JOHANN FRICK

Contractualism and Social Risk

I. INTRODUCTION

Contemporary nonconsequentialism is a family of views united less by a positive doctrine than by skepticism toward central tenets of consequentialist ethical thought. One such tenet, which is embraced by most consequentialists but opposed by many nonconsequentialists, is the notion of interpersonal aggregation.[1] Ethical theories, like classical utilitarianism, that defend interpersonal aggregation hold that in evaluating an action, we should sum the benefits and losses it imposes on different people to obtain an aggregate quantity; this represents the overall goodness of the action's consequences. The rightness or wrongness of the action depends not on how it affects each individual, but on the net balance of benefits over losses.

Aggregative reasoning of this kind often yields counterintuitive implications, especially in cases where it enjoins us to let a few people suffer

In writing this article I have accumulated many debts. For their very helpful comments on earlier drafts, I thank Arthur Applbaum, Ralf Bader, Eric Beerbohm, Selim Berker, Dan Brock, Norman Daniels, Tom Dougherty, Nir Eyal, Marc Fleurbaey, Mira Frick, Daniel Halliday, Louis-Philippe Hodgson, Aaron James, Stephen John, Frances Kamm, Gregory Keating, Paul Kelleher, Seth Lazar, Marc Lipsitch, Jeff McMahan, Andreas Mogensen, Ekédi Mpondo-Dika, David O'Brien, Laura Ong, Derek Parfit, David Plunkett, Tim Scanlon, Lucas Stanczyk, Victor Tadros, Gerard Vong, Dov Waisman, Alec Walen, and Daniel Wikler, as well as two anonymous referees for *Philosophy & Public Affairs*. I would also like to thank audiences and participants at the Moral and Political Philosophy Workshop at Harvard University, the New York Early Career Ethics Workshop, the Lawrence S. Rockefeller Seminar at Princeton, the Moral Philosophy Seminar at the University of Oxford, and the Workshop on the Ethics of Social Risk at the Centre de recherche en éthique, Montreal.

1. For a dissenting voice in the nonconsequentialist camp, see S. Matthew Liao, "Who Is Afraid of Numbers?" *Utilitas* 20 (2008): 447–61.

large losses so that many people can enjoy smaller benefits. A well-known illustration is T. M. Scanlon's

> *Transmitter Room Case*: Jones has suffered an accident in a TV broadcasting station and is receiving extremely painful electrical shocks. If we turn off the power to save him, billions of viewers will miss the last half hour of the World Cup final.[2]

Many of us share Scanlon's intuition that it would be wrong not to save Jones from his agony, *regardless* of how many people are watching the game. The benefit of watching a soccer match is trivial compared to the agony of suffering strong electrical shocks. No matter how large the *sum* of these benefits, it would be wrong to keep the power on. Defenders of interpersonal aggregation, by contrast, seem committed to the view that at some point the combined benefits to the viewers must become large enough to morally outweigh Jones's agony.

A common nonconsequentialist diagnosis for the inadequacy of aggregative views is that they do not respect what John Rawls has called "the separateness of persons." Collections of people are not super-individuals. Unlike a single individual, who may rationally choose to make some sacrifice in order to receive a stream of benefits, a group of people lacks the requisite unity such that imposing a significant harm on one person could straightforwardly be offset by conferring sufficient benefits on others.

Seeking to formulate an account that better respects the separateness of persons, leading nonconsequentialists have put forward so-called competing claims models of moral rightness, according to which morality requires us to identify, by a series of pairwise comparisons, the action or policy that satisfies the strongest *individual* claim or—its flipside—generates the weakest *individual* complaint.[3]

The fullest development of this antiaggregative approach is the contractualist moral theory defended by T. M. Scanlon. According to Scanlon's contractualist formula, an action is morally right if and only if

2. Adapted from T. M. Scanlon, *What We Owe to Each Other* (Cambridge, Mass.: Harvard University Press, 1998), p. 235.

3. See, for instance, T. Nagel, "Equality," in *Mortal Questions* (Cambridge: Cambridge University Press, 1979), pp. 106–27; and T. Nagel, *Equality and Partiality* (Oxford: Oxford University Press, 1991); see also T. M. Scanlon, *What We Owe to Each Other*, chap. 5.

it is justifiable to each person, that is, if a principle licensing the action could not be reasonably rejected by any *single* individual for personal reasons. Personal reasons, for Scanlon, are "reasons that are . . . tied to the well-being, claims, or status of individuals in [a] particular position."[4] They exclude appeals to how an action will affect *other* individuals, as well as to *impersonal* concerns, such as the goodness of the overall outcome that the action will produce. Scanlon's contractualist formula thus rules out interpersonal aggregation in two distinct senses: On the one hand, it excludes the "axiological" aggregation characteristic of classical utilitarianism. That is, it rejects the thought that the rightness of an action is a function of the *overall* goodness of its consequences, which in turn is determined by summing its good and bad effects for different individuals. On the other hand, the contractualist formula also rules out aggregation of a "normative" kind, in that it looks only at the personal reasons of individuals and does not allow for the pooling of moral claims or complaints across different individuals.

In addition, Scanlon affirms the following:

> *Greater Burden Principle*: "It would be unreasonable . . . to reject a principle because it imposed a burden on you when every alternative principle would impose much greater burdens on others."[5]

Scanlonian contractualism yields the intuitively correct answer in Transmitter Room. Turning off the power is the only course of action that is justifiable to everyone, since Jones's personal reasons for rejecting a principle that allowed us to continue the transmission are much weightier than any individual's reasons against interrupting the broadcast. The fact that there are a large number of people with personal reasons for not wanting the broadcast to be interrupted is deemed

4. Scanlon, *What We Owe to Each Other*, p. 219. For Scanlon, the personal reasons invoked in reasonably rejecting a principle must be what he calls *generic* reasons. Generic reasons are based not on the idiosyncratic interests and features of actual individuals, but on what persons occupying this particular position would have reason to want "in virtue of their situation, characterized in general terms." Scanlon, *What We Owe to Each Other*, p. 204.

5. T. M. Scanlon, "Contractualism and Utilitarianism," in *Utilitarianism and Beyond*, ed. Amartya Sen and Bernard Williams (Cambridge: Cambridge University Press, 1982), p. 111.

morally irrelevant, since contractualism rejects the aggregation of individual claims.[6]

Cases like Transmitter Room, however, are quite rare. As Barbara Fried remarks, outside the context of crime and war, most instances of serious harm do not occur in situations where someone trades off *certain* harms to *known* individuals against certain benefits to other known individuals.[7] (Let us call such situations *interpersonal trade-offs under certainty*.) Rather, most harm to others results from *risks* that some agent imposed on them, or failed to eliminate, in pursuit of some benefit for them, for the agent, or for others. The question I wish to raise in this article is whether contractualists can provide a plausible account of such risky actions and omissions, while honoring their antiaggregative commitments.

Most of my article will focus on one class of risky actions and omissions that, prima facie, are particularly troubling to the contractualist. These actions are characterized by the following four features:

(1) The risky action or omission will affect a large number of individuals. Because of this, it is virtually certain that *some* people will end up being burdened by it.
(2) The individual losses to those who are burdened (relative to the baseline of some available alternative) are considerably greater than the individual gains for those who are benefited.
(3) The action-type in question is rare, or rarely affects the same people twice; as a result, we cannot assume that over time almost

6. Scanlon does allow that the numbers may count when we must choose between preventing a smaller or a larger number of people from suffering an *equivalent* burden (for example, saving one from death vs. saving two from death). But even here his reasoning remains nonaggregative. In a choice between saving A or saving both B and C, it would be wrong to flip a coin—but not because two deaths represent a greater *total* loss, or because B and C can pool their claims to form a stronger collective claim. Rather, following Frances Kamm, Scanlon argues that a principle which did not allow numbers to "break the tie" in this case could be reasonably rejected by both B and C for failing to take account of the value of saving their life, "since it permits the agent to decide what to do in the very same way that it would have permitted had he not been present at all, and there was only one person in each group." Scanlon, *What We Owe to Each Other*, p. 232, following F. Kamm, *Morality, Mortality*, vol. 2 (Oxford: Oxford University Press, 1993), pp. 116–17.

7. B. Fried, "Can Contractualism Save Us from Aggregation?" *Journal of Ethics* 16 (2012): 39–66.

everyone will benefit from a principle that permits actions of this type to be performed.[8]

(4) The risky action or omission is intuitively permissible.

Many actions in public policy, particularly in the domains of public health and of risk management, share these four features. (I will provide a range of stylized examples over the course of this article.) The challenge for the contractualist is to explain, without resorting to interpersonal aggregation, how such risky actions could be permissible, given features (1), (2), and (3). For want of an established label, I shall dub this the problem of *social risk*.

As we shall see, cases of social risk force the contractualist to confront a question about his theory that, I argue, has not received a fully satisfactory answer to date: Are an individual's personal reasons for rejecting a principle permitting the imposition of a given social risk a function of the *prospect* that the risky action gives to each person ex ante, allowing us to discount both benefits and harms by their improbability of occurring? Or is justifiability to each person a function of the action's outcomes ex post? In that case, what would matter morally is not the ex ante likelihood of any *given* individual's being benefited or harmed by the action, but rather the near certainty that, ex post, *some* persons will turn out to have been harmed by the action while others will have been benefited.

The contractualist faces an apparent dilemma: if he rejects the ex ante view of justification, as Scanlon explicitly does in *What We Owe to Each Other*, and appeals instead to an ex post view, as other authors have urged,[9] he seems committed to the intuitively unappealing conclusion that most instances of social risk, as defined above, are unjustifiable. On

8. This feature blocks an appeal to what Scanlon calls "*intra*-personal aggregation," that is, "aggregation *within* each person's life, summing up all the ways in which a principle [prohibiting that kind of risky action] would constrain that life, rather than aggregation *across* lives, adding up the costs or benefits to different individuals." Scanlon, *What We Owe to Each Other*, p. 237. Intrapersonal aggregation may allow the contractualist to justify the performance of *routine* risky actions. While, on any given occasion, such a risky action may impose significant individual losses on some while providing only relatively minor benefits to others, over time almost *everyone* will come out ahead under a principle that permits risky actions of this type to be performed. The kinds of risky actions that I am concerned with in this article are not routine in this sense.

9. See, in particular, Sophia Reibetanz Moreau, "Contractualism and Aggregation," *Ethics* 108 (1998): 296–311. Elizabeth Ashford, "The Demandingness of Scanlon's Contractualism," *Ethics* 113 (2003): 273–302, also assumes an ex post view of justification.

the other hand, if he embraces an ex ante view, Scanlon fears that this will move his contractualist theory too close to those aggregative views that he wished to escape from in the first place.[10]

In this article, I argue that this dilemma can be overcome. I defend a version of ex ante contractualism that gives a satisfactory response to the problem of social risk, while also avoiding the excesses of aggregative moral theories. My article proceeds as follows: To get a better grip on the problem of social risk, Section 2 begins by introducing a stylized case of social risk, and points out similarities with the more familiar problem of interpersonal trade-offs under certainty. Section 3 presents a simple yet seductive argument—the "Argument from Irrelevant Information"—that supports an ex post view of contractualist justification. According to this argument, the problem of social risk can be morally assimilated to that of interpersonal trade-offs under certainty. Section 4 challenges this view, arguing instead for an ex ante contractualism according to which the problem of social risk is more akin to one of *intra*personal risk taking. Section 5 returns to the Argument from Irrelevant Information, distinguishing three ways of fleshing it out. I argue that ultimately we must reject the Argument from Irrelevant Information on any of these interpretations, and with it the ex post contractualism that it supports. In Section 6, I critically examine T. M. Scanlon's reasons for eschewing ex ante contractualism, by considering what I call the problem of *ex ante rules*. An ex ante rule is any rule the adoption of which, at some time t_1, is in everyone's individual interest, but which licenses or requires some agent to act at a later time t_2 in a way that benefits some but significantly burdens others. As we shall see, the worry that ex ante contractualism may support morally indefensible ex ante rules is at the heart of Scanlon's objections to this view. Moreover, I show that existing attempts by ex ante contractualists, such as Rahul Kumar, to overcome

10. This, at least, was Scanlon's view in *What We Owe to Each Other*. In a more recent article, "Reply to Zofia Stemplowska," *Journal of Moral Philosophy* 10 (2013): 508–14, Scanlon writes that his earlier opposition to ex ante contractualism was "a mistake," and credits an earlier draft of my article for changing his mind.

That ex ante contractualism cannot "save us" from aggregation is also the conclusion of Barbara Fried, "Can Contractualism Save Us from Aggregation?" Two more optimistic assessments of the prospects for ex ante contractualism, which can be read as complementing the case I make in this article, are given by Aaron James in "Contractualism's (Not So) Slippery Slope," *Legal Theory* 18 (2012): 263–92; and Rahul Kumar, "Risking and Wronging," *Philosophy & Public Affairs* 43 (2015): 27–49.

this difficulty are not fully satisfactory.[11] However, by adopting what I call the *Decomposition Test*, according to which an action is justifiable if and only if the actions it licenses us to perform are justifiable to each person *at each temporal stage*, I believe that Scanlon's misgivings can be overcome.

Unfortunately, ex ante contractualism's success at dealing with the problem of social risk comes at a theoretical price, which hitherto has received scant attention from defenders of this view. In Section 7, I draw out some of the implications that ex ante contractualism has for a related problem: the question of "identified" vs. "statistical" lives. I argue that, contrary to received wisdom, ex ante contractualism can lend a measure of normative support to our common psychological propensity to give greater weight to the saving of "identified" over "statistical" lives. The problem, as I show in Section 8, is that, under certain circumstances, ex ante contractualism would lead us to privilege the saving of one identified life over *any* number of statistical lives.[12] The most promising way of avoiding this extreme conclusion, I argue in Section 9, is by scaling back the ambitions of contractualism as a moral theory. Instead of providing a complete account of what it *is* for actions to be right or wrong, the contractualist competing claims model is better understood as capturing an important *class* of pro tanto moral reasons that contribute to *making* actions right or wrong, but that do not by themselves determine an action's rightness *all things considered*.

II. SOCIAL RISK AND INTERPERSONAL TRADE-OFFS UNDER CERTAINTY

Consider the following stylized case of social risk:

> *Mass Vaccination (Unknown Victims)*: One million young children are threatened by a terrible virus, which is certain to kill all of them if we do nothing. We must choose between mass producing one of two vaccines (capacity constraints prevent us from producing both):
>
> - Vaccine 1 is certain to save every child's life. However, the vaccine will not provide complete protection against the virus. If a child

11. Kumar, "Risking and Wronging."

12. James, "Contractualism's (Not So) Slippery Slope," and Kumar, "Risking and Wronging," the two main proponents of ex ante contractualism in the literature, both fail to discuss this problem.

receives Vaccine 1, the virus is certain to paralyze one of the child's legs, so that he or she will walk on crutches for the rest of his or her life.

- Vaccine 2 is risky. It gives every child a 999/1000 chance of surviving the virus completely unharmed. However, for every child there is a 1/1000 chance that Vaccine 2 will be completely ineffective and that the child will be killed by the virus. (Assume that the outcomes for different children are probabilistically independent.) Call the children who end up dying the *luckless* children.

Whichever of the two vaccines we choose to produce will be administered to all one million children.[13]

Before I go on, a brief note about the notions of "chance" and "probability," as I employ them in this and subsequent cases. Notwithstanding the posits of contemporary quantum physics, I shall assume that natural processes at the *macro*-physical level that we deal with in medicine and most other domains of risky human activity are, for all intents and purposes, deterministic. When using the terms "probability" or "chance" in this and all following examples, I shall therefore assume that we are speaking not about *objective* indeterminacy at the level of physical reality itself, but about *epistemic* probability, which reflects our incomplete knowledge of the state of the world and the laws of nature. The notion of epistemic probability that I will be employing is *evidence*-relative: To say that, from the perspective of an agent, some event *e* has an epistemic probability *p* of occurring is to say that, given the evidence available to the agent, her *rational* degree of credence in the proposition "*e* will occur" would be *p*. (To keep things simple, I shall assume that the decision makers' *actual* degrees of credence track what the available evidence makes it rational for them to believe.) In Section 5.C, I revisit this assumption and argue that, for the purposes of morally evaluating risky actions, it typically *makes no difference* whether the risk in question corresponded to objective or to "merely" epistemic chances.

13. I specify that the patients are young children in order to bracket the complicating factor of patient autonomy; an adult patient would presumably have the right to refuse either treatment.

How ought a contractualist to go about deciding how to act in Mass Vaccination (Unknown Victims)? One way of thinking about this case is by analogy with the following interpersonal trade-off under certainty:

> *Mass Vaccination (Known Victims)*: The threat situation is as in Mass Vaccination (Unknown Victims) above. This time, we must choose between mass producing one of the following two vaccines:
>
> - Vaccine 1, as above, is sure to save every child's life, at the cost of paralyzing one of their legs.
> - Vaccine 3 is sure to allow 999,000 children to survive the virus completely unharmed. However, because of a known particularity in their genotype, Vaccine 3 is certain to be completely ineffective for 1,000 identified children. These *doomed* children are sure to be killed by the virus if we choose Vaccine 3.

From what was said in the previous section, it is clear how a Scanlonian contractualist ought to think about Mass Vaccination (Known Victims). The individual burden of becoming paralyzed in one leg, though significant, is not even close to that of losing one's life at a young age. Since contractualism does not allow the aggregation of claims, the Greater Burden Principle implies that no one could reasonably reject a principle requiring us to produce Vaccine 1, whereas a principle licensing us to choose Vaccine 3 could be reasonably rejected by each of the 1,000 doomed children.

Now return to Mass Vaccination (Unknown Victims). Note that, due to the Law of Large Numbers, we can predict with a high degree of confidence that the overall pattern of outcomes from administering Vaccine 2 would look a lot like that of administering Vaccine 3 in my second case. Although no *particular* child is certain to die, it is a statistical certainty that *some* children will be killed by the virus if we administer Vaccine 2, and the likelihood that roughly 1,000 (±100) will die is greater than 99 percent.[14]

14. The probability that at least *one* child will be killed by Vaccine 2 is $1-\left(\frac{999}{1{,}000}\right)^{1{,}000{,}000}$, which is very close to 1. By Chebyshev's Inequality, the likelihood that roughly 1,000 children (±100) will be killed is > 0.99. Given that contractualism prohibits aggregation, only the former number really matters to the argument.

Moreover, the harm of death is just as great for a luckless child as for a doomed child. Holding constant the fact of actually suffering a harm, it is typically no better for a person that she had some chance of escaping the harm rather than being certain to suffer it. (I return to this point in Section 5.)

III. THE ARGUMENT FROM IRRELEVANT INFORMATION AND EX POST CONTRACTUALISM

Given these two parallels—a near-identical overall pattern of outcomes and equivalent individual harms—a number of nonconsequentialist philosophers have maintained that cases of social risk can, under certain conditions, be morally *assimilated* to interpersonal trade-offs under certainty. In the following, I will concentrate on an argument by Sophia Reibetanz Moreau.[15] Elsewhere,[16] I have contended with similar arguments by Marc Fleurbaey and Alex Voorhoeve.[17]

All of these philosophers deploy versions of the following basic argument, which, applied to my two examples, runs as follows:

> *The Argument from Irrelevant Information*: In Mass Vaccination (Known Victims), selecting Vaccine 3 is unjustifiable to the doomed children, since this would impose on them a greater individual burden than Vaccine 1 imposes on anyone. But in Mass Vaccination (Unknown Victims), we know that, if we choose Vaccine 2, some children will end up in the *same* position that the doomed children occupied in Mass Vaccination (Known Victims). They, too, will lose their lives, whereas under Vaccine 1 the greatest harm that any child would have had to suffer is a paralyzed leg. To be sure, we do not yet know who these luckless children are. But surely, the identity of the eventual victims is *irrelevant information*. Whoever will lose their life through Vaccine 2 will have as strong a complaint against the decision maker

15. Reibetanz Moreau, "Contractualism and Aggregation."

16. Johann Frick, "Uncertainty and Justifiability to Each Person: Response to Fleurbaey and Voorhoeve," in *Inequalities in Health: Concepts, Measures, and Ethics*, ed. Nir Eyal, Samia Hurst, Ole Norheim, and Daniel Wikler (New York: Oxford University Press, 2013).

17. Alex Voorhoeve and Marc Fleurbaey, "Decide as You Would with Full Information! An Argument against *Ex Ante* Pareto," in Eyal, Hurst, Norheim, and Wikler, *Inequalities in Health*.

> as a doomed child has against Vaccine 3. Therefore, Vaccine 1 is the only option that is justifiable to each person in *either* case.

For proponents of the Argument from Irrelevant Information, the strength of an individual's personal reasons for rejecting a risky action depends not on how likely *they themselves* were to suffer a burden, but on how likely it was that *someone* would. Thus, Sophia Moreau writes:

> As long as we know that acceptance of a principle will affect *someone* in a certain way, we should assign that person a complaint that is based upon the full magnitude of the harm or benefit, even if we cannot identify the person in advance. It is only if we do not know whether acceptance of a principle will affect anyone in a certain way that we should allocate each individual a complaint based upon his expected harms and benefits under that principle.[18]

Moreau subscribes to what is called an *ex post view of justification*, according to which the relative strength of an individual's harm-based complaint turns not on the ex ante prospect it offered her individually, but on the foreseeable distribution of outcomes across individuals that the action will produce ex post. Paired with a ban on the aggregation of claims, this ex post view strongly privileges the perspective of the individual who turns out to be most burdened under a given principle. It is to her that justification must be addressed, typically by showing that any alternative principle would have ended up imposing an even greater burden on her or on someone else.

Although the Argument from Irrelevant Information is seductive, I believe that ex post contractualism is hard to accept. Note, in particular, that it would dramatically contradict many of our ordinary moral convictions: in real life, we often impose social risks that closely resemble that of choosing Vaccine 2 in Mass Vaccination (Unknown Victims). Thus, it is commonly deemed morally unproblematic to systematically inoculate young children against certain serious but nonfatal childhood diseases where there is a remote chance of fatal side effects from the inoculation itself. But this is not because it is unlikely that inoculation will ever lead to disaster for some unlucky children. Given the large

18. Reibetanz Moreau, "Contractualism and Aggregation," p. 304.

number of children inoculated each year, it is a statistical certainty that some small number of them will develop fatal complications. If ex post contractualism were correct, a policy of systematic inoculation would be very hard to defend. Indeed, it would be as hard to defend as a policy that sacrificed the lives of some small *known* group of children each year in order to save the others from contracting nonfatal childhood diseases.

IV. THE ARGUMENT FROM THE SINGLE-PERSON CASE AND EX ANTE CONTRACTUALISM

Given these problems with ex post contractualism, it is worth exploring a different approach. A natural thought is that, instead of assimilating Mass Vaccination (Unknown Victims) to an interpersonal trade-off under certainty, we could compare it to a *single*-person decision under risk. Consider:

> *Individual Vaccination*: A single young child (call her Clara) is threatened by a terrible virus, which is certain to kill her if we do nothing. As Clara's guardians, we must decide which of two available vaccines to give her. As before,
>
> - Vaccine 1 is certain to save Clara's life, but at the cost of one of her legs becoming paralyzed by the virus.
> - Vaccine 2 will give Clara a 999/1000 chance of surviving the virus completely unharmed; however, there is a 1/1000 chance that Vaccine 2 will be completely ineffective for Clara and that she will be killed by the virus.

How ought we to act? In von Neumann-Morgenstern's expected utility theory, the effect of a utility loss on a person's overall expected utility is weighted by the probability of its occurrence; this means that large but unlikely losses can be equivalent in expected utility to smaller but more likely losses (likewise for gains). Hence, according to expected utility theory, the prospect offered to Clara by Vaccine 1 is worse than that offered by Vaccine 2 if and only if her utility from becoming paralyzed and walking on crutches for the rest of her life is less than the weighted sum of her utility loss from death (weighted by a factor of 1/1000) plus the benefit of a life in full health (weighted by a factor of 999/1000). Many of us would

agree that this is the case, and hence that it would be morally right for Clara's guardians to choose Vaccine 2 on her behalf.

This result is not an artifact of embracing expected utility theory. Some versions of *prioritarianism*, for instance, maintain that as Clara's guardians we ought not necessarily select the treatment that gives her the greatest expected utility ex ante. Rather, we should assign disproportionate normative weight to avoiding an outcome that is very bad for Clara (in this case, her death).[19] Thus, even if Clara's expected utility from receiving Vaccine 2 were *somewhat* greater than her expected utility from Vaccine 1, a prioritarian might still opt to give her Vaccine 1, because this is sure to avert the outcome that is worst for Clara. In the present case, however, Clara's expected utility from receiving Vaccine 2 is arguably *considerably* greater than her expected utility from Vaccine 1. Given this, I claim, a reasonable prioritarian would concur that it would be justifiable for Clara's guardians to choose Vaccine 2 on her behalf. For a young child, escaping certain paralysis in one leg for the rest of her life is worth a 1/1000 risk of death.

Our response in Individual Vaccination points the way to a different way of thinking about the imposition of social risks. I call this the *Argument from the Single-Person Case*:

> *Argument from the Single-Person Case*: For contractualists, the rightness of an action is a function of each individual's personal reasons for rejecting a principle that licenses the action. An act is wrong if and only if there is someone who can complain that we failed to treat her in a way that was justifiable to *her*, not because its consequences were *impersonally* bad. In assessing the rightness of a risky action, we must focus on how it affects each individual, compared to the possible alternatives. To better see how Mass Vaccination (Unknown Victims) affects each person individually, this case might be analytically decomposed into one million single-person gambles that occur in parallel. But what do we see if we do this? The prospect offered to each child by

19. I set aside the question whether this version of prioritarianism, which applies the view not just to *inter*personal trade-offs but to *intra*personal gambles as well, is the most plausible way of fleshing out the view. For discussion of this point, see Michael Otsuka and Alex Voorhoeve, "Why It Matters That Some Are Worse Off Than Others: An Argument against the Priority View," *Philosophy & Public Affairs* 37 (2009): 171–99; and Derek Parfit, "Another Defense of the Priority View," *Utilitas* 24 (2012): 399–440.

> administering Vaccine 2 is exactly the same as the one given to Clara in Individual Vaccination, and the alternative option is the same as well. Hence, by mass producing and administering Vaccine 2, we do for each child what a guardian, concerned solely with *that* child's interests, ought to have done for her in a single-person case. Given this, how could any child in Mass Vaccination (Unknown Victims) reasonably reject a principle that licenses us to choose Vaccine 2?[20]

Underpinning this argument is a different way of thinking about the justification of social risk. According to *ex ante contractualism*, the strength of someone's personal reasons for rejecting a principle licensing a risky action depends on the quality of the prospect that the action gave her *ex ante*. A person's harm-based complaint against a loss she suffers must, therefore, be *discounted* by her ex ante unlikelihood of suffering a loss and by her ex ante likelihood of benefiting from the risky action. If the person's ex ante prospect from the risky action was good enough, she may have no reasonable complaint under the Greater Burden Principle, even if she is unlucky and ends up suffering a greater loss than anyone would have suffered under an available alternative.[21]

20. Alternatively, consider the following *sequential* decomposition of Mass Vaccination (Unknown Victims): A doctor must, over a period of time, see one million young children, each in the same situation as Clara in Individual Vaccination. Due to the Law of Large Numbers, the doctor can foresee that, if he gives each child Vaccine 2, at some point misfortune will strike and some child will be killed by the virus. Does this mean that, while it would be justifiable to give Clara Vaccine 2 in Individual Vaccination, it would be unjustifiable for the doctor to adopt a policy of giving every child Vaccine 2 in the sequential case? This is what ex post contractualism implies, but it seems scarcely credible. The potential benefits for each child of receiving Vaccine 2 in the sequential case are just as great as they were for Clara in Individual Vaccination, and the risk faced in exchange is no greater than what was deemed acceptable to Clara in that case. It seems implausible that a policy of giving every child Vaccine 2 should be unjustifiable to a child, not because of how it is likely to affect that child, but just because, under that policy, *other* children will, at different points in time, be exposed to the same probabilistically independent risk. For an argument along similar lines, see Tom Dougherty, "Aggregation, Beneficence, and Chance," *Journal of Ethics and Social Philosophy* 72 (2013): 1–19.

21. This is not to claim that prospects have nonderivative value or disvalue in themselves. As I will argue in Section 5, people are not made worse off by the mere risk of suffering a loss, nor do they benefit from the mere chance of receiving a benefit, holding everything else constant. But while what ultimately matters to individuals' well-being is whether they end up being benefited or harmed, what might matter morally is to ensure a justifiable distribution of prospects. I owe this point to Derek Parfit, "Justifiability to Each Person," *Ratio* 16 (2003): 368–90.

Ex ante contractualism, as I am proposing it, must be sharply distinguished from the contractualist rule-utilitarian view espoused by John Harsanyi.[22] According to Harsanyi, in order to decide an interpersonal trade-off under certainty, such as Mass Vaccination (Known Victims), in a manner that is justifiable to each person, we should imagine all affected parties placed behind a hypothetical veil of ignorance that deprives them of all knowledge of their personal identity. The right course of action is that which all parties would, in this situation, endorse on grounds of self-interest.

To Harsanyi, the hypothetical veil of ignorance device offers two principal attractions: First, it promotes impartial and unbiased decision making by modeling our moral principles on the choices of hypothetical deliberators who lack the ability to tailor principles to their personal traits and needs. Behind the veil of ignorance, even selfish deliberators are forced to give fair consideration to the interests of every person concerned—for all they know, they might turn out to *be* that person. Second, the veil of ignorance device promises to reduce a difficult interpersonal trade-off to a more tractable problem of *intra*personal choice under risk. Since none of the affected parties possesses information that would give them reason to prefer a different course of action than anyone else, we can expect deliberation behind the hypothetical veil of ignorance to render a unanimous verdict.

In Mass Vaccination (Known Victims), this verdict would be to produce Vaccine 3. Given no information about her identity, Harsanyi maintains, each child must assume that there is a 999/1000 chance that Vaccine 3 will work for her and only a 1/1000 chance that it will not (if she turns out to be one of the doomed children).[23] This intrapersonal gamble is equivalent to that of Individual Vaccination, and there we concluded that it would be in a child's interest to take the risk. Therefore, according to Harsanyi, choosing Vaccine 3 can be justified to each person and is morally right.

22. See, for instance, John Harsanyi, "Morality and the Theory of Rational Behavior," in Sen and Williams, *Utilitarianism and Beyond*, pp. 39–62.

23. John Rawls's more famous use of the veil of ignorance device differs from Harsanyi's in this regard. Whereas Harsanyi assumes that every deliberator behind the veil has an equal epistemic probability of occupying any position beyond the veil, Rawls himself eschews this equiprobability assumption in favor of a "thick" veil of ignorance, which deprives the deliberators of all probabilistic knowledge.

Scanlon, however, rejects Harsanyi's hypothetical veil of ignorance approach, arguing that it mischaracterizes the nature of justification to individuals. By modeling interpersonal trade-offs on analogy with a single individual's choices under risk, the hypothetical veil of ignorance mechanism fails to give each person the separate and individual concern that she is due. That it is often rational to take risks in pursuit of greater benefits in our own lives does not, ipso facto, justify making someone suffer burdens for the good of other people. Thus, "the question of what *everyone* could reasonably agree to or what no one could reasonably reject" is "a quite different question [from] what would maximize the expectations of a single self-interested person choosing in ignorance of his true position."[24] For a moral principle to be valid is for it to be justifiable to each person from her *own* point of view, without artificial informational restrictions.

Unlike Harsanyi's *hypothetical* veil of ignorance method, which works by depriving people of information that they in fact have, ex ante contractualism of the kind I am proposing comes into its own in situations where our inability to foresee the individual outcomes of our actions places us behind what might be called a *natural* veil of ignorance. It is natural uncertainty itself, not the strictures of a hypothetical veil of ignorance, that forces people to deliberate in the absence of complete knowledge about how they will be affected by the action in question. To claim that *such* situations can be illuminatingly analyzed by analogy with intrapersonal gambles under risk is not to commit ourselves to a view that reduces interpersonal justification in general to a question of self-interested choice under informational restrictions.

Indeed, if ex ante contractualism is sound, it suggests a sharp distinction between the two Mass Vaccination cases. Choosing Vaccine 3 in Mass Vaccination (Known Victims) means the doomed children face the prospect of certain death. By contrast, administering Vaccine 2 in Mass Vaccination (Unknown Victims) gives no child a comparably bad prospect. While it is certain that *some* children will die, there is *no child* who is certain to die if we administer Vaccine 2. Indeed, given the significant burden of partial paralysis under Vaccine 1, a course of action that

24. Reibetanz Moreau, "Contractualism and Utilitarianism," p. 122.

promises to avoid this certain burden at the cost of a tiny risk of death seems in every child's interest. Given this, selecting Vaccine 2 is justifiable to every child, according to ex ante contractualism.

A. The Case of Knowable Victims

In discussing the contrast between Mass Vaccination (Known Victims) and Mass Vaccination (Unknown Victims), I have so far focused on two classes of cases: cases where we *know* the identities of individuals who will suffer if we perform a certain action and cases where, given the available evidence and current scientific know-how, we *cannot know*, at the time of acting, who will suffer if we act in a certain way. But this distinction is not exhaustive.

Consider, briefly, an intermediate case, which is like Mass Vaccination (Known Victims), except that the children who will be killed by the virus if we choose Vaccine 3 are not currently known but merely know*able*, in the sense that we have the necessary evidence and wherewithal to determine their identities *prior* to administering Vaccine 3. For instance, suppose we know that Vaccine 3 will be completely ineffective for a child if and only if she carries a rare gene G. Moreover, we know that there are 1,000 such children among the one million threatened by the virus. What we do not currently know is *which* children carry G. However, a simple and costless genetic test would allow us to find out.[25] Call this case Mass Vaccination (Knowable Victims).

Intuitively, the fact that the victims in this case are not known but merely know*able* does not make a moral difference. "Surely," the thought goes, "if vaccinating would be unjustifiable in Mass Vaccination (Known Victims) because we already know the identities of the doomed children, our ignorance of the doomed children's identities in Mass Vaccination (Knowable Victims) would be a very poor excuse for choosing Vaccine 3. For this lack of information about the doomed children's identities is one that we are entirely free to remove by carrying out the simple genetic test."

25. Alternatively, imagine that the test has *already* been carried out, and that its results—though not yet known to us—are stored in a filing cabinet where we could easily consult them.

Ex ante contractualism supports this intuitive verdict.[26] To see why, return first to the contrast between Mass Vaccination (Unknown Victims) and Mass Vaccination (Known Victims). According to the Argument from the Single-Person Case, the crucial difference between these cases is that in the former we can say to each child: "To the best of our knowledge, choosing Vaccine 2 is highly likely to benefit you, and has only a tiny chance of turning out to your disadvantage. Indeed, we could justify giving you this vaccine even in a *single*-person case, where furthering your interests was our sole concern." By contrast, we could *not* say this to each child in Mass Vaccination (Known Victims), in attempting to justify Vaccine 3. For in that case, there are some individuals—the doomed children—of whom we *know* that they will die if we choose Vaccine 3. For a doomed child, choosing Vaccine 3 does *not* correspond to what we ought to do in a case where her interests were our sole concern.

I maintain that in the respects highlighted by the Argument from the Single-Person Case, Mass Vaccination (Knowable Victims) is equivalent to Mass Vaccination (Known Victims). As in Mass Vaccination (Known Victims), we could not justify Vaccine 3 by saying to each child: "To the best of our knowledge, choosing Vaccine 3 is highly likely to benefit you, and has only a tiny chance of turning out to your disadvantage. We could justify giving you this vaccine even in a single-person case." Both halves of this statement are false.

Consider the first half first. Granted, it is true that

(A) In Mass Vaccination (Knowable Victims), given our present ignorance about which children have gene G, there is no child of whom we *know* that will die if we choose Vaccine 3.

Indeed,

(B) In Mass Vaccination (Knowable Victims), given our present ignorance about which children have gene G, it is rational to believe, of any given child, that it is highly likely to benefit if we produce Vaccine 3, and very unlikely to be burdened.

Yet both (A) and (B) are things we can affirm *only* as long as we do not avail ourselves of all the freely available evidence, namely, by carrying

26. I thank an anonymous referee for pressing me to clarify my reasoning about this point.

out the simple and costless genetic test. However, something that we can only affirm without taking into account freely available evidence is not something we can affirm "to the best of our knowledge." Hence, neither (A) nor (B) can support

(C) In Mass Vaccination (Knowable Victims), it is true of every child that, to the best of our knowledge, choosing Vaccine 3 is highly likely to benefit the child, and has only a tiny chance of turning out to her disadvantage.

Nor is it true that, by choosing Vaccine 3 in Mass Vaccination (Knowable Victims), we are treating each child in a way that would have been justifiable in a single-person case. Consider a single-person analogue involving only Clara and her doctor: The doctor can administer a genetic test, which will determine whether Clara carries G or not. Given that this test is simple and costless, and will help the doctor to select the appropriate treatment, there is no justification for failing to administer the test. But if the test indicates that Clara carries gene G, it will not be justifiable for the doctor to administer Vaccine 3. He must administer Vaccine 1. Consequently, we cannot say to each child in Mass Vaccination (Knowable Victims) that by giving him Vaccine 3, we are treating him in a way that would have been justifiable in a single-person case. This will be false for those children who carry G.

I conclude that if ex ante contractualists are right to use the Argument from the Single-Person Case to draw a moral line between cases like Mass Vaccination (Unknown Victims) and Mass Vaccination (Known Victims), then cases involving victims whose identities are not presently known but are know*able* should often be treated on analogy with cases involving *known* victims—at least if it is simple and costless to ascertain their identities.

Suppose, by contrast, that although we could find out who carries gene G in Mass Vaccination (Knowable Victims), doing so would be *extremely costly* for us. Intuitively, this seems to bring this case closer to Mass Vaccination (Unknown Victims), where information about the identities of the victims is simply unobtainable prior to acting.

I believe that ex ante contractualism has the resources to underwrite this intuition as well, though I can only sketch a tentative proposal. Once again, it is helpful to first consider a single-person analogue of the case in

question. Suppose that, in principle, the doctor has the ability to determine whether Clara carries the rare gene G, but that doing so would be extremely costly to the public purse (Clara herself cannot bear these costs). In this case, the doctor may have a valid justification for not carrying out the genetic test, and for instead giving Clara the treatment that gives her the best prospect, *given ignorance about her exact genotype*, namely, Vaccine 3. He might say to Clara: "To the best of my knowledge, given justifiable limits on the resources we can be expected to expend in gathering further information about your particular case, Vaccine 3 is highly likely to benefit you, and has only a tiny chance of turning out to your disadvantage."

This reasoning carries over to Mass Vaccination (Knowable Victims). Unlike in the *costless* case we discussed above, choosing Vaccine 3 in the *costly* case may be justifiable to every child. For in this case, by giving all children Vaccine 3, we would be doing for them the same thing that would have been justifiable to each in a single-person case.

V. REJECTING THE ARGUMENT FROM IRRELEVANT INFORMATION AND EX POST CONTRACTUALISM

Since ex ante contractualism avoids some of the implausible implications of the ex post view and seems independently attractive, it pays to take a closer look at the ex post contractualist's Argument from Irrelevant Information, to see whether we can find fault with it.

Consider the following reconstruction of the Argument from Irrelevant Information:

(1) An action is permissible if and only if it is justifiable to each person.
(2) In Mass Vaccination (Known Victims), selecting Vaccine 3 cannot be justified to the doomed children.
(3) In Mass Vaccination (Unknown Victims), we know that, if we select Vaccine 2, the luckless children will suffer the same losses that the doomed children suffer under Vaccine 3, although we cannot know the identities of the luckless children.
(4) The fact that we cannot know the identities of the luckless children is irrelevant to the question of whether selecting Vaccine 2 is justifiable to each person.

(5) In Mass Vaccination (Unknown Victims), there are some children to whom selecting Vaccine 2 cannot be justified, although we cannot know the identities of these children. [from (2), (3), and (4)]

∴ Selecting Vaccine 2 is impermissible in Mass Vaccination (Unknown Victims). [from (1) and (5)]

The crucial step in this argument is premise (4). Proponents of the Argument from Irrelevant Information assume that the lack of information about individual outcomes that characterizes cases of social risk is a merely epistemic phenomenon, devoid of moral significance. It is morally irrelevant, they claim, that we cannot know who will be harmed by the risky action, as long as we know that some persons will be harmed. Instead of affecting whether Vaccine 2 is justifiable to each person, our lack of information about individual outcomes merely masks the identity of those children to whom the choice of Vaccine 2 cannot be justified.

Premise (4), however, directly contradicts what defenders of ex ante contractualism affirm. And they, as we saw, have an argument for rejecting premise (4): given lack of information about individual outcomes, it is in each patient's ex ante interest to take the gamble that Vaccine 2 offers. By administering Vaccine 2, we thus do for each child what we would do on her behalf in a single-person case, where that child's well-being was our only concern. The same cannot be said in an interpersonal trade-off under certainty.

Lest they beg the question against their opponent, defenders of ex post contractualism thus need an independent argument for affirming premise (4). Let me review what seem to me the three strongest contenders.

A. The Argument from Certain Loss

The first and, I believe, the most natural way of reading the passage from Moreau that I cited above is as proposing what I shall term the *Argument from Certain Loss*. According to this argument, information about individual outcomes is not needed in cases of social risk, because whoever ends up being burdened by the risky action will have a complaint that is augmented by the fact that it was certain that *someone* would be burdened. Recall Moreau: "As long as we know that acceptance of a

principle will affect *someone* in a certain way, we should assign that person a complaint that is based upon the full magnitude of the harm or benefit, even if we cannot identify the person in advance."

The Argument from Certain Loss, however, implies an odd asymmetry between single-person gambles and cases of social risk-imposition. A proponent of the Argument from Certain Loss would agree that a principle licensing us to choose Vaccine 2 in Individual Vaccination is justifiable to Clara. Even if Clara is unlucky and the gamble turns out badly for her, she has no weighty complaint. After all, in this single-person case, the likelihood that *someone* will be killed is the same as the likelihood that Clara will be killed. By contrast, the choice of Vaccine 2 is said to be unjustifiable in Mass Vaccination (Unknown Victims), since here it is a statistical certainty that *some* child will be killed. What, therefore, makes the moral difference between the single-person and the social case is our enhanced knowledge of the *overall* pattern of outcomes—not any difference in the attractiveness of the gamble offered to each individual.

But why, we may wonder, should a luckless child have a greater complaint than Clara, just because, even if *he* had not been unlucky, it is statistically certain that some *other* child would have died instead? This looks suspiciously like a new form of interpersonal aggregation: the combination of complaints by different individuals *at different possible worlds*, depending on who happens to be unlucky at that possible world. But if contractualists reject the combination of claims by different people at the *same* possible world, then a fortiori they ought not to allow different people to aggregate their complaints across different possible worlds.

B. The Argument from Lack of Concern

According to the next argument, the lack of information about individual outcomes in Mass Vaccination (Unknown Victims) does not matter, because we know that whoever ends up being burdened could make the following complaint: given that we knew for certain that some children would die if we chose Vaccine 2—an outcome that we could have avoided by opting for Vaccine 1 instead (which would not have imposed a similarly severe loss on anyone)—our willingness to nonetheless opt for Vaccine 2 betrays a lack of concern for the eventual losers.

This argument also misses its mark: it might be claimed that our action shows too little concern for avoiding that *there will be* losers. This,

however, expresses a preoccupation with the overall *shape* of the outcome, as might arise from an impersonal moral principle such as telic egalitarianism or the priority view. As I argued above, contractualism's exclusive focus on *personal* reasons does not allow us to appeal to such impersonal principles.

By contrast, what no luckless child can claim is that the risky action showed too little concern that *he will be* a loser. Again, our action takes no greater risks with the well-being of any child than its analogue in Individual Vaccination, where defenders of ex post contractualism agreed that taking the gamble on Clara's behalf did not show insufficient concern for *her*. Unequal outcomes are not always evidence of unequal concern or treatment, especially when they are produced by a chancy causal process.

C. The Argument from Determinism

Let us consider one final argument for the ex post perspective. We are conducting our discussion under the working assumption that at the macro-physical level that we deal with in most cases of social risk-imposition, nature is, for all intents and purposes, deterministic. All probabilities we encounter in these contexts are merely *epistemic*, that is, owed to our incomplete knowledge of the state of the world and the laws of nature that govern it, not *objective*, that is, due to indeterminacy at the level of physical reality itself.[27]

But if this assumption is true, might it not spell trouble for the ex ante contractualist? After all, if determinism is true, then someone who loses

27. How exactly to characterize the notion of "objective" probability remains a hotly debated topic in contemporary metaphysics. The intuitive contrast with the "epistemic" notion of probability that I have been working with is, nonetheless, clear enough. Recapitulating from above: to say that, from the perspective of an agent, some event e has an epistemic probability p of occurring is to say that, given the evidence available to the agent, her rational degree of credence in the proposition "e will occur" would be p. Note that, given incomplete knowledge of the world, p can be > 0 even for events that are, as a matter of fact, physically impossible. By contrast, to claim that e has an objective probability p of occurring is to locate probability "in the world." Here, p characterizes the physical propensity or disposition of a given type of physical situation to yield event e or to yield a long-run relative frequency p of such an outcome-type. To say that e has an objective probability p of occurring is thus to imply, minimally, that it is in fact possible that e will occur. For a seminal discussion of the distinction between epistemic and objective probability, see David Lewis, "The Subjectivist's Guide to Objective Chance," in *Philosophical Papers*, vol. 2 (Oxford: Oxford University Press, 1986).

a gamble was, in a sense, always going to lose it. In terms of objective, if not of epistemic, probability, there was *no* chance that the gamble would turn out in her favor.

From this angle, it may now seem that the problem with imposing social risks is not just, as the Argument from Certain Loss assumed, that the agent knows with statistical certainty that *someone* will end up losing if she imposes the risk. Rather, the agent knows that whoever ends up losing was in fact *certain* to lose, in objective terms. But surely, the argument goes, a principle allowing us to condemn someone to an *objectively certain* loss, so that others may enjoy smaller individual benefits, is one that the loser could reasonably reject. On this view, there is an important difference between cases of social risk such as Mass Vaccination (Unknown Victims) and single-person gambles like Individual Vaccination: In cases of social risk, we know with statistical certainty that there will be some losers (by the Law of Large Numbers and the assumption that individual outcomes are probabilistically independent). Moreover, if determinism is true, we know that those who end up losing were, in fact, *objectively* certain to lose. In single-person gambles, we also know that *if* the person loses, she was objectively certain to lose. However, unlike situations of social risk, we have no statistical reasons to believe that we *are* in a case in which someone will lose. Indeed, in Individual Vaccination, the available evidence makes it rational for us to believe, with credence 0.999, that Clara will win her gamble (indeed, that she is *objectively certain* to do so). This may make it permissible to give her Vaccine 2, even if doing the same in Mass Vaccination (Unknown Victims) would not be justifiable to each person. Call this the *Argument from Determinism*.

If sound, the Argument from Determinism would have strongly revisionary implications. Ex ante reasoning could continue to be employed for social risk-imposition under conditions of causal indeterminacy. However, in all domains in which the assumption of causal indeterminacy is not plausible (and this, I have said, may include most risky human activity), contractualist decision makers who cared about making their actions justifiable to each person would have to resort to an ex post model of justification. We have already seen that this would have a profoundly confining effect on the conduct of social policy, which often involves an ineliminable element of risk. Given these stakes, we should take a hard look at the Argument from Determinism to examine whether

the fact of causal determinism really has the profound normative significance that this argument attributes to it.

In so doing, it is helpful to first go back to a single-person case like Individual Vaccination, and to ask whether, in *that* situation, the presence or absence of objective chance makes any difference to Clara's reasons for wanting to take Vaccine 2. Let us distinguish two scenarios: in the *indeterministic* scenario, the success of Vaccine 2 is objectively chancy, contrary to what I have been assuming so far. That is, given Vaccine 2, Clara has a 999/1000 objective chance of surviving the virus unharmed and a 1/1000 objective chance of dying. In the *deterministic* scenario, the outcome of Vaccine 2 is merely epistemically chancy. Given the available evidence, there is a 999/1000 epistemic chance that Vaccine 2 is *objectively certain* to allow Clara to survive the virus unharmed, and a 1/1000 epistemic chance that she is *objectively certain* to die if we give her that medicine.

Now, I cannot see any good reason why Clara should hope to be in the indeterministic rather than the deterministic scenario. In both situations, Clara's rational degree of credence that she will actually avoid death is the same. And from the perspective of prospective self-interest, this is all that appears to matter. (If you are not yet persuaded, the following analogy might help: All else equal, is there any greater self-interested reason to buy a ticket for a lottery where the prize will be allocated by a genuinely indeterministic mechanism [a "quantum randomizer"] than to buy a ticket for a lottery with the same epistemic odds where a ticket is either sure to win or sure to lose, in the objective sense [a scratch-card lottery with a preagreed winning number]? Again, it is hard to think of any such reason. If you had a ticket to the indeterministic lottery, and someone offered to trade it for an equivalent ticket to the deterministic lottery *plus* a small amount of cash, there would be no good reason to refuse this offer.)

Nor could it be said that if the gamble goes badly, it is *worse* for Clara to become disadvantaged as the result of a causal process that was deterministic as opposed to indeterministic. The prudential value of receiving a chance of some good, whether epistemic or objective, is strictly *parasitic* upon the value of the good itself. It is no benefit, in and of itself, to have had an objective rather than a merely epistemic chance of receiving the good. If, therefore, the gamble does not end up giving Clara the benefit she sought, it would be in no way *better* for her to have "at least"

had an objective rather than a merely epistemic chance of obtaining the benefit. Once we hold fixed the fact of Clara's death, there is no retrospective reason for her to care whether this happened due to a process that was objectively or merely epistemically chancy.

I conclude that there is no reason, neither a prospective nor a retrospective one, why a single patient like Clara should care whether she is in the deterministic or the indeterministic scenario.

These arguments carry over to the case of social risk. Here, too, the only thing each child has reason to care about is whether he is being treated in a way that will allow him to actually survive the disease unharmed. But, for each child, our rational degree of credence that Vaccine 2 will actually allow that child to survive unharmed is just as high in the deterministic as in the indeterministic scenario.

Of course, we know that not all children will, in fact, survive if we choose Vaccine 2. Some will die because the vaccine fails to protect them from the virus. But this is due to the Law of Large Numbers, paired with the assumption that individual outcomes are probabilistically independent, not the issue of determinism vs. indeterminism. Moreover, as I argued in discussing the Argument from Certain Loss, this fact does not increase any child's personal reasons for rejecting a principle licensing us to select Vaccine 2. To claim it does would be to engage in the interpersonal aggregation of claims across different possible worlds.

All told, it is hard to see how the fact that we can predict the overall pattern of outcomes in Mass Vaccination (Unknown Victims) would give anyone a greater harm-based complaint against Vaccine 2 than in a single-person gamble. I conclude that the Argument from Irrelevant Information fails. By assimilating cases of social risk-imposition to interpersonal trade-offs under certainty, the ex post contractualist conflates the moral significance of "we know that someone will be harmed" with that of "there is someone who we know will be harmed."

Let us take stock: if my arguments for ex ante and against ex post contractualism have been successful so far, I have shown how the epistemically chancy nature of many public-policy decisions may allow a contractualist to justify the imposition of social risks that will foreseeably result in a pattern of benefits and burdens that his theory would have forbidden him from bringing about through an interpersonal trade-off under certainty.

A catchy way of summarizing the appeal of ex ante contractualism is that it permits us to "count the numbers without aggregating": For a population of a given size, the greater the number of people who would foreseeably experience some harm H as the result of some socially risky action, the worse the prospect that this action gives to each individual (and mutatis mutandis for benefits). In this sense, an ex ante view "counts the numbers," because the number of those who will end up being harmed by the action is an indicator of the prospect that it gives to each individual. At the same time, ex ante contractualism does not aggregate, because the structure of justification remains individualistic. Overall outcomes matter because of what they say about individual prospects, not because of their total goodness or the combination of claims.

VI. EX ANTE RULES AND THE DECOMPOSITION TEST

The intermediate result I have reached stands in sharp contrast with the position defended by T. M. Scanlon in *What We Owe to Each Other*. There are a number of passages in which Scanlon explicitly rejects ex ante contractualism. He 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.[28]

Scanlon concedes 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." However, he argues that this does not 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." Rather,

> 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

28. Scanlon, *What We Owe to Each Other*, p. 208.

> as an indicator of the care that the agent has to take to avoid causing harm.[29]

Scanlon's central objection to ex ante contractualism is closely connected with what I shall call the problem of *ex ante rules*. Showing how an ex ante contractualist can meet this problem will allow me to introduce a crucial refinement to the view.

By "ex ante rule" I mean any rule the adoption of which, at some time t_1, is in everyone's interest, but which licenses or requires some agent to act at a later time t_2 in a way that benefits some but significantly burdens others. The locus of social risk, in the case of ex ante rules, is not in the *effects* of the agent's actions at t_2, which may be perfectly predictable. Rather, it is the result of not knowing, ex ante, whether one will be among those benefited or burdened by the *application* of the rule at t_2. This could be because the ex ante rule itself incorporates some randomizing mechanism, like a lottery, or because it is triggered by circumstances that are imperfectly predictable. The main objection that Scanlon marshals against ex ante contractualism in *What We Owe to Each Other* is that it could support ex ante rules that are intuitively unacceptable. He writes:

> Consider any principle licensing us to impose very severe hardships on a tiny minority of people, chosen at random (by making them involuntary subjects of painful and dangerous medical experiments, for example), in order to benefit a much larger majority. A contractualist would want to keep open the possibility that such a principle could reasonably be rejected because of the severe burdens it involves. But this would be effectively ruled out on the [ex ante view], according to which the weight given to these burdens, as grounds for rejecting the principle, would be sharply discounted because only a very small fraction of the population would actually suffer them.[30]

Adopting the rule in Scanlon's example ("conduct medical experiments on randomly chosen subjects to advance medical knowledge") may be in

29. Ibid., p. 209.
30. Ibid., pp. 208–9.

everyone's ex ante interest, since one's probability of being harmed in a medical experiment may be small enough to be outweighed by the prospect of benefiting from greater medical knowledge. Nonetheless, such a rule seems intuitively unacceptable to most nonconsequentialists.

In order to avoid the possibility of unacceptable ex ante rules, Scanlon argues, in *What We Owe to Each Other*, that ex ante contractualism must be rejected.[31] But neither does he embrace the kind of ex post contractualism that I have been arguing against so far. Instead, he proposes an account of the morality of risk imposition that focuses on "reasonable precautions" and on the cost of avoiding the harm in question. To complete the above quotation:

> 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. Our reactions to the medical experiment case, . . . for example, depend heavily on whether the harm in question is directly inflicted on particular people or whether it occurs "by accident," that is to say, occurs despite the fact that reasonable precautions have been taken. . . . The difference between these two versions of the medical experiment examples does not have to do with the cost to the victims: the harm is just as bad when suffered "by accident" as when it is inflicted. The difference rather lies in the cost of avoiding these ways of bringing harm. I believe our reactions to these cases reflect the view that, except for a few very unusual kinds of cases, we can accept a prohibition against intentionally inflicting serious harm on others. But the cost of avoiding all behavior that involves risk of harm would be unacceptable.[32]

I have misgivings about Scanlon's way of distinguishing between the two kinds of medical experiments that he mentions. To see why, it will help to consider two more fleshed-out cases:

31. But see note 10.
32. Ibid., p. 209.

> *Risky Treatment*: At time t_1, a doctor administers a risky treatment to 100 paraplegic children. He knows that, for each child, the treatment has an 85 percent chance of curing her paraplegia, while foreseeing that approximately 15 children will die at t_2 as a result of receiving the treatment.

Compare this to the following case involving an ex ante rule:

> *Human Experiment*: At t_1, a doctor organizes a lottery among a group of 100 paraplegic children, which selects 10 of them by a random process. At t_2, he conducts medical experiments on these 10 children, foreseeing (but not intending) that they will die in the course of the experiments. He knows for certain that the knowledge gained in this way will allow him to cure the remaining 90 children of their paraplegia.

Like Scanlon, I think that it may be permissible for the doctor to undertake the risky treatment in the former case; yet an ex ante rule licensing us to perform lethal human experiments on 10 randomly selected children seems clearly impermissible—despite the fact that Human Experiment gives every child a *better* ex ante prospect than Risky Treatment.

As Scanlon tells it, the difference between Risky Treatment and Human Experiment must be due to a difference in the cost of avoiding harm. The thought is that avoiding the intentional imposition of harms is usually not very costly, whereas entirely eliminating any risk of harm can be prohibitively costly and, hence, too confining. But while this is often true, *need* there be a difference in cost between avoiding harm in Risky Treatment and Human Experiment? And if there were not, would that render these two cases morally on a par?

Suppose that in both cases we have already taken maximum possible precautions to avoid unnecessary harm. Given the available technology, it simply is not feasible to further reduce the likelihood of harm in either case, however much money or effort we throw at the problem. Moreover, suppose that reaching this maximum level of protection has cost the same amount in Risky Treatment and Human Experiment. Still, the intuition that most nonconsequentialists have is that the latter case remains impermissible, whereas the former is permissible.

Scanlon may retort that in Human Experiment, there is an easy way to avoid causing any harm, namely, by simply not conducting the experiment. But this does not solve our problem. That same option is available in Risky Treatment. Here, too, the doctor can avoid the risk that anyone will die by not treating any of the paraplegic children. *Pace* Scanlon, it is hard to think of a non-question-begging explanation of why it would be too costly or confining to avoid carrying out the risky treatments in the former case, but not the human experiments in the latter case.

Indeed, depending on how we understand "costs" and "confinement" in this context, the costs of not causing any harm may actually be higher in Human Experiment than in Risky Treatment. If we meant by "cost" the "social opportunity cost" of refraining from the experiments, the cost of not conducting the human experiment is greater than that of not undertaking any risky treatments. Five *more* children, who could have been cured, will remain paraplegic if the doctor does not apply the ex ante rule in Human Experiment than if he abstains from the risky treatment.

Hence, if we take the permissibility of potentially harmful actions to be a function of the costs of avoiding any harm, either there should be no difference between Risky Treatment and Human Experiment, or it is the latter, not the former, that should be permissible.

I turn now to what strikes me as a more promising explanation for the intuitive difference between the two cases: The moral ideal underpinning contractualism is that we regulate our behavior in such a way that our actions are, at all times, justifiable to each other. The problem with the ex ante rule in Human Experiment, I submit, is that it fails to live up to the "at all times" part of this ideal. In general, I propose that contractualists ought to adopt the following:

> *Decomposition Test*: If a rule or procedure can be decomposed into a sequence of distinct causal stages, each of which involves the voluntary action of some agent (or of a surrogate for human agency, such as a programmed machine), then it is permissible to adopt and act on this rule or procedure only if the actions it requires at every stage are justifiable to each person *at that time*.

The Decomposition Test implies that if an ex ante rule requires us to act, at a later stage, in a manner that cannot be justified to every person *at*

that time, then the ex ante rule as a whole cannot be justified to each person.[33] Its justifiability to each person unravels by backward induction. I shall refer to the ex ante view supplemented by the Decomposition Test as *stage-wise ex ante contractualism*.

This refined version of ex ante contractualism captures the difference between Risky Treatment and Human Experiment. Even though it may be in every child's ex ante interest to be signed up for *either* scheme, Human Experiment is not justifiable to each person because it fails the Decomposition Test.

Risky Treatment is what I call a *single-stage* procedure, because a human agent intervenes only once, namely, in administering the treatment at t_1. This action, I have argued, can be justified to each child as being in her ex ante interest. That any particular child ends up being harmed as a result of the risky treatment is not because of anyone's choice or agency, but due to the vagaries of the natural lottery. Given the available technology, scientific know-how, and so on, the risk of being harmed by the risky medicine is the unavoidable consequence of taking the medicine at t_1.

Human Experiment, by contrast, is a *multistage* procedure and must therefore pass the Decomposition Test. The doctor intervenes *twice*: once at t_1, to organize the lottery, and then again at t_2, to conduct experiments on the 10 selected children, which will foreseeably result in their demise. This latter intervention, however, fails the test of justifiability to each person at that time. It violates the Greater Burden Principle, since it imposes on the selected children the grave and certain harm of death, in order to spare the others from the smaller burden of paraplegia.

This conclusion might be undermined if the patients had *alienated* their right to life through some ex ante agreement or *waived* it through contemporaneous consent. But, by assumption, the patients are young children who lack the moral power to enter into binding agreements or waive their rights. Given this, the mere fact that participation in Human Experiment was in all children's ex ante interest at some point in the

33. A somewhat similar idea is mooted by James Lenman, "Contractualism and Risk Imposition," *Politics, Philosophy and Economics* 7 (2008): 99–122.

past does not suffice to justify the doctor in carrying out the lethal experiments at t_2.[34]

Although the Decomposition Test tracks our intuition that administering the risky treatment is permissible while carrying out the human experiments is not, some philosophers may doubt whether it does so for the right reasons. One such philosopher is Rahul Kumar, another proponent of ex ante contractualism.

According to Kumar, the correct explanation of why it is impermissible for the doctor to conduct human experiments is not that his action fails the Decomposition Test. Indeed, for him, "the grounds for reasonably rejecting any principle permitting the kind of medical experimentation in question have nothing to do with the imposition of either harm or the risk of harm."[35] Rather, according to Kumar, the correct objection is that

> the experimentation requires that it be permissible to involuntarily involve individuals in it as experimental subjects. Any such permission would make it the case that the authority to make decisions concerning how an individual's body may be used has, at least in part,

34. The distinction between "alienating" and "waiving" a right or claim is due to Joel Feinberg, "Voluntary Euthanasia and the Inalienable Right to Life," *Philosophy & Public Affairs* 7 (1978): 93–123. There are many situations in which, by entering into an agreement, you confer on others the right, given certain conditions, to treat you in ways *at a later time* that, absent your advance consent, would not have been justifiable to you (for example, by making financial demands on you, as in insurance contracts; by ignoring certain interests or desires of yours, as in surrogate motherhood arrangements; or by preventing you from speaking publicly about an issue, as in nondisclosure agreements). In a second group of cases, we believe that persons lack the moral power to *alienate* their claims against certain kinds of treatment by advance consent. Here, we think, a person's right not to be treated in certain ways can only be temporarily *waived* by *contemporaneous* consent, not alienated in advance. (Think of consent to sexual intercourse.) In a third class of cases, we think that even contemporaneous consent cannot render the action in question morally permissible. These are action-types that, in all but supreme emergencies, we consider categorically impermissible—whether or not the person contemporaneously consents to such treatment. For instance, even if, at that very moment, I volunteer to be killed and eaten by you for your gustatory pleasure, my consent is not "morally effective"; that is, it does not have the power to make your action morally permissible. (Incredibly, such things do happen: http://news.bbc.co.uk/2/hi/europe/3286721.stm.) My intuition, for what it is worth, is that, even if the patients in Human Experiment are consenting adults, not young children (as I have assumed), this case would fall at least into the second and possibly into the third category. However, I do not have the space to pursue this question further.

35. Kumar, "Risking and Wronging," p. 37.

> been ceded to others. Each individual has good reason to want this kind of decision-making discretion to be solely her own.[36]

As Kumar writes further:

> An individual's sense of herself as an independent agent is intimately tied up with her having sole decision-making authority over how her body is to be used. Things may happen to a person's body that she does not choose, either through chance events or as the foreseeable result of what others do. But that is just part of what it is to be an embodied agent in the world; maintaining complete control over what happens to one's body is not possible. Having things just happen to you is not, however, the same as another having the right (whether or not it will ever be made use of) to make a decision about how your body is to be used, or what will be done to your body, without your having a say in the matter. Such a right would undermine the decision-making authority over the use of her body that is partly constitutive of a person's self-conception as an autonomous agent.[37]

On Kumar's account, the key moral difference between Risky Treatment and Human Experiment is that a principle permitting the doctor's actions in Human Experiment would have to grant him the authority to *use* the bodies of those children selected by the lottery without their having a say in the matter. By contrast, the foreseeable deaths in Risky Treatment are merely something that "happens" to the unlucky children; no child is used against her will.

I do not wish to rule out that the factor emphasized by Kumar's account plays a role in explaining the moral difference between Risky Treatment and Human Experiment.[38] It is nonetheless a mistake to conclude that the Decomposition Test does no work in explaining the moral

36. Ibid., pp. 37–38.

37. Ibid., p. 37.

38. One source of skepticism is the emphasis that Kumar places on the fact that the patients in Human Experiment are being used *against their will*. This suggests that if the patients in this case had been adults, they could have made the doctor's actions permissible by providing consent to being used in lethal medical experimentation. As I state in note 34, I am not convinced that, exceptional circumstances aside, people have the power to provide morally effective consent to being used in this way. However, settling this difficult question is beyond the scope of this discussion.

difference between Risky Treatment and a multistage process like Human Experiment. Suppose Human Experiment was impermissible solely in virtue of (i) using the bodies of some children against their will, as Kumar suggests, but not, more generally, because (ii) it requires us to act at t_2 in a way that allows two children to suffer the certain and uncompensated burden of death in order to spare others a less severe burden. This would suggest that a variant of Human Experiment that featured (ii) but not (i) would not be impermissible. But this, I believe, is not the case. Consider:

> *Preliminary Treatment*: At t_1, a doctor administers a preliminary treatment to 100 paraplegic children. The treatment has a 90 percent chance of working for each child. If the preliminary treatment works, it will allow the doctor to cure that child of her paraplegia at t_2 by giving her a small dose of some drug D. (If a child who has become "treatable" in this way does not receive drug D at t_2, she will remain paraplegic.) However, if the preliminary treatment does *not* work, the child will become poisoned and die, unless the doctor saves his life at t_2 by giving him a larger dose of drug D.
>
> Because he possesses only a limited supply of drug D, the doctor can foresee that if he administers the preliminary treatment at t_1, he will face the following choice at t_2: he can *either* (1) cure the 90 treatable children of their paraplegia, leaving the 10 poisoned children to die, *or* (2) save the 10 poisoned children's lives, with the result that all 100 children continue to be paraplegic.

Like Human Experiment, Preliminary Treatment differs from Risky Treatment in the way highlighted by the Decomposition Test: If the doctor administers the preliminary treatment at t_1, he will face a *separate* decision at t_2 whether to save the 10 poisoned children from death, or to cure the other 90 of paraplegia. But unlike in Human Experiment, if the doctor opts to cure the 90 treatable children, no child's body will be *used* against its will at t_2, including those children who are not treated. Not treating a child is not a way of using his body. Nonetheless, I believe a contractualist ought to conclude that it would not be justifiable to each person for the doctor to administer the preliminary treatment and then choose option (1). At t_2, the doctor faces a choice between saving 10 children from certain death and curing 90 from paraplegia, which is a

considerably less severe burden. It is hard to deny that a child's individual claim to be saved from certain death is stronger than that to be cured from paralysis. Hence, under the competing claims model, we must conclude that it would be morally impermissible for the doctor to administer the preliminary treatment and then opt to cure the 90 treatable children, even though doing so would not use any child against her will. Kumar's account is unsatisfactory, because it fails to capture this.[39]

Let me end this section by considering a final intricacy. It might be suggested that we could make it the case that Human Experiment satisfies the Decomposition Test by eliminating the need for any agent to intervene at time t_2. What if we could set up a fully automatic mechanism that takes the decision of whether or not to conduct the lethal experiments at t_2 out of our hands? Consider:

> *Automatic Experiment*: At t_1, a doctor sets in motion an unstoppable device, which first randomly selects 10 paraplegic children from among the 100, and then, at t_2, automatically conducts lethal medical experiments on the 10 chosen children, without the need for any further human involvement. The knowledge gained in this way is certain to allow us to cure the remaining 90 children.[40]

It might be thought that this proposal satisfies the Decomposition Test, since the only stage at which a human agent intervenes is at t_1, to set in

39. Note, incidentally, that Preliminary Treatment also functions as a counterexample to the claim that the key moral difference between Human Experiment and Risky Treatment is that in the former case, some children are being *harmed as a means* to curing others of their paraplegia. If that were the key difference between these cases, then in Preliminary Treatment it would not be wrong to use drug D to cure the 90 treatable children of their paraplegia at t_2, leaving the 10 poisoned children to die. For in doing so, the doctor would not be harming any child as a means to benefiting the others.

40. In *Morality, Mortality*, 2:303, Frances Kamm considers an example that raises similar issues. She asks whether it would be permissible for a community to construct an ambulance in which the brakes would automatically be disabled whenever the driver is rushing a large group of patients to the hospital, thereby preventing the driver from braking for pedestrians. It is foreseeable that, under this scheme, more people would be run over by ambulances. But suppose that the number of patients whose lives would be saved by the scheme is *greater*, so that everyone's probability of survival would be increased *ex ante*. Would it be permissible to construct such an ambulance? Like my Automatic Experiment, Kamm's example uses a technical fix to tie the agent's hands at a later point in time. And, like me, Kamm denies that this fact would make it permissible to operate such a scheme.

motion the device. But *this* action is in every child's ex ante interest at that time. Hence, a principle licensing us to set in motion the automatic device could not be reasonably rejected by any child.

I do not think that stage-wise ex ante contractualism is committed to this counterintuitive conclusion. Begin by noting that I have worded the Decomposition Test in a way that already takes account of the possibility of replacing human agency at t_2 by a preprogrammed device. A procedure counts as multistage under the Decomposition Test, not only if it can be decomposed into a sequence of distinct causal stages, each of which involves human agency, but also if some of these stages involve the actions of *surrogates* for human agency, such as machines that human agents have set up to perform certain actions in their stead. The question thus remains: is the procedure contemplated in Automatic Experiment justifiable at each stage?

The answer to this question is "no." Instead of avoiding a conflict with stage-wise ex ante contractualism, Automatic Experiment merely pushes the moral problem back to an earlier point, namely, to the decision, at t_1, to set in motion an unstoppable device that will automatically perform the lethal experiments at t_2. The only rationale for setting up such a device, however, is precisely to avoid the need to perform an action at t_2 that, we know, could not be justified to each person at that time. In effect, the doctor would be saying: "Since, in Human Experiment, it would be unjustifiable for me to follow through on my ex ante rule and kill the 10 chosen children at t_2, I will set in motion a device that will take this decision out of my hands, while delivering the result I desire." This, you will agree, is not a convincing solution. In Human Experiment, the killing at t_2 is unjustifiable, not because it is carried out by a human agent rather than a machine, but because it avoidably places the uncompensated and severe burden of death on 10 children in order to cure 90 others from the less serious burden of paraplegia. And this feature, of course, carries over to Automatic Experiment. Hence, if it is morally unjustifiable for the doctor to kill the 10 chosen children at t_2, it would likewise be unjustifiable to set in motion a device at t_1 that would deliver the same result automatically. Demands for justification cannot be evaded by getting machines to do our dirty work for us.

I conclude that, if we amend ex ante contractualism in the way I have suggested, namely, by incorporating the Decomposition Test, the view can successfully distinguish between Risky Treatment, on the one hand,

and Human Experiment and its permutations, on the other, and thus successfully addresses Scanlon's worry.

VII. THE PROBLEM OF IDENTIFIED VS. STATISTICAL LIVES

Unfortunately, we are not home and dry yet. As I shall now argue, the Achilles' heel of ex ante contractualism is not, as Scanlon assumed, the problem of ex ante rules, but a different problem—what I shall dub the problem of identified vs. statistical lives. Consider the following illustration:

> *Miners (1 vs. 100)*: Gareth, a miner, is trapped in a collapsed shaft. If we do not save him, he is virtually certain to die within days. However, a rescue will be costly. Suppose we must choose between the following two options:
>
> - *Rescue*: Spend all our available funds to rescue Gareth.
> - *Prevention*: Spend our available funds to improve safety at this mine, reducing the risk of future accidents. If we choose this option, the risk of death for each of the other 100 people working at this mine of dying in a future accident will be reduced from 3 percent to 1 percent. We expect that this will save two lives (though we cannot know whose). However, Gareth will die.

Human beings have a well-documented psychological propensity to consider rescuing identified individuals in imminent peril more important than preventing the loss of "statistical lives," that is, lives that will predictably be lost to known risk factors in the future unless we intervene, but whose identities we cannot know, at least at present. But is this propensity anything more than an empirical fact about how humans tend to think and act? Is there any reason to believe that it also corresponds, at least partially, with what we have *moral* reason to do? In recent years, many philosophers, economists, and lawyers have greeted this notion with skepticism.[41]

41. The label "identified" vs. "statistical" lives is due to Thomas Schelling, "The Life You Save May Be Your Own," in *Problems in Public Expenditure Analysis*, ed. Samuel Chase (Washington, D.C.: Brookings Institute, 1968). Another seminal discussion that is critical of the "identified victims bias" is Charles Fried, "Value of Life," *Harvard Law Review* 82 (1969):

Interestingly, ex ante contractualism departs from mainstream opinion on this score. As I shall now show, ex ante contractualism suggests a principled rationale for privileging the saving of identified over statistical lives in cases like Miners (1 vs. 100). I believe that this feature of ex ante contractualism contains at once an important kernel of insight, yet also makes pure ex ante contractualism untenable as an account of what it is for an action to be morally right, all things considered.

To see how an ex ante contractualist would think about this issue, begin by noting that, in two important respects, Miners (1 vs. 100) differs from the social risk cases that have preoccupied us thus far. First, the choice between Rescue and Prevention is what I call *competitive at the ex ante stage*. As with interpersonal trade-offs under certainty, there are two distinct groups—Gareth, the miner presently trapped underground vs. the other miners, who are at risk from future accidents—whose interests are opposed from the start. Gareth is obviously in need of Rescue and will not benefit from our choosing Prevention; all other miners will not (personally) benefit from Rescue, and have an interest in Prevention, since this reduces their own risk of losing their lives in a future mining accident. This sets Miners (1 vs. 100) apart from Mass Vaccination (Unknown Victims). In that case, our ignorance of the eventual individual outcomes is *complete*, so that, while we know that some individuals will end up gaining and others losing ex post, at the ex ante stage it is in *everyone's* interest to see the risky action performed.

Second, in Miners (1 vs. 100), individual outcomes are unknowable in only one of the two possible courses of action: if we opt for Rescue and against Prevention, any given miner may or may not perish in a future mining accident. We cannot know whom this fate will befall in advance. In this sense, the lives that are saved by Prevention, and will be lost if we choose Rescue, are "statistical lives." By contrast, if we do not choose Rescue, Gareth, the miner currently trapped underground, is sure to die.

1415–37. A recent influential discussion of the problem in the context of the "treatment vs. prevention" debate in bioethics is Dan Brock and Daniel Wikler, "Ethical Challenges in Long-Term Funding for HIV/AIDS," *Health Affairs* 28 (2009): 1666–76. I respond to Brock and Wikler in Johann Frick, "Treatment versus Prevention in the Fight against HIV/AIDS and the Problem of Identified versus Statistical Lives," in *Identified versus Statistical Lives: An Interdisciplinary Perspective*, ed. Glenn Cohen, Norman Daniels, and Nir Eyal (New York: Oxford University Press, 2015). That article anticipates some of the points I develop in this and the next two sections.

This will be a loss of an "identified life." The natural veil of ignorance, we might say, only covers outcomes under one of the two possible courses of action in this case.

It is clear how an act-consequentialist would reason in the Miners (1 vs. 100) case. All else equal, two lives saved, taken together, is a better outcome than one life saved. The fact that we do not at present (and perhaps never will) know the identities of the two statistical lives that are saved by Prevention does not affect the impersonal goodness of this outcome, and hence is morally irrelevant. We ought to choose Prevention.

Ex post contractualism concurs with this verdict. As I pointed out, although Scanlonian contractualism does not permit the combination of complaints, Scanlon does allow that numbers can break ties when individual complaints of equivalent strength are opposed.[42] Since we assume that *someone* will lose his life whether we choose Rescue or Prevention, the strongest individual complaint that can be made against either option, according to the ex post contractualist, is the same: someone has an undiscounted complaint against losing his life prematurely. The only difference is that if we choose Rescue, there will be two such complainants, whereas under Prevention there will only be one, namely, Gareth. Hence, if the ex post contractualism was correct, Scanlon's tie-breaker view would apply in Miners (1 vs. 100) and favor Prevention.

Ex ante contractualism disagrees with both these views. However, the *justification* for the ex ante contractualist perspective must differ from the one I offered in discussing Mass Vaccination (Unknown Victims). Because in Miners (1 vs. 100) there are two distinct groups whose interests are competitive ex ante, the Argument from the Single-Person Case, which I employed in defending the risky action in Mass Vaccination (Unknown Victims), cannot be appealed to here. Whichever course of action we select, we will not be able to tell every person concerned that this is exactly what we would have done, had we been guided solely by a concern for *his* interests. We thus require a different defense of ex ante contractualism for cases in which the natural veil of ignorance is only partial.

To motivate this new argument, consider a feature of Miners (1 vs. 100) which is typical of most instances where we must choose between

42. See note 6.

preventing the loss either of identified or of statistical lives: Saving Gareth, the identified miner, means saving someone who but for our intervention was certain (or at least very likely) to die in short order. By contrast, preventing the loss of statistical lives in Miners (1 vs. 100) comes about by *slightly* reducing an *already quite small* risk of death for each member of a larger group.

This difference, I believe, is morally significant. In particular, it seems hard to deny that a single individual has a stronger claim to be rescued from a *certain* (or very likely) harm of size H than a different individual has to have her risk of suffering a harm of equal size slightly reduced, especially if her risk of suffering this harm is quite low to begin with. Suppose, for instance, that we face the following choice:

> *Miners (1 vs. 1)*: We can either save Gareth from certain death or reduce the risk of death of *one* other miner (call him Morgan) from 3 percent to 1 percent.

I think it is clear that Gareth has a much stronger claim to our assistance than Morgan, and a much greater complaint if we do not help him. This is because a person has a much stronger *prudential interest* in avoiding certain death than in having his risk of death reduced from 3 percent to 1 percent. This is not to suggest, of course, that if Morgan is unlucky and his 3 percent risk of death materializes, he will not lose just as much as Gareth would, if left to certain death. The claim is merely that, since Morgan is so much less likely to die than Gareth is, and since we can reduce his risk of death by so much less, Morgan's claim to our assistance is weaker than Gareth's.

In conjunction with the contractualist prohibition against combining claims, this judgment about the 1 vs. 1 case gives rise to a simple yet powerful argument for privileging the saving of identified over statistical lives in cases like Miners (1 vs. 100). I call this the *Pro Identified Lives Argument.*

I have just argued that

(1) In the Miners (1 vs. 1) case, given a choice between saving Gareth and reducing Morgan's risk of death from 3 percent to 1 percent, Gareth has a much stronger claim to our assistance than Morgan.

Consider now the choice that we actually face in Miners (1 vs. 100):

> *Miners (1 vs. 100):* We have to choose between saving Gareth from certain death and reducing the risk of death to each of 100 other miners from 3 percent to 1 percent.

It seems true that

(2) In Miners (1 vs. 100), each miner—considered as an individual—would be treated no differently by a principle that licensed us to save Gareth than Morgan was treated in the 1 vs. 1 case. That is, each miner would remain exposed to a 3 percent risk of death when they could instead have had their risk of death reduced to 1 percent.

Hence, it seems plausible that

(3) In Miners (1 vs. 100), none of the miners has a stronger individual claim to be protected from their 3 percent risk of death than Morgan has in the Miners (1 vs. 1) case.

Furthermore, according to Scanlon's individualist restriction,

(4) In deciding which course of action to pursue, we must not *combine* the claims of different individuals. Rather, the right action is that which satisfies the strongest *individual* claim.

But from (1), (3), and (4) it follows that

(5) In Miners (1 vs. 100), we ought to save Gareth from certain death rather than reduce the risk of death to each of 100 other miners from 3 percent to 1 percent, for this is what satisfies the strongest individual claim.

If the argumentative strategy of the Pro Identified Lives Argument is sound, it would provide a qualified defense for prioritizing the saving of identified over statistical lives.[43] On this account, the distinction between identified and statistical lives is morally relevant *when and because* it

43. The argument would have to be slightly modified in order to deal with cases in which the potential losses to identified and statistical victims are not equal in size. Suppose, in particular, that the certain loss of the identified victim was smaller in size than that of the statistical victims. In that case, for the argument to go through, it would have to be established that the identified victim has a stronger claim to be protected from her certain loss than any potential statistical victim has to be protected from a mere risk of

coincides with the distinction between saving those who are certain (or very likely) to die unless we intervene and helping those who face a smaller individual risk of death that we can reduce by a lesser amount.

By contrast, ex post contractualists, who deny premise (3) of the Pro Identified Lives Argument, once again seem committed to an implausible asymmetry. As we have seen, a principle licensing us to save Gareth in Miners (1 vs. 100) treats no individual miner any worse than Morgan would be treated by a corresponding principle in Miners (1 vs. 1). Nonetheless, ex post contractualists hold that in Miners (1 vs. 100), any miner who ends up dying would have a full, undiscounted complaint, even though his own risk of harm was no greater than that of Morgan in Miners (1 vs. 1). Somehow, the fact that, if we save Gareth, it is foreseeable that *someone* from the group of 100 will die in a future accident is thought to strengthen the complaint of whoever turns out to be harmed. As I argued above in rebutting the Argument from Certain Loss, this move involves a tacit appeal to the interpersonal combination of complaints across different possible worlds, and ought, for that reason, to be rejected by Scanlonian contractualists.

I note in passing that if a contractualist should discount harm-based complaints by their ex ante improbability even in cases that are competitive at the ex ante stage, this would help avoid the main worry in Elizabeth Ashford's influential piece "The Demandingness of Scanlon's Contractualism."[44] Ashford argues that contractualism is implausibly demanding in that it would seem to prohibit activities that come with a tiny risk of very serious harms, such as commercial air travel, when some of these risks are borne by people who do not stand to benefit from the risky activity itself: the Amish farmer living under a busy flight path is at some small risk of being killed by a falling plane, but does not himself partake in the benefits of modern air travel. Ashford's argument is premised on the assumption (which she takes over from Scanlon) that it is improper for a contractualist to discount harm-based complaints by their improbability of occurring. Once we give up this assumption, as I have argued we should, Ashford's problem goes away. If we discount the

---

suffering some greater loss. I believe that this, too, is often the case. However, I will not pursue this possibility in the following.

44. Elizabeth Ashford, "The Demandingness of Scanlon's Contractualism," *Ethics* 113 (2003): 273–302.

reasons that an Amish farmer who bears the risk of being hit by a falling plane has for rejecting a principle licensing commercial air travel by the tiny likelihood of such an accident occurring, and compare them to the *certain* and significant burden that banning air travel would place on many people, it is plausible that a principle licensing air travel passes the test of justifiability to each person.[45]

VIII. A PROBLEM FOR CONTRACTUALISM

Let us return to Miners (1 vs. 100). Unfortunately, there is a sense in which the Pro Identified Lives Argument succeeds *too well* at capturing a moral difference between preventing identified and statistical losses. In fact, the argument's conclusion is *not at all* sensitive to the number of statistical deaths that we allow to happen by saving Gareth from certain death. Suppose that instead of 100 there are 1,000 people working at the mine, each of whose risk of death we could have reduced

45. In "Risking and Wronging," Rahul Kumar offers a similar ex ante contractualist solution to Ashford's case, arguing that a representative Amish person, "Jeb," could not reasonably reject a principle licensing air travel. This is so, he writes, because the strength of Jeb's reasons for rejection must be discounted by the improbability of his ever being harmed in an accident: "It seems plausible to take the risk of ending up harmed as a consequence of an aviation accident imposed on someone in Jeb's circumstances to be quite low (aviation accidents happen, but not that often)" (p. 48). However, Kumar then makes the following puzzling addendum: "Say the Amish experience a population explosion, such that their number grows exponentially, making the territory in question very densely populated. The risk of an Amish person being harmed would then no longer be quite low, as the probability of the area in which, for example, a bit of falling aircraft debris lands being occupied will have gone up substantially. This can be taken to be a respect in which the numbers are relevant to determining what is permissible, but not in the way ruled out by the contractualist strictures" (p. 48). This, however, is confused. The risk to Jeb, or any other representative Amish person, does not increase as a result of the population explosion: airplane debris can only fall in one location, and the representative person's chances of finding himself in precisely this spot when the accident occurs is no higher after the population explosion than before. What *has* gone up is the likelihood that *some* Amish person (de dicto, not de re) will be harmed by falling airplane debris. But since this does not correspond to any *individual* experiencing a greater risk of harm, no Amish has stronger personal reasons for objecting to air travel than before. *Pace* Kumar, ex ante contractualism therefore cannot take the numbers into account in the way he suggests. You may feel that this represents a worrying limitation of the view. Notwithstanding the fact that no *individual* person is more likely to be killed by falling airplane debris than before, the greater likelihood that *someone* will be killed should, intuitively, affect the balance of reasons for or against air travel, even if it does not decisively tip the scales. And you would be right. As I argue in the following section, ex ante contractualism's insensitivity to numbers in cases that are competitive *ex ante* is the major shortcoming of this view.

from 3 percent to 1 percent. Call this Miners (1 vs. 1,000). In this case, choosing to save Gareth rather than improving mine safety will come at the cost of 20 statistical deaths. The Pro Identified Lives Argument, however, is impervious to this information: as long as Gareth's claim to be saved from certain death is stronger than that of any of the 1,000 other miners to be protected from their small individual risk of death, the logic of the argument requires us to save his life.

Most people will find this conclusion hard to accept. It is consistent with many people's intuition that we should give *somewhat* greater weight to the saving of identified over statistical lives (so that we *might* prefer to save Gareth from certain death in Miners [1 vs. 100] even at the cost of allowing two statistical deaths to happen). However, a view that would, in principle, allow *any number* of statistical lives to be lost in order to save one identified individual from certain death seems deeply implausible.

IX. TOWARD A PLURALIST ACCOUNT OF MORAL RIGHTNESS

Here is where we stand: As I argued in Sections 2 to 5, ex post contractualism is implausible, because it lacks the resources to draw a moral distinction between intuitively permissible instances of social risk, such as Mass Vaccination (Unknown Victims), and morally questionable interpersonal trade-offs like Mass Vaccination (Known Victims). By contrast, ex ante contractualism handled these cases with aplomb. Now, however, it seems that ex ante contractualism faces serious problems of its own, because its tendency to privilege the avoidance of identified over statistical losses can yield verdicts about cases like Miners (1 vs. 1,000) that are implausibly extreme. Is there a way out of this impasse?

I believe that there is, but only by scaling back the ambitions of contractualism as a moral theory. In *What We Owe to Each Other*, Scanlon claims that his contractualist formula gives an account of what it *is* for an action to be right or wrong, in the domain of interpersonal morality. Given this theoretical ambition, however, the topic of risk poses a formidable stumbling block to the contractualist: as we have seen, there appears to be no proposal for fleshing out the notion of contractualist justification—neither the ex ante view, nor the ex post view, nor Scanlon's own proposal—on which Scanlon's contractualist account of wrongness does not yield strongly counterintuitive verdicts about some cases involving risk.

I believe that the first step to escaping this impasse is for Scanlon to embrace a suggestion made, for broader reasons, by Derek Parfit in *On What Matters*. According to Parfit's proposal, being disallowed by a principle that no one could reasonably reject should be understood not as what it *is* for an action to be wrong, but rather as a *wrong-making property* of actions. More precisely, it is a *higher-order* wrong-making property, under which certain first-order wrong-making properties can be subsumed.[46]

Second, I believe Scanlon ought to accept that not *all* wrong-making properties of actions are captured by the higher-order property of being disallowed by a principle that no one could reasonably reject. According to this *pluralist* proposal, while the contractualist formula captures an important *class* of pro tanto moral reasons that contribute to making actions right or wrong in the domain of interpersonal morality, reasonable rejectability in the contractualist sense is not the only relevant consideration in determining whether an action is right or wrong, *all things considered*.[47]

46. Derek Parfit, *On What Matters*, vol. 1 (Oxford: Oxford University Press, 2011), pp. 368–70. In more recent writings, Scanlon appears willing to consider this proposal; see T.M. Scanlon, "Replies," *Ratio* 16 (2003): 424–39.

47. It is important to understand how my pluralist proposal differs from Scanlon's own view, which could, in a different sense, be characterized as "pluralist" as well. The intended scope of Scanlon's contractualist theory is the domain of interpersonal morality. That is, his contractualist formula seeks to capture what it is for an action to be right or wrong in our dealings *with other persons*, that is, in the part of morality "having to do with our duties to other people, including such things as requirements to aid them, and prohibitions against harming, killing, coercion, and deception." Scanlon, *What We Owe to Each Other*, p. 6. Following Frances Kamm, "Owing, Justifying, and Rejecting," in *Intricate Ethics: Rights, Responsibilities, and Permissible Harm* (Oxford: Oxford University Press, 2007), pp. 455–90, I shall refer to this part of morality as "M1." However, Scanlon does not claim that his contractualist account of right and wrong *exhausts* morality "in the broad sense in which most people understand it." Scanlon, *What We Owe to Each Other*, p. 6. For instance, an action may be morally criticizable, in a *broader* sense, because it fails to respond appropriately to the impersonal value of some wonder of nature, like the Grand Canyon (ibid., pp. 218–23), or to the badness of animal suffering (ibid., pp. 177–87). (Call the part of morality not captured by Scanlon's formula "M2.") In this sense, Scanlon may be called a "pluralist about the moral," since he holds that different accounts of moral rightness and wrongness are appropriate for M1 and M2. By contrast, my proposal denies that the contractualist formula provides an adequate account of what it is for an action to be right, all things considered, even *within* the domain of M1. It is pluralist in a stronger sense, because it holds that the contractualist formula captures an important—but not the *only*—class of pro tanto moral reasons that contribute to making actions right or wrong in the

If not moral rightness tout court, is there a different moral notion that subsumes the specific class of pro tanto reasons captured by the contractualist formula? While no candidate fits the bill perfectly, our concept of *equity* perhaps comes closest. Being equitable is a higher-order right-making property of actions; yet an action's being equitable does not entail that it is necessarily morally right, all things considered, nor is an action that fails to treat everyone equitably necessarily morally wrong, all things considered. Second, equity is an "individualistic" moral notion. Being equitable is a property that attaches to actions, not in virtue of their overall or aggregate effects, but in virtue of how they treat each person individually. This matches the antiaggregative character of contractualist thinking embodied in Scanlon's Individualist Restriction. Third, there is a comparative dimension to equity: whether an action treats someone equitably often depends on how others are being treated by comparison. This is reflected in Scanlon's Greater Burden Principle and the comparative character of the competing claims approach.

Reading the contractualist formula as offering an account of equity also meshes well with the particular specification of the contractualist idea—stage-wise ex ante contractualism—that I have argued for in this article. To review some of the cases we have considered: It is inequitable to let Jones suffer painful electrical shocks in order to spare the football spectators the much smaller burden of having their match interrupted. Likewise, it is inequitable to subject two children (even chosen by a random process) to lethal medical experiments in order to cure others of paraplegia. Finally, it seems inequitable to abandon Gareth to a certain death in order to protect the other miners from a much smaller risk of suffering this fate. Ex ante contractualism explains all three of these judgments. By contrast, there seems to me no inequity in choosing to produce Vaccine 2 in Mass Vaccination (Unknown Victims), whereas choosing Vaccine 3 in Mass Vaccination (Known Victims) appears grossly inequitable. Again, ex ante contractualism explains the moral difference, whereas act-utilitarianism and ex post contractualism are committed to viewing both actions as being morally on a par.

---

domain of *interpersonal* morality. I thank an anonymous referee for prompting me to clarify this point.

One payoff of reading the competing claims model as offering an account not of what it is for actions to be right or wrong, all things considered, but of a particular class of pro tanto moral reasons is that it makes conceptual space for the thought that an action may satisfy the strongest individual claim and yet be wrong, all things considered. Likewise, an action may be right, all things considered, without satisfying the strongest individual claim. There is, I believe, *at least* one further moral consideration that ought to guide our actions, and that must sometimes be traded off against the reasons of equity captured by the competing claims model: I believe that, alongside reasons of equity, we have *independent* pro tanto reasons that stem from the effects of our action on people's *well-being*.

Return once more to the case of Miners (1 vs. 1,000). If ex ante contractualism is correct, the fact that there will foreseeably be 20 statistical deaths if we opt to save Gareth does not give any of the other miners a strong personal complaint. No miner, considered individually, could claim to have been treated inequitably if instead of further reducing his own small risk of death, we opted to save Gareth from certain death. But does this mean that it could not be wrong, all things considered, to save Gareth rather than to avert 20 statistical deaths? Surely not. That there will be a much greater loss of life if we save Gareth is a fact that seems to matter *in its own right*, in a way that is not captured by the competing claims model. On the pluralist view that I propose, an action's consequences in terms of people's well-being constitute an *independent* right- or wrong-making property of that action, irrespective of whether they also give any individual a weighty personal complaint against the action.

In some cases, reasons of equity and reasons of well-being point in the same direction, and consequently there is no tension between these two components of the pluralist view. This will often be the case in situations such as Mass Vaccination (Unknown Victims) or Risky Treatment, where the risky action or policy is in the ex ante interest of all concerned.

By contrast, situations involving trade-offs between identified and statistical victims are ones in which no available action or policy will be in the ex ante interest of everyone concerned. Here, there is the potential for reasons of equity and reasons of well-being to pull in opposite directions. Giving equity its due in such cases means that when the two options do not differ significantly in terms of their expected

consequences, reasons of equity may outweigh a concern with people's well-being. This may be the case in Miners (1 vs. 100), though nothing rides on endorsing this particular judgment. By contrast, when the difference in the expected consequences is as large as in Miners (1 vs. 1,000), it is plausible that our concern with equity is swamped by our reasons for averting an outcome in which many more lives are lost.[48]

Somewhere in between these poles, there is a tipping point at which reasons of equity begin to be outweighed by reasons of well-being. I cannot say precisely where this tipping point lies, nor can I set forth a general algorithm by which to determine the relative weights one should place on equity and well-being in other cases. But nor do I consider it my role to do so. The aim and ambition of moral philosophy should be to inform our judgment, by making us alive to the relevant ethical considerations, not to abolish the need for judgment altogether.

48. Of course, this sketched solution leaves many details to be filled in. For instance, what exactly determines the strength of our reasons of well-being for choosing one option over another? The answer depends, in part, on the appropriate procedure for *combining* the reasons of well-being that we have for helping different individuals. In a situation like Transmitter Room, is it the case, as classical utilitarians assume, that the very weak reasons that we have for preventing each of the millions of spectators from suffering mild frustration *together* give us a stronger reason for action than our weighty reason for sparing Jones his agony? I do not think that they do. Following Frances Kamm's suggestion in *Intricate Ethics*, I believe that, in a context where we have a very weighty reason of well-being to prevent someone from suffering a serious harm, the much weaker reasons to prevent other individuals from suffering a very minor harm are not "relevant" and cannot together outweigh our reason to prevent the serious harm. See Kamm, *Intricate Ethics*, pp. 33–37. If this is the case, then saving Jones would be the right thing to do, not just because it satisfies the reasons of equity captured by the competing claims model, but also because it is what we have the most reason of well-being to do. Contrast this with Miners (1 vs. 1,000), where what is at stake on *both* sides of the trade-off are individual losses of equivalent magnitude (deaths), which are therefore "relevant" to one another as far as our reasons of well-being are concerned. Since, moreover, we could save a much larger number of deaths by focusing on enhancing mine safety, this is the option that we have the most reason of well-being to pursue. Reasons of well-being and reasons of equity thus pull in opposite directions in this case, whereas they arguably do not in Transmitter Room. Unfortunately, I cannot pursue these points in more detail here.