualism is not then a *theory*, properly speaking. Yet it is of a piece with the open-ended character of contractualism generally. Taking the theory on its own terms, the proper test of adequacy is not whether it could have determinacy of a kind it does not strive for but rather whether its general structure helps us reason through various specific cases in perspicuous ways.⁷¹

To elaborate, Scanlon's theory is notoriously open about what counts a comparatively "reasonable" complaint, that is, about which objections are sufficient to defeat which other objections. No less important, however, is that the theory is open about what counts as a *relevant* objection in the first place. We are called upon to make a *relevance judgment*, in the case at hand, about any number of possible factors, including the relevance or weight of preferences or interests, of entitlements, rights, and statuses (all defeasible), and of objections to arbitrary treatment or to the agent's attitudes (his or her intentions, motives, etc.).⁷² The idea is not simply that some facts of the case are of no normative relevance at all, as in Parfit's famous case of future Tuesday indifference (which shows that it is not relevant to the badness of someone's pain that it occurs on a Tuesday). For even the fact that people indeed have certain reasons for action can be disqualified as morally irrelevant in the context in question (e.g., my real interest in spending a day at the beach is irrelevant when I have promised to be elsewhere).

Moreover, all "inputs" into moral reasoning are codependent: what counts as relevant is to be decided in part based on what sorts of considerations would add up to a sufficient, "reasonable" complaint. We do not first settle the relevant inputs and then settle what objections are comparatively sufficient. The underlying parameters are open to adjustment according to what

70. The lottery could run in two stages: a first stage that decides whether someone will be executed, and a second stage that decides who. The first stage is heavily weighted against anyone being chosen, and by chance has never yet moved to the second selection stage.

71. But will contractualism *guide* decision-making if it is not made less judgment-dependent somehow? It arguably still can. First, policy-making is invariably shaped by value-judgment. Second, contractualism can in any case offer principles that frame and in that sense guide policy choice. Third, the possibility of an aggregative cost-benefit calculus that provides "determinacy" without judgment is largely illusory: unless any decision rule can be appropriately morally grounded, in part by way of moral judgment, it is no more justified than a determinate toss of a coin.

72. Although Scanlon has come to reject the relevance of intention to permissibility in SCANLON, MORAL DIMENSIONS, *supra* note 9, ch. 1.

yields the best characterization of the moral situation overall. We therefore are offered not a way of deriving moral common sense from something general and abstract but at most a holistic way of shaping and articulating our commonsense intuitions and judgments in a way that enhances our moral understanding while leaving room for critical engagement. Since the general framework of contractualist reasoning does at least shape what specific sorts of judgments are required, it cannot be fairly said to leave us *wholly* in the dark. Judgment is required but also guided.

There is no obvious general problem here for *ex ante* contractualism as opposed to its *ex post* or hybrid version: all such contractualisms will require independent but guided judgment about what the interests and outcomes are, whether actual or expected, in a given setting—unless, of course, we do not feel we can be very confident about how risks are to be identified and assessed, even with what help the larger framework of reasoning provides, whereas we *can* be relatively confident about the goodness or badness of outcomes for peoples lives. Ultimately, the success of the *ex ante* contractualist project depends on showing that we do not have to feel that way about the imposition of risk. I say above that there is no slippery slope here into utilitarianism. Yet the project can hardly rest easy: the task remains to show in credible detail that we can indeed get a grip on how and to what extent we can justify the risks that actually pervade our social lives.