Journal of Ethics and Social Philosophy
VOL. 19, NO. 2 · FEBRUARY 2021
https://doi.org/10.26556/jesp.v19i2.1024

CONTRACTUALISM, COMPLAINTS, AND RISK

Bastian Steuwer

One of the most prominent and forceful objections against utilitarianism is that it fails to respect the separateness of persons. Utilitarianism aggregates all benefits and burdens of an action in order to decide whether or not the action is permissible. It seems as though the utilitarian treats all benefits and burdens an action produces as if they were the benefits and burdens of one entity or one system of ends. In doing so, utilitarianism fails to respect the separateness of persons as individuals and as systems of ends of their own.[1]

In response to utilitarianism's failure to respect the separateness of persons, nonconsequentialists have proposed conceptions of morality that are based on the competing claims or complaints that individuals can raise. Placing the commitment to individual claims or complaints at the heart of morality seems a promising route to ensure respect for the separateness of persons. The most systematic of these proposals is contractualism as developed by T. M. Scanlon. Scanlon argues that an act's rightness or wrongness depends on its justifiability to each person. As a test for justifiability, Scanlon proposes that the permissibility of an act depends on whether it follows from a principle that no one can reasonably reject. An act is permissible only when no one can reasonably reject a principle that entails the permissibility of that act. One natural idea is that the individual with the largest complaint has most reason to reject a principle. It then appears that a principle can be reasonably rejected only when the largest complaint is larger than the complaint anyone else could bring forward against any alternative principle.[2] Recently Scanlonian contractualism has received scrutiny for the way it deals with cases where risks, rather than certainties of harm and benefit, are at stake.[3] My discussion in this article will focus on Scan-

1 See Gauthier, *Practical Reasoning*, 125–26; Nagel, *The Possibility of Altruism*, 133–40; Rawls, *A Theory of Justice*, 23–26; Nozick, *Anarchy, State, and Utopia*, 32–33; Nagel, *Mortal Questions*, ch. 8; and Nagel, *Equality and Partiality*, chs. 4–8.

2 See Scanlon, "Contractualism and Utilitarianism," and *What We Owe to Each Oth*er, ch. 5.

3 See Reibetanz, "Contractualism and Aggregation"; Ashford, "The Demandingness of Scanlon's Contractualism"; Lenman, "Contractualism and Risk Imposition"; Fried, "Can

lonian contractualism, but my conclusions may apply more widely to any moral theory that places the idea of justifiability and individual complaints or competing claims at the heart of morality.

The debate around contractualism and risk is typically framed as a debate between two opposing views. *Ex ante* contractualism is concerned with prospects while *ex post* contractualism is concerned with outcomes.[4] I believe that this framing is unhelpful. What can it mean to say that a theory of risk impositions is concerned with outcomes when it is designed to provide guidance in cases where we are uncertain about the outcome? With the help of a sequence of thought experiments from Michael Otsuka, I provide a more helpful way of understanding what is at stake between different contractualist approaches to risk (section 1).[5] In addition, the sequence allows me to propose a new view on contractualism and risk, which I call *objective ex ante contractualism* because of the special importance it gives to objective, as opposed to epistemic, probability. My version of contractualism focuses on the complaints of would-be victims whose fates are already determined. After discussing the sequence, I will show that a natural extension of the sequence highlights that two conditions that *ex post* contractualism should ideally fulfill are inconsistent with one another (section 2). In section 3, I will present the defense of my objective *ex ante* view by arguing that it provides us with the best model of the key contractualist idea of acting in ways that are justifiable to each. Section 4 responds to objections.

1. OTSUKA'S SEQUENCE

> *Dust*: A comet is en route to the midwestern United States carrying a pathogen that will soon lead to millions of people being infected and dying. The government is briefed on two alternative ways of containing the pathogen. The first option has the side effect that a different hazard will

Contractualism Save Us from Aggregation?"; James, "Contractualism's (Not So) Slippery Slope"; Fleurbaey and Voorhoeve, "Decide As You Would with Full Information!"; Frick, "Uncertainty and Justifiability to Each Person"; Scanlon, "Reply to Zofia Stemplowska"; John, "Risk, Contractualism, and Rose's 'Prevention Paradox'"; Frick, "Contractualism and Social Risk"; Kumar, "Risking and Wronging"; Otsuka, "Risking Life and Limb"; Horton, "Aggregation, Complaints, and Risk"; and Rüger, "On *Ex Ante* Contractualism."

4 For the former, see James, "Contractualism's (Not So) Slippery Slope"; John, "Risk, Contractualism, and Rose's 'Prevention Paradox'"; Kumar, "Risking and Wronging"; and Frick, "Contractualism and Social Risk." For the latter, see Fleurbaey and Voorhoeve, "Decide As You Would with Full Information!"; Otsuka, "Risking Life and Limb"; and Rüger, "On *Ex Ante* Contractualism."

5 Otsuka, "Risking Life and Limb," 77–88.

> be released over Florida. It is known that it would cause Bob Johnson, a resident of Boca Raton, to lose one leg. Unfortunately, Bob Johnson cannot be evacuated in time. The second alternative has the side effect that the hazard will have to be released in a dust cloud over California. Each of forty million Californians faces a small risk of death, and it is known that exactly one Californian will die. The Californian who will die has a genetic predisposition that will cause his or her death upon being subjected to the dust.

Intuitively, the right course of action here would be to release the hazard over Florida and cause Bob Johnson to lose a leg. But it appears that contractualism struggles to explain this intuitive answer. Bob Johnson's complaint against choosing to release the hazard is not discounted. It is certain that he will suffer. The complaints of the Californians, however, should be discounted. For each of the forty million Californians, the likelihood of being the one who dies is only a one in forty million. Although death is terrible, a one in forty million chance of death is not altogether that terrible. We often incur similar risks when crossing the road, cooking with gas, or swimming in the ocean. The complaint against the imposition of the risk of death would suddenly be a rather trivial moral complaint. How can such a trivial moral complaint outweigh Bob's quite serious complaint of losing his leg?

One way for contractualism to accommodate the case is by pointing out that all the complaints combined add up to something significant: a complaint of the magnitude of certain death. But this response leads to highly counterintuitive results in other cases.

> *Transmitter Room*: Jones, a worker in a TV transmitter room, has had an accident. He is now lying on the floor and suffering extremely painful electric shocks. There is only one way to save Jones, namely by interrupting the current transmission signal for about fifteen minutes. This in turn will cause millions of viewers who want to see the football World Cup match that is in progress to be upset.[6]

If we add up the complaints due to inconvenience and upset of all the millions of viewers, it seems that they will outweigh Jones's complaint against being subject to pain. But here it is clear that we should not let Jones suffer for the relatively mild loss of missing fifteen minutes of a football match. We should not aggregate morally trivial complaints so that they outweigh serious moral complaints of single individuals.

6 Scanlon, *What We Owe to Each Other*, 235.

Otsuka, in his discussion of Dust, resists this solution and instead points to a different feature of the case. Unlike in Jones's case, in Dust there is one person who will experience grave harm. The aggregated complaints add up to the real-life predicament of one person. We do not need to imagine a social entity that experiences the harms of dying, but there is an individual made out of flesh and blood who will die. It is merely a fact concerning our informational limitations that prevents us from identifying that person in the same manner we were able to identify Bob Johnson. Yet we can still say something about the individual who is going to die. The person who is going to die is "the Californian with the genetic predisposition." The complaint of the Californian with the genetic predisposition is nondiscounted. Her (or his) complaint would outweigh Bob Johnson's complaint.

Now is the complaint of the Californian with the genetic predisposition a complaint *ex ante* or *ex post*? *Ex post* contractualism can account for this complaint. We know that the result of the action will be one person dying. Since the outcome distribution of the action is already known to us, an *ex post* contractualist can peek ahead, anticipate this distribution, and assign complaints to those affected by it.

But can *ex ante* contractualism? I think it can. The Californian with the genetic predisposition is a person with a determinate identity when we make the decision. Regardless of what happens and regardless of our action, the Californian with the genetic predisposition will always be the same person. If we limit our attention to only those possible worlds that are possible outcomes of our action, then we can say that "the Californian with the genetic predisposition" rigidly designates over this restricted domain of discourse. Since only those possible worlds that constitute possible outcomes of our actions are of interest to us, I will simply refer to such descriptions as "rigid designators."[7] Releasing the hazard over California will impose the certainty of death on this existing person with a determinate identity. From the *ex ante* perspective, the Californian with the genetic predisposition can object to the imposition of a 100 percent risk of death. We do not need to appeal to the outcome of the action *ex post* to make this claim.

This means that our understanding of *ex ante* contractualism should be broader. The classical version of *ex ante* contractualism focuses on the risks as faced by individuals with proper names, or otherwise identifiable individuals. But not all

7 This definition also includes an element of temporality in the *ex ante*/*ex post* distinction. The possible worlds that are possible outcomes of the action are those possible worlds that coincide in their history until the point of action. Rigid designators are descriptions that refer to information that is contained in the shared history. Nonrigid designators are descriptions that refer to information about the future where the possible worlds no longer coincide.

versions of *ex ante* contractualism focus on these risks. The version of *ex ante* contractualism I defend focuses on the complaints that rigidly designated individuals can raise. The two forms of *ex ante* contractualism differ thereby in whose complaints they focus on. This in turn is linked to a distinction between two kinds of risk: epistemic risks (credences) and objective risks (chances).[8] The distinction I am relying on here classifies some probability functions as expressing our uncertain degrees of belief or confidence about the world. These are epistemic probability functions, also called credence functions. By contrast, objective probability functions express a mind-independent idea of probability. The objective probability function, a chance function, reflects information about the world and not about our knowledge of the world. If there are nontrivial objective probabilities, then there are truly "chancy" events. While there are various theories on what chances are, the differences between them are not important for my arguments.[9] What I rely on is solely the contrast between chances and credences.

In Dust, we only have epistemic probabilities for the risks that each identifiable Californian faces. However, we can give objective probabilities for the risk that the Californian with the genetic predisposition faces. This suggests an important link between the question of whose complaints we are interested in and what kind of risk we are interested in. By focusing on rigidly designated individuals, objective *ex ante* contractualism gives primacy to objective risk assessments over epistemic risk assessments. Objective *ex ante* contractualism holds that in a case like Dust where the uncertainty is merely a matter of failing to identify the victim, we should choose descriptions that reveal the objective risks that individuals are facing. This is the "objective" component in objective *ex ante* contractualism.[10]

8 I follow here the orthodox tradition in the philosophy of probability dating back to Rudolf Carnap, who distinguished between two concepts of probability (frequentist and evidential), which are examples of the broader approaches of chance and credence. See Carnap, "The Two Concepts of Probability," 516–25; Eagle, "Chance versus Randomness," sec. 1; and Hájek, "Interpretations of Probability," sec. 3.

9 The most common approaches are frequentism, propensity views, and best systems approaches. In addition, some philosophers embrace a "no theory" approach to chances according to which objective probabilities are not reducible to anything else like frequencies or propensities. For an overview, see Hájek, "Interpretations of Probability"; for the no theory approach, see Sober, "Evolutionary Theory and the Reality of Macro-Probabilities," 148–54. Actual frequentist views are an exception to my claim that my use of objective chance is neutral between the different theories of chance. According to actual frequentist views, objective probabilities only refer to actually occurring frequencies. Under such a view, objective probabilities only represent statistical facts about reference groups and have no obvious moral significance.

10 Importantly, the two kinds of risks are linked in a manner that should guard us from identi-

Let me move on to the next case in the sequence:

> *Wheel*: The case is structurally similar to Dust. Again, we have a comet en route and a disaster about to occur. Again, one of our options is to release the hazard over Florida and cause Bob Johnson's loss of a leg. But now our second option changes. As a side effect of averting the disaster, each Californian will be placed under a gigantic roulette wheel in the sky. The wheel will spin indeterministically and release a roulette ball that will kill exactly one person.

Otsuka reports his intuitive judgment that in Wheel, as in Dust, we should still prefer to release the hazard over Florida, causing the loss of Bob Johnson's leg. But here we cannot rely anymore on the description of "the Californian who is genetically predisposed." Instead, we would need to rely on a description like "the Californian who would be hit by the roulette ball" or "the Californian who would be most harmed by the decision." These descriptions are nonrigid designators since different persons may die due to the falling ball. While the complaints of rigidly designated individuals have to be discounted, the complaints of nonrigidly designated individuals do not. The probability of someone being harmed by the wheel is one. We can peek ahead and assign a complaint to that person. We may think that such statistical persons are still actual persons worthy of respect and with claims that ought to be taken into consideration.[11]

This cannot be reconciled with the *ex ante* perspective. The complaint of the Californian most harmed by the decision is not a complaint of any person with a determinate identity prior to the action. There is no token individual for whom it is true that she has imposed on her a 100 percent risk of death. Accordingly, my objective *ex ante* view holds that releasing the hazard over California is permissible in Wheel.

Anticipating the strongest complaint *ex post* is easy in a case like Wheel. We know for certain how the benefits and burdens will be distributed in the outcome. We only lack information about who will be in which position. I now move on to a case where certainty about the resulting distribution is absent.

---

fying epistemic or objective *ex ante* contractualism exclusively with one kind of risk. Whenever we have an objective probability for a given event (such as Charlotte Williams's being harmed), we should adjust our credence (i.e., our epistemic probability) to match the objective probability. The next case in the sequence is an example of this. This widely accepted claim is an implication of David Lewis's Principal Principle. See Lewis, "A Subjectivist's Guide to Objective Chance."

11 See Daniels, "Can There Be Moral Force to Favoring an Identified over a Statistical Life?" 116; and Otsuka, "Risking Life and Limb," 85–86.

> *Guns*: In this case, we have the option to shoot down the comet with an automated weapons system. Unfortunately, the system also has guns in the sky pointed at each Californian. Each gun is operated by an indeterministic randomizer. The chance for each gun to fire and kill the person it is aimed at is one in forty million. The guns, and thus the risks each gun imposes, operate independently of one another.

The objective risk for each Californian is the same as in Wheel, one in forty million. Any assessment of rigid designators that relies on objective risks will be the same between Wheel and Guns. However, the assessment for nonrigid designators like "the Californian who will be most harmed" changes. Here we move away from certainty about the distribution that will come about and introduce risk as well. There is a 63 percent chance that at least one Californian will die, a 26 percent chance that at least two Californians will die, an 8 percent chance that at least three will die, and so on. What should *ex post* contractualists say about a case like this?

One answer is that Guns highlights the limits of *ex post* contractualism. Under this version of *ex post* contractualism, we should draw a distinction between two types of cases. In some cases, like Dust or Wheel, we know that the risk imposition will lead to harm while in Guns the harm is not guaranteed. Anticipating the complaint of the eventual victim is permitted in Dust and Wheel but not permitted in Guns according to this view. Since we do not know for certain that someone will be harmed, we cannot anticipate this complaint.[12]

The problem with this version of *ex post* contractualism is that it relies on a distinction between risky cases that is morally dubious.[13] Cases with guaranteed harms can easily be transformed into cases without guaranteed harm without changing anything of moral relevance. Take the example of a coin flip with inversely correlated harms and benefits. If the coin lands heads, *A* benefits and *B* is harmed. If the coin lands tails, *A* is harmed and *B* benefits. This is a case of guaranteed harm. *Ex post* contractualism would sometimes rule out this kind of risk even if it is in the antecedent interests of both *A* and *B*. But what if the coin lands on the edge? This would be a freak accident but is nonetheless a possibility. Let us assume that no one will be harmed if the coin lands on the edge. The case is now one without guaranteed harm. If we are not allowed to anticipate any complaint *ex post*, we should do what is in the antecedent interests of both. Sim-

12 Sophia Reibetanz Moreau defends such a view ("Contractualism and Aggregation," 302–4). Victor Tadros, in a different context, argues that these two kinds of risks are distinct ("Controlling Risk," 148–54).

13 Otsuka makes a similar argument in "Risking Life and Limb," 88.

ilar things hold for a version of Wheel. If we allow only a tiny chance that no one will be harmed, the restricted *ex post* view would allow the risk imposition since this case would no longer involve guaranteed harm. Yet if we are convinced that imposing the risk in Wheel is impermissible, it should be impermissible even in this varied scenario. We need a different version of *ex post* contractualism.

Earlier I mentioned that in Guns we only know facts about what distributions of harms are to occur with which likelihood. For example, we know that the chance that at least one Californian will die is about 63 percent. One possibility for *ex post* contractualists is to translate these facts about distributions into complaints. Imagine we specify a ranking of all persons affected. The main ranking criterion is how strong each individual complaint against the action is. In cases where individuals are equally affected, we need other tiebreaking criteria. This way we can assign each individual a unique place in the ranking. Then we repeat this for all possible outcomes. We can now construct fictional characters or "statistical persons" based on these rankings. "The worst-off Californian" refers to the first-ranked person in each of the outcomes. "The second worst-off Californian" refers to the second-ranked person and so on. In cases of objective risk imposition, these designators are nonrigid since they refer to different individuals in different possible worlds. This construction allows us to assign unique complaints to individuals instead of being limited to talking about distributions of harms. Speaking of the complaints of nonrigidly designated persons brings the *ex post* perspective closer to the theoretical core of contractualism. It can provide a model of justifiability to each that an analysis of different distributions of harms cannot offer. *Ex post* contractualists should therefore accept the following principle.

> *Ex Post Discounting*: When assessing the complaints of individuals, we should discount the complaints of nonrigidly designated individuals such as the worst-off, the second worst-off, and so on by the improbability of harm.

As mentioned, in our case of Guns, this means that the complaint of the worst-off Californian is a discounted complaint against death rather than a nondiscounted complaint as in Wheel. The complaint is discounted by the 37 percent probability that the worst-off will not be harmed. But now the second worst-off Californian has a discounted complaint as well, as has the third worst-off, and so on. Should this difference matter?

Victor Tadros believes that it should. He gives the following argument based on an example that is a simpler version of the contrast between Wheel and Guns.[14] Imagine we have two options. If we choose the first option, then it is

14 Tadros, "Controlling Risk," 153–54.

guaranteed that one and exactly one person will die. If we choose the second option, then there is only a 75 percent chance that someone will die but there is also a 25 percent chance that two persons will die. Whatever we do, the risks to each rigidly designated individual are the same. Under one view, the options are equally choiceworthy. If we choose the second option, there is a possibility that no one will die, but this is balanced by the possibility that more than one will die. Tadros, however, argues that we should choose the second option because we have special reason to prevent a situation where harm will definitely occur. We should not regard the loss of two lives as twice as hard to justify than the loss of one life. This is because the two lives are separate and not part of one aggregate that suffers a double loss.

But it is hard to see why the separateness of persons should give us a special reason to avert definite harm. Tadros's argument implies that we have less reason to prevent an additional second death. Attaching special significance to the fact that harm will occur means attaching special significance to an isolated harm as opposed to a harm that occurs alongside many other harms. Yet deaths should have the same disvalue regardless of whether they are part of an action in which only one, two, or many people die. The death is just as tragic and severe for this person regardless of how many other people have died.[15] Respect for each individual and for her separateness would seem to indicate that we should treat her loss by itself and not accord it more or less moral force because of the number of other people who have died. If this is true, then we should treat both options in Tadros's example as equally choiceworthy. The *ex post* contractualist should then regard Guns and Wheel as equally hard to justify. What should matter to us is the expected number of lives lost and not how the risk is distributed across nonrigid designators. This gives us a second principle that *ex post* contractualism would want to fulfill.

> *Equal Treatment for Equal Statistical Loss*: We should treat cases alike if in both cases there is the same expectation of statistical loss and the only difference is the distribution of possible losses across possible outcomes.

2. A PROBLEM FOR *EX POST* CONTRACTUALISM

Consider:

> *Gas*: We receive yet another option to prevent the catastrophe. This time we have to release a gas in the air that will travel to California. Scientists tell us that there is the possibility that in California the gas will react by

15 See also Otsuka, "Risking Life and Limb," 88–92.

> means of an indeterministic process with another substance and become toxic. If that happens, all Californians will die. However, they assure us that this is very unlikely. The objective probability of this occurring is only one in forty million.

In one way, Gas is a continuation of Wheel and Guns. In all three cases, each rigidly designated Californian faces an objective risk of one in forty million. The cases differ, however, in the distribution of risk across nonrigid designators. In Wheel, the distribution represents one extreme. All risk is concentrated in the likelihood of one person dying. In Guns, the distribution is spread out across all forty million nonrigid designators ranked from the worst-off to the best-off. The risks for those higher up the list are very high; for those lower down the list, they are minute. Now in Gas we face the opposite extreme. The risks are spread out perfectly evenly across all nonrigid designators. All nonrigid designators are tied, because whatever will happen, everyone in California shares the same fate. What is particularly interesting about Gas is that the distribution of discounted complaints is the same for rigid and nonrigid designators. Whether we use rigid or nonrigid designators to determine the justifiability of our action does not matter since both will yield the same result.

This is challenging for the *ex post* contractualist for the following reason: I have argued that *ex post* contractualists should accept the following two principles. They should accept *Ex Post* Discounting. This allows *ex post* contractualism to be applied to cases where harms are not guaranteed, and it provides the *ex post* perspective with a model of justifiability to each. Second, they should accept Equal Treatment for Equal Statistical Loss. This means that in Wheel and Guns what matters is the number of expected lives lost. The principle follows from accepting the claim that the disvalue of a given harm should not vary depending on how many other people will be harmed. The possibility that no person may die should be balanced by the possibility that more than one person may die.

My case Gas shows how these two principles can conflict. The number of expected lives lost in Gas is one, just like in the other two cases. If Wheel and Guns are on a par, then so is Gas. But Gas contains only heavily discounted complaints by nonrigidly designated persons. This is because the complaint of the worst-off Californian is based on only a one in forty million chance of death, a morally trivial complaint. Following *Ex Post* Discounting, it should be these discounted complaints that determine the justifiability of the risk imposition. If we want to follow Equal Treatment for Equal Statistical Loss and hold that the risk imposition in Gas is impermissible, we would need to aggregate the complaints in Gas. But whichever way we calculate the complaints, the complaints in Gas seem very

close to the complaints by the many in Transmitter Room. The complaint of Bob Johnson resembles the complaint of Jones, the worker in the transmitter room. As it turns out, the strongest version of an *ex post* view leads to a case that is very much like Transmitter Room. If we allow aggregating the complaints in Gas, then why can we not aggregate the complaints in Transmitter Room?

One proposal is that while individual and nonaggregated complaints matter, aggregative considerations can determine whether it is reasonable to reject principles.[16] Following this proposal, it is still individual complaints that matter. But their strength would be magnified by the number of people having the same complaint.

> *Ex Post Discounting (Multiplied)*: When assessing the complaints of individuals, we should discount the complaints of nonrigidly designated individuals such as the worst-off, the second worst-off, and so on by the improbability of harm. The strength of their complaint is determined by multiplying the strength of their individual complaint by the number of nonrigidly designated individuals who will be equally affected.[17]

According to this proposal, it would be unreasonable for Bob Johnson to insist on his complaint given that there are so many complaints on the other side. The strength of the individual complaint opposing Bob Johnson is magnified by the number of people who would be similarly affected. Yet Jones is equally faced with many complaints on the other side. Why should we not be allowed to multiply the individual complaint of a single football fan by the number of football fans that are equally affected? If we are allowed to magnify this individual complaint, then it would be unreasonable for Jones to reject a principle that allows the World Cup match to be broadcasted. The proposal to allow individual and nonaggregated complaints to be amplified reintroduces aggregative reasoning through the back door. So what could distinguish between Gas and Transmitter Room? Why should we understand Bob Johnson's insistence on his individual complaint as unreasonable while Jones's insistence is reasonable?

Perhaps it is the following: In Transmitter Room, the small complaints stem from mere annoyance. In Gas, the small complaints are derivative of a very serious moral claim, namely the claim not to die. This very serious claim becomes

16 This is suggested by T. M. Scanlon in his "Contractualism and Justification." Véronique Munoz-Dardé had earlier presented the idea that in some cases agents with strong complaints cannot reasonably reject principles. Munoz-Dardé invokes the idea of a threshold of reasonable demands that one can make on others. This allows for the possibility that a person with a stronger individual complaint may not be able to reasonably reject a principle ("The Distribution of Numbers and the Comprehensiveness of Reasons," 208–15).

17 I owe this proposed revised principle to an anonymous reviewer.

less important to each individual taken separately, due to the sharp discounting of their complaints by the likelihood of death occurring. Maybe Bob Johnson's insistence is unreasonable while Jones's is not because in Jones's case the opposing complaints are not complaints of the right kind. The trivial joy of watching football is not relevant to Jones's torture, while the risk of death, even if small, is relevant to Bob Johnson's lost leg. This proposal is coherent with what I wrote earlier about the opposition to aggregation. I wrote that "we should not *aggregate morally trivial complaints* so that they *outweigh serious moral complaints* of single individuals" (emphasis added). Trivial complaints should not outweigh serious complaints regardless of the numbers involved. But this leaves open that complaints of similar magnitude or qualitative significance could outweigh each other depending on the numbers.[18]

In line with the earlier distinction between the complaints of the Californians and the complaints of the World Cup viewers, we could think of complaints as being qualitatively different for different levels of actual or possible harm. Following this idea, heavily discounting a complaint against being killed does not make this complaint morally trivial. The complaint is still qualitatively on a different level than the complaint against mere annoyance. This allows us to distinguish the aggregation in Gas from the aggregation in Transmitter Room.

One problem with the idea that risks of death are qualitatively different from very small certain harms is that the same answer is available to the *ex ante* contractualist. If we stop believing that heavily discounted risks of death are morally trivial, then we could engage in a limited form of aggregation in cases like Wheel too. And then *ex ante* contractualism can account for the same answer. In other words, once we adopt the view that heavily discounted harms are not morally trivial, we lose a key motivation for adopting *ex post* contractualism.

Second, treating risks of death as qualitatively different from small certain harms fails Equal Treatment for Equal Statistical Loss in a central case. It cannot treat identified victims and statistical victims alike, even though equal respect for identified and statistical victims was one of the key motivations for *ex post* contractualism. Suppose that in a one-versus-one confrontation, a complaint against missing fifteen minutes of a World Cup match is as strong as a complaint against a risk of death of one in forty million. If we can save either one person from missing part of the match or one person from this risk of death, we should

18 The idea that complaints can only be aggregated in some circumstances is called limited aggregation. The view is suggested by Scanlon, *What We Owe to Each Other*, 238–41; and also endorsed and defended by Kamm, *Morality, Mortality*, 1:156–61, and *Intricate Ethics*, 31–40; Temkin, *Rethinking the Good*, ch. 3; and Voorhoeve, "How Should We Aggregate Competing Claims?" I set out my own view of limited aggregation in Steuwer, "Aggregation, Balancing, and Respect for the Claims of Individuals."

be indifferent. If, however, there were two people subjected to this risk of death, we should save them at the expense of the person missing parts of the World Cup match. Now what if there are many people who would be missing fifteen minutes of the World Cup match? It seems that here numbers should matter. Otherwise we would give undue importance to small risks. We should rather spare a million people from missing the World Cup match than to reduce a one in forty million risk of death to a single person. In other words, here we should be allowed to aggregate the complaints against missing parts of the World Cup match. If this is so, then we should be allowed to aggregate both the complaints against the risk of death and the complaints against missing fifteen minutes of the World Cup match. If there are many complaints against small risks, similar to my Gas case, then these might add up to one expected life lost. But since we are also allowed to aggregate the complaints of the World Cup viewers, these complaints might be decisive. However, if we contrast a single identified person with the World Cup viewers, as in Transmitter Room, we are required to save the identified person. Distinguishing between different kinds of harm cannot, therefore, treat cases where a statistical life is lost the same as cases where an identified life is lost.

Third, the idea that heavily discounted complaints against serious harm remain morally significant is also implausible in its own right. One downside of this view is that it has a problem analogous to Kamm's Sore Throat case. In Kamm's original case, we have a choice between saving one life and saving another life *and* saving someone from a sore throat. Kamm wants to say that here we should not decide in favor of saving the second person's life solely on the grounds that we can also save someone from a sore throat.[19] Now imagine that the tiebreaker is not the sore throat but the imposition of a tiny risk of death, for example, by calling an ambulance. Not only is it the case that we would be *permitted* to save the person who does not need the ambulance on the grounds that her rescue does not impose a trivial risk. But also we would be *required* to save her. It would be impermissible *not* to use the trivial risk as the deciding factor. Together with the insufficient motivation for treating equally strong complaints differently, I think this gives us grounds to treat equally strong complaints as either relevant or irrelevant. What we should accept, however, is that complaints can be aggregated when their strength is relevant to the strength of the complaints with which they are competing.

Since the *ex post* contractualist cannot distinguish between the aggregation in Gas and the aggregation in Transmitter Room, she should accept the risk imposition in Gas as permissible. She then cannot accept the principle of Equal

19 Kamm, *Morality, Mortality*, 1:146–47.

Treatment for Equal Statistical Loss. This is bad news for the *ex post* contractualist for two reasons. First, she must reject the plausible claim that harms have the same disvalue regardless of how many other people will also be harmed. The risk that one person will be harmed will receive greater weight than the risk that any additional victim over and above the first victim will be harmed. Second, a version of *ex post* contractualism that accepts the risk imposition in Gas includes a bias against statistical lives, a charge *ex post* contractualists usually raise against their *ex ante* colleagues. In some cases, like Gas, a statistical life will not be saved even though an identified life would have been. This criticism against the *ex ante* view becomes less convincing since the two theories differ only in the degree to which they are biased against statistical lives.

3. WHAT WE OWE . . . TO WHOM?

My discussion of the sequence has revealed two things. First, it has shown that two plausible principles that an *ex post* view would want to fulfill cannot be jointly fulfilled. Second, it has given us a better way of understanding *ex ante* and *ex post* views. We can understand these views as answering the question of whose complaints we should be concerned with as contractualists. Should we appeal to the complaints of identifiable individuals (epistemic *ex ante*)? Should we appeal to the complaints of rigidly designated individuals (objective *ex ante*)? Should we appeal to the complaints of nonrigidly designated individuals (*ex post*)? In what follows I will argue in favor of objective *ex ante*. The concern with the complaints of rigidly designated individuals expresses the best model of acting in ways that are justifiable to each separate person. As I explained earlier, such a concern with rigidly designated individuals means that we should draw a distinction between cases involving epistemic risk and cases involving objective risk. In a second step, I argue that this is a virtue of objective *ex ante* contractualism since it illuminates the distinction between luckless and doomed victims.

3.1. Justifiability to Each Separate Person

The core idea of contractualism is that actions must be justifiable to each. Moreover, in order to respect the separateness of persons, our actions must be *justifiable* to each as a *separate* person. This guiding idea, I argue, supports the view that our justification should address rigidly designated individuals rather than identifiable individuals or nonrigidly designated individuals. In other words, the basic idea of contractualism supports objective *ex ante* contractualism.

Consider the difference between the following three statements made by the US president after deciding on which option to take. The three statements mirror

the three options for who the ideal addressee of justification is. In each scenario, the president addresses a victim and tries to justify the imposition of the burden on her.

1. To Charlotte Williams, born on the first of June 1975, resident of Santa Barbara, who is going to die from this measure, I can only say that I am deeply sorry but your complaint against the measure was outweighed by other complaints. Even though it is hard to accept, I am convinced the measure is justifiable to you too.
2. To the Californian with the genetic predisposition, whoever he or she *may be*, I hope that you hear me. I can only say that I am deeply sorry but your complaint against the measure was outweighed by other complaints. Even though it is hard to accept, I am convinced the measure is justifiable to you too.
3. To the Californian who is going to die from the measure, whoever he or she *turns out to be*, I can only say that I am deeply sorry but your complaint against the measure was outweighed by other complaints. Even though it is hard to accept, I am convinced the measure is justifiable to you too.

Should we believe that there is an important moral difference between justification 1 and justification 2? Epistemic *ex ante* contractualists like Johann Frick believe that there ought to be. Frick, for example, holds that our ability to identify a given individual with a complaint makes a difference. Should it be impossible or overly burdensome to identify which person is going to die from the proposed policy, then we ought to treat this as a case of many discounted complaints against killing.[20] I disagree. Frick's argument relies on an idea about what we can justify to each person. But this, I think, misrepresents the core idea of contractualism. Contractualism is about *justifiability* rather than *actual justification*. Justifiability is already an idealized concept. It requires us to take into account all effects of actions on everyone concerned and to take into account all complaints everyone may have. It also requires us to take into account complaints that no one in fact has or will raise. The ideal of justifiability is one of acting in accordance with principles that would sustain a hypothetical and ideal form of justification. Since we have already idealized, it is difficult to see why we should not idealize epistemic limitations as well.

Therefore, I believe that we should think of 1 and 2 as equally good justifications. In both cases, the president is justifying her behavior to the victim. Both speeches are meant for one person alone, and address and justify the action to

20 Frick, "Contractualism and Social Risk," 193–94.

one person alone. The only difference is that speech 1 includes more detail that allows us to identify the individual. While identifiability is important for Frick, he does not discuss what is required to identify an individual. Taking a cue from Casper Hare, we can think of "identifying" an individual by knowing more personal information about that particular person.[21] We might then have identified a victim without knowing their name as long as we know enough distinctive personal information. But whether the president is able to include more detail in the description, such as name, birth date, place of residence, or other identifying information, is morally irrelevant. We are not interested in token individuals because of names or other personal information such as appearance, tastes, or talents that allow us to identify them. This information is morally superfluous. We are interested in token individuals because of their particular situation and predicament. The description "the Californian with the genetic predisposition" conveys everything that is morally important. Objective *ex ante* contractualism bases its complaints only on morally relevant information about a person's situation. This ensures that we do not confuse justifiability, which is at the heart of contractualism, with actual justification.

Even more so, at times additional information that allows us to identify individuals can even distort our moral reasoning. Imagine a doctor who has to decide on which treatment to administer to two unconscious patients, Deborah and Eric.[22] Out of expediency, the doctor has to administer the same treatment for both, even though they have two different diseases, *x* and *y*. On the one hand, the doctor can think of the prospects that Deborah and Eric have. Without any further information, the doctor would assign a fifty-fifty probability that Deborah has either of the two diseases. (And the same for Eric.) The trade-off between the two diseases will then be regarded as an intrapersonal trade-off where Deborah's and Eric's interests are the same. On the other hand, the doctor could think of the interests of "the patient with disease *x*" and "the patient with disease *y*." In this way, she would regard the trade-off as interpersonal. This way of regarding the case is superior. The doctor knows that she is dealing with an interpersonal trade-off; she knows that the interests of her two patients are not aligned. Doing one act will harm one and benefit the other. The doctor should not deceive herself into thinking that this is a choice without a conflict of interests.

Rather than between 1 and 2, we ought to hold that there is an important difference between the justifications in 2 and 3. While the contrast between 1 and 2 has shown the importance of justifiability as opposed to actual justification, the

21 Hare, "Should We Wish Well to All?" 467–71.

22 The case is a variation of one by Anna Mahtani ("The *Ex Ante* Pareto Principle," 310–11.) Mahtani credits Caspar Hare as her inspiration.

contrast between 2 and 3 shows how important it is that justifications have to be addressed to separate persons. In statement 2 (and 1), the president addresses and talks to one person alone, while in 3 the president does not address any specific person. At the time of the president's address, the words are not addressed to one individual alone. The first two speeches constitute a private channel of communication between the president and the victim. The communication and the justification are one to one. If what the president says is correct, then she would have succeeded in justifying her action to this person.

In the third speech, however, the words cannot address only one person. The justification cannot be private or one to one in the same sense. At best, the president will have addressed a person once the policy is applied, but this does not make it the case that the president did address this person prior to the action or when acting.[23] It is thus difficult to see how the justification in 3 conforms to the contractualist ideal of justifying one's action to each. Justification is owed to each separate person. But the discourse in 3 does not address persons separately. The appeal of a justification like 3 stems from the way we assimilate this thought with justifications given along the lines of my proposed speech 2. In these cases, the "someone" refers to a given individual. But this is not the case in 3. In 3, the justification addresses a compound of different individuals across different possible worlds.[24]

We can see this even more clearly when we consider cases where the complaint of the Californian who is going to die outweighs the complaint of a rigidly designated individual, such as Bob Johnson. Bob Johnson could rightly ask who the person is that can reasonably reject the proposal that would get him off the

23 The formulation here implies a rejection of the view that future contingents already have truth values. But my argument is not restricted to this metaphysical view. Some philosophers believe that future contingents already have truth values and that this view is compatible with indeterminism (see Belnap and Green, "Indeterminism and the Thin Red Line"; or Lewis, *On the Plurality of Worlds*, 206–9). If this is true, then it is the case that the president's justification does actually address one individual even though the identity depends on the objectively risky event. However, this only holds *if* the president *actually* acts this way. Should the president decide not to act this way, we have to assess a counterfactual rather than a future contingent. Under most standard views of counterfactuals, these counterfactuals will be open counterfactuals without a truth value (see Hare, "Obligations to Merely Statistical People," 380–82). This means that the model of justifiability used in 3 and whether it addresses a person will depend on what the decision maker ends up doing. But this puts the cart before the horse. An action should not be more or less justifiable based on what the agent actually does. The fact that alternative actions will be open counterfactuals also means that the model of justification used in 3 cannot be applied to help decide between different alternatives, since all but one of the alternatives include an open counterfactual.

24 See also Frick, "Contractualism and Social Risk," 196.

hook. It cannot be that we determine the identity of said person only after the fact. Even more so, *ex post* contractualism makes it impossible for us to know or determine who that person would be. It would be morally impermissible to perform the only actions that could determine the identity of this person. It will never be determined who the person was for whose sake we sacrificed Bob Johnson's leg.

Indeed, there is a compelling justification for imposing risks in cases like Wheel, even though we know one person will be harmed. Note that the victim in cases like Wheel would not have been permitted to save herself over Bob Johnson. She was facing only a small risk of death, a risk small enough that she would have been required to bear this risk. We can give the following powerful reason to the victim: you were not allowed to save yourself even accounting for your partiality toward yourself. So, you cannot complain to a third party that was not allowed to be partial toward you that she did not save you.[25]

The fact that speech 3—and thereby the model of justifiability that *ex post* contractualism employs—fails to address a particular person can also be seen clearly in a different context. By carrying the logic of speech 3 forward, *ex post* contractualism makes the permissibility of risk impositions dependent on mere population size. For this, see the following case:

> *Water (County Level)*: There is a toxic pollutant in the groundwater all over California. The pollutant will lead to every Californian losing the small finger of the right hand if nothing is done. Scientists have developed a chemical that will neutralize the pollutant. However, the chemical is still in development and thus is risky. The scientists have reduced the risk of death considerably to only one in forty million. The risks are objective and probabilistically independent for each Californian. While the pollutant affects the groundwater of all of California, the water systems are separate for each county. Each local authority has to make the decision.

Let us take as an example Santa Barbara County, which has only about 450,000 residents in contrast to the forty million residents of California as a whole. The objective risk for each individual to die is still one in forty million. But while the likelihood of at least one person dying is significant across California, the likelihood of at least one person dying in Santa Barbara County is lower. The probability is only slightly over 1 percent. Perhaps discounting the harm of death by 99 percent makes the harm less grave than the loss of the finger. (If you do not believe the harm is discounted enough, just reduce the population size further.)

25 See also Voorhoeve, "How Should We Aggregate Competing Claims?" 74.

If this is the case, then *ex post* contractualism allows releasing the chemical for Santa Barbara County. If all the other counties are of a similar or smaller size than Santa Barbara, the risk imposition would be permissible in those counties too.[26]

This leads to an absurd conclusion. *Ex post* contractualism needs to hold the following. If the government of California were to decide, releasing the chemical would be impermissible in the contractualist sense; it would not be justifiable to each. If each local government were to decide, releasing the chemical would be permissible in each case. It would be justifiable to each. Even though every single person is affected in the very same manner, the policy would turn out to be unjustifiable to one of them if the decision were taken at a different level. *Ex post* contractualism somehow generates a person with a complaint from a group of persons without a complaint. The absurdity is even clearer if we accept that unjustifiable risk impositions wrong an individual.[27] While none of the county governments would be wronging an individual if they released the chemical, the government of California would be wronging an individual. But who would be wronged? This example reveals that *ex post* contractualism fails to give us a model of acting in ways that are justifiable to *separate* persons.

3.2. The Luckless and the Doomed

Objective *ex ante* contractualism draws a distinction between cases like Dust in which the risk imposition is epistemic and cases like Wheel in which the risk imposition is objective. This is because in cases of epistemic risk, like Dust, we can identify a rigidly designated individual who is certain to be harmed while in cases of objective risk, like Wheel, we cannot. This distinction may seem suspect, and none of the other authors writing on contractualism has considered it relevant.[28] However, I believe that distinguishing between epistemically risky cases and objectively risky cases is far from being a defect of the view. To the contrary, it is a virtue of it. The reason is that this distinction tracks another distinction about the moral relevance of luckless and doomed victims. In epistemically risky cases like Dust, there is going to be one doomed victim, and in objectively risky cases like Wheel, there is going to be one luckless victim. While the effect on

26 Some counties of California are comparatively large, for example, Los Angeles County with over 10 million people. We can imagine that in those counties more local authorities have to make the decision.

27 See, e.g., Oberdiek, *Imposing Risk*, 126–53. Frances Kamm has argued for the more radical claim that Scanlon's account for wrongness should generally be understood as an account of wronging (*Intricate Ethics*, 461–68).

28 Indeed, Frick argues against its relevance in "Contractualism and Social Risk," 197–201.

both is the same, we can see that there is a significant difference between having doomed a person who ends up dying and having given that person a very favorable chance of survival.

John Broome in his discussion of fairness makes the following remark about persons who lose out in the allocation of a scarce good.[29] Whoever loses out has grounds for complaint. But the person would have an even bigger ground for complaint if it were never even in the cards for her to have received the good. We cannot justify our allocation to this person by saying that we gave her a fair shot at receiving the good. Losing out for this person is not tough luck but, worse, an inevitable feature of our decision. The fact that she might have won, that it was once in the cards for her to win, mitigates her complaint against missing out. In short, after the allocation, a luckless loser has a less strong complaint than someone who was doomed to lose. The lottery example shows how the kind of risk that is at play in allocating the good matters for the complaints that individuals can raise. In a lottery that employs epistemic risks, it was never in the cards for anyone other than the winner to win. In an objectively risky lottery, this is not the case. Every person stood a chance of getting the good. The lottery is fair because it is the luck of the draw that decides who gets it.[30] Objectively risky lotteries are such that we can say to the person that she could have received the good. We designed the lottery such that it could have easily gone the other way and she might have won.[31]

These points about fairness in allocating goods are not limited to the allocation of benefits. They should also apply to the allocation of burdens or harms. Common examples to illustrate lottery fairness include such cases. The draft lottery to select soldiers for the Vietnam War is a paradigm example. The cases I have discussed are similar. In all cases, harms are avoidable only at the expense of a moral catastrophe. We have to decide about the allocation of harm. This means that we can say to those who are luckless that they could have avoided the harm, whereas those who would have been doomed would not have had any

29 Broome, "Fairness," 98.

30 This idea is even invoked by critics who account for lottery fairness in a different manner. George Sher and Michael Otsuka give accounts of lottery fairness of merely epistemic lotteries since both doubt that lotteries with objective risks exist. Sher mentions the "luck of the draw" interpretation as the most obvious rationale for lottery fairness, but adds that it is incomplete because it cannot account for the fairness of lotteries that do not employ objective risks. Otsuka argues that objectively risky lotteries would be fairer than epistemically risky lotteries, if it were possible to run them. Sher, "What Makes a Lottery Fair?" 203–4; Otsuka, "Determinism and the Value and Fairness of Equal Chances."

31 I owe this point to Kai Spiekermann. He explores this idea in connection to lottery fairness and social risk in "Good Reasons for Losers."

such chance. It is a virtue of objective *ex ante* contractualism that it can distinguish in this manner between luckless and doomed victims.

While the previous considerations on fairness illustrate the importance of the distinction between the luckless and the doomed in giving reasons after the risk materializes, there are also reasons to care about the distinction before the action. Consider the following case narrated by Anatol Rapoport.[32] In the Second World War, an allied air base in the South Pacific faced the problem that most of their planes did not survive their allocated missions. The chance of survival was only one in four. An alternative but rejected policy would have increased the chances of survival. Only half of the planes would fly missions with increased bomb load. The increased load would mean that less fuel would be available, and the pilots could not return to safety and would crash. Instead of giving everyone a chance of one in four, the policy would fate half the pilots to certain death. The repulsion against and failure to adopt the policy is best explained by an objection against dooming individuals to death.[33]

However, the difference between doomed and luckless victims goes beyond cases where the victims know their fate. Assume a small variation of this case where, in order to ensure compliance with the order to fly, after the selection by lot all pilots board a plane. Neither commanders nor pilots know which planes are loaded and which carry empty loads. Pilots who fly an empty plane have orders to return to a different base when they realize their plane is empty at the first target. At the decision to order the pilots to fly, every pilot faces an epistemic risk of death of 50 percent. This variation is no less objectionable than the initial plan. By distinguishing between doomed and luckless victims, objective *ex ante* contractualism can account for this. The doomed pilots are certain to die whereas under the ordinary protocol all pilots face a three-quarters objective risk of death. By contrast, epistemic *ex ante* contractualism may justify the order to fly given that it reduces the epistemic risk each pilot faces. *Ex post* contractualism in turn would justify the order to fly given that it reduces the number of expected lives lost. Only objective *ex ante* contractualism can account for the answer that is both the actual decision at the base and the intuitively correct one.

One might object to my analysis of the case of the pilots. Assuming that the selection by lot were random, every pilot would have faced a 50 percent objective risk of death under the alternative policy as opposed to a 75 percent objective risk of death under the standard policy. However, it is not accurate to draw

32 Rapoport, *Strategy and Conscience*, 88–90. Rapoport presents this case as a real-life case but could not vouch for its authenticity.

33 Jonathan Glover reports that the horror of certain death motivates the refusal to accept the policy of one-way missions in Rapoport's example (*Causing Death and Saving Lives*, 212–13).

the conclusion that objective *ex ante* contractualism would therefore endorse the alternative policy. The problem here is similar to the problem of medical experimentation discussed by Frick. In the example of medical experimentation, there is an *ex ante* selection of persons to be experimented on. At the stage of selection, the policy of experimenting is beneficial to all, but after the selection is made, severe hardship is imposed on some. Objective *ex ante* contractualism can avail itself of the same reply as epistemic *ex ante* contractualism and adopt what Frick calls the Decomposition Test.[34] The Decomposition Test imposes a requirement to always act, in each action, in ways that are justifiable to each. The policy of selecting people at random first and then imposing severe hardships on them does not meet this test. This holds for the case of medical experimentation as well as for the case of the pilots. When sending out the pilots, some pilots are doomed to certain death. Objective *ex ante* contractualism prohibits this.[35]

Our objection to dooming the pilots to certain death is linked with our intuitions about risk concentration and risk dispersal. Take, for example, our reaction to a now debunked story about the Coventry Blitz, the horrendous bombing raid of Nazi aircrafts on the city of Coventry. According to the story, Churchill knew about the impending devastating attack on Coventry and could have averted it. To avoid revealing military intelligence, Churchill sacrificed Coventry for the sake of the overall war effort and reduction of the overall death toll. When the story was published, it was perceived as a grave accusation and moral flaw for Churchill to have acted this way.[36] Distinguishing between doomed victims in Coventry and unlucky victims elsewhere in the United Kingdom can explain why. Rapoport's pilot case and the Coventry Blitz reveal that our intuitions about concentrating and dispersing risks are sensitive to what kind of risk we are talking about. The plan to fly one-way missions disperses and reduces epistemic risks, but this does not make the plan very appealing given that objective risks are concentrated. There is little point in dispersing epistemic risks if we know

34 Frick, "Contractualism and Social Risk," 201–12.

35 Nir Eyal has suggested that what is problematic about Rapoport's case is not that the pilots are doomed, but rather that they are doomed by their commanders. The commanders, as opposed to enemy fire, would be killing the pilots by adopting the policy. See Eyal, "Concentrated Risk, the Coventry Blitz, Chamberlain's Cancer," 105–7. However, I believe that this part of the story is not central. My reaction would not change if some of the planes had insufficient fuel due to sabotage and the commanders had the choice of aborting the mission and calling the planes back. (Imagine that bombs are loaded automatically according to overall weight.) The commanders would still doom some pilots to certain death, even if the pilots would not be killed by the commanders.

36 See Eyal, "Concentrated Risk, the Coventry Blitz, Chamberlain's Cancer," 94–95. Eyal seeks to vindicate Churchill's imagined reasoning.

that it is already carved in stone who will die. However, dispersing objective risks is a genuine sense in which burdens are shared and additional burdens are spread more widely.

Thus far I have argued that part of the reason why the distinction between objective and epistemic risks is meaningful is because it can explain the moral difference between luckless and doomed victims. This allows me to respond to one concern about my view. Imagine a vaccine that we know carries a certain small risk of serious harm. Whether the foreseen harms of mass vaccination are a reason against the mass vaccination will depend, on my view, on the specific mechanism by which the risk manifests itself. If the mechanism is a random mutation, then it is a small objective risk, whereas if the mechanism relies on genetic predispositions, then it is a small epistemic risk but a large objective risk. Why should this mechanism matter? In response: the mechanism matters because in the case of the random mutation, the harmed victim is luckless, whereas in the case of the genetic predisposition, we would doom the victim to be harmed. As I have argued, there is an important moral difference between being luckless and being doomed, and this moral difference makes the otherwise uninteresting-seeming difference in the biological mechanism of the vaccine relevant. While we often do not know with certainty what mechanism applies, we often have information about whether our applied case is more like the case of random mutations or more like the case of genetic predispositions. This, I believe, rightly influences how we ought to act in the case.

The distinction between objective and epistemic risks is also important for another reason. It can illuminate the importance of hypothetical consent. An important and familiar reason for rejecting *ex post* contractualism is that it makes actions impermissible even if these actions would receive the hypothetical consent of all affected parties. For each individual, it is sometimes rational to take small risks of death for moderate gains. For example, it would be rational to take a vaccine against a disease that is not life threatening even if there is a risk of a lethal allergic reaction. If such risks are imposed on a large scale, then we can be virtually certain that some person will die from the risk. Not only are these risk impositions intuitively permissible, but we can give a strong argument in favor of them. Frick has called this the Argument from the Single Person Case.[37] If the risk imposition were to affect only a single person, it would be permissible. In such a case, it seems reasonable that we should do what is in that person's rational self-interest. Now in a second step, we learn that there is a second person in

37 Frick, "Uncertainty and Justifiability to Each Person," 133–34; and Frick, "Contractualism and Social Risk," 186–88. Similar arguments are made by Tom Dougherty ("Aggregation, Beneficence, and Chance," 8–11) and Caspar Hare ("Should We Wish Well to All?" 455–67).

an identical position as the original person. The risky treatment is available at no additional cost for that person too. The case is still relevantly similar to deciding for one person. It does not involve any competing claims. We can add more and more people. Individually, we would always favor giving them the treatment. Yet *ex post* contractualism needs to hold that for a sufficiently large group the risk imposition becomes impermissible.

Is there anything the *ex post* contractualist could say to reject the Argument from the Single Person Case? The best response seems to be the following. The hypothetical consent that each person would give is vitiated because each person is imperfectly informed.[38] If we knew that a person would only consent because she was insufficiently informed, it is less plausible to assign moral weight to this hypothetical consent. Imagine that you are a guardian charged with that person's interest. If you were fully informed and knew that the risk imposition was in that person's interest only because of imperfect information, you would not assign moral importance to that fact about self-interest. A close variation of this case is a case where you are in charge of various persons' interests. You may not know which person is going to lose out, but you still know the related fact that one of the persons whose interests you look after is going to lose out. As a fully informed guardian, you would therefore object to the action. In epistemically risky cases like the vaccine case, this is the case. Somewhere in the chain, there is a person for whom it is not in their fully informed self-interest that the risk will be imposed. The chain of single person cases is no longer fully symmetrical under conditions of full information. Since we can anticipate this, we have grounds to object to the risk imposition.

The reply to the Argument from the Single Person Case helps us refine the importance of hypothetical consent. Unlike with actual consent, we have no reason to give moral significance to hypothetical consent that arises due to imperfect information. Yet this challenge does not impede giving significance to hypothetical consent that is not tainted in this manner. Such untainted hypothetical consent is at stake in objectively risky cases. Remember the Water case introduced earlier. In Water, every Californian faces the same problem for deliberation. Either they will lose their small finger or they will incur a minute risk of death. The gamble is in the self-interest of each Californian; each would hypothetically consent. In this case, the response that hypothetical consent arises only out of imperfect information has no bite. Even if all Californians knew all relevant facts about themselves, it would nonetheless be in their self-interest to take the gamble. The Argument from the Single Person Case stands. Distinguishing between epistemic and objective risks helps us understand that the Argument from the

38 See Fleurbaey and Voorhoeve, "Decide As You Would with Full Information!"

Single Person Case is compelling in some cases while unconvincing in others. By distinguishing between these cases, objective *ex ante* contractualism retains what is attractive in the Argument from the Single Person Case while avoiding the charge that hypothetical consent is vitiated due to imperfect information. In the revised case, all risk impositions are independent from one another. There is no conflict over the resource that gives everyone a favorable prospect for their lives. Since there is no connection between the risks, there is no reason why it should not be permissible to impose all of them at once. Consequently, objective *ex ante* allows us to impose all of them at once.

4. OBJECTIONS

I will consider two main lines of objection to my version of *ex ante* contractualism that discounts objective, rather than epistemic, risk. The first line of objection stems from the possibility that determinism is true. The second line of objection criticizes an identified victim bias in my position.

4.1. Determinism

My view distinguishes between objective risks and epistemic risks. There is a worry that even if this distinction would be of moral importance, it is irrelevant in the real world. If determinism is true, the worry goes, then there is no such thing as objective risk. There might be actually observed frequencies but no objective risk in a robust sense that could be morally relevant. The view that the truth of determinism implies the absence of objective chances was once taken as the orthodox view in the philosophy of probability. Recently, however, there has emerged a growing literature in the philosophy of probability that argues that objective chance or objective probability is compatible with determinism.[39]

A first reason to think that objective probabilities are compatible with determinism stems from the existence of probabilistic laws in science. To give some examples, classical statistical mechanics, evolutionary theory, Mendelian genetics, meteorology, and the social sciences all include probabilistic laws. In fact, it appears that deterministic laws are largely confined to just one branch of science, namely the physical sciences. The probabilities posited by the laws of the special sciences, including parts of the physical sciences like classical statis-

39 See Loewer, "Determinism and Chance"; Hoefer, "The Third Way on Objective Probability"; Glynn, "Deterministic Chance"; Eagle, "Deterministic Chance"; Lyon, "Deterministic Probability"; Strevens, "Probability out of Determinism"; Emery, "Chance, Possibility, and Explanation"; Frigg and Hoefer, "The Best Humean System for Statistical Mechanics"; List and Pivato, "Emergent Chance."

tical mechanics, do not appear to be epistemic. For example, the process of ice cubes melting when being put in water is a probabilistic process according to classical statistical mechanics. It appears that classical statistical mechanics can, by virtue of this probabilistic law, *explain* why the ice cube is melting. Indeed, if we believe that special sciences above the microphysical level are able to explain phenomena, then they explain these phenomena by reference to probabilistic laws. This makes it difficult to conceive of such laws as being concerned with epistemic probabilities. The laws of classical statistical mechanics cannot both incorporate our ignorance about deterministic processes and at the same time explain why ice cubes are melting or why the climate system is changing. Our ignorance cannot explain.

So how can we accommodate both the fact that laws of the special sciences posit objective chances and the idea that the universe is deterministic at the microphysical level? One rationale for the compatibility of objective chance and determinism at the microphysical level is that the descriptions of "chance" and "determinism" are level specific.[40] It is imprecise to talk about whether or not the world is deterministic. The real question is whether or not the world is deterministic *at a specific level*. A helpful test to see whether the world is deterministic at a given level is to ask whether knowing the entire history of the world described at that level determines a future event. Those who argue that the world is deterministic at the microphysical level mean to say the following: if we knew all the laws of nature as well as the initial conditions of the universe described in microphysical language, then the only chances of an event happening are zero or one. But this does not say anything about whether or not the world is deterministic at some macrolevel. It does not follow that, at the macrolevel, the history of the world already determines the event. In other words, determinism at the microphysical level can coexist with indeterminism at some macrolevel. This way, macrolevel events like melting ice cubes or coin tosses will have their own macrolevel chances.

For the purposes of moral theorizing, we are predominantly concerned with the agential level, the level at which we describe agents and their actions. The agential level is the appropriate level for the moral decision-making of agents. What would rule out the possibility of objective chances in the relevant sense is, therefore, not determinism at the microphysical level but rather *determinism at the agential level*. Yet there is no reason to think that our world is deterministic at the agential level. To the contrary, all indications of our best available (social) science at the agential level tell us that the world is *indeterministic* at the agential level. Even if we knew the entire history of the universe described at the level

40 Glynn, "Deterministic Chance"; List and Pivato, "Emergent Chance."

of agents and macro-objects, like coins, together with all laws of human behavior, we would not be able to predict, say, the outcome of the next presidential election. Arguments for determinism rely on information about microphysical particles and their properties, something that is inadmissible when thinking about whether the world is deterministic at a higher level. The level-specific approach to determinism and chance retains the ability to draw a distinction between objective chance and epistemic credence at each level of description.[41] Imagine an agent is about to toss a fair coin. The odds of the coin landing heads are 0.5. These are objective chances since the prior history of the world, at the level of coin tosses, does not determine this event. After the coin toss, the agent is covering the coin and asks again what the odds are of the coin having landed heads. The answer would seem to be 0.5. But this statement about probabilities is clearly different from the earlier one. The second odds are credences, the first are chances. Thus, the level-specific view can retain the distinction between chances and credences at every level. This distinction in turn means that while agents can create objective chances, they can also create merely epistemic risks. A lottery based on whose birthday is earliest in the year would create epistemic risks if the birthdays of participants are unknown, but it would not create objective risks for the participants.

We can see the point of the level-specific view in another way. Consider again the coin flip. Assume that we hold all other factors constant except for the force exerted on the coin. The following conditionals might all be true:

> If I flip the coin with a force between 0.18345 and 0.18348 N, it will land heads.
>
> If I flip the coin with a force between 0.18349 and 0.18352 N, it will land tails.
>
> If I flip the coin with a force between 0.18353 and 0.18356 N, it will land heads.

And so on. But what about the conditional "If I flip the coin, it will land heads"? Or the conditional "If I flip the coin, it will land tails"? The antecedents of these conditionals are underspecified. They do not tell us with which force the coin is flipped, and the deterministic laws of physics tell us that small changes in the force applied to the coin lead to different outcomes. The antecedent of the underspecified conditionals describes a set of possible worlds. In this set, there are some possible worlds where the coin lands heads and some possible worlds where the coin lands tails. What we can give for the underspecified conditional

41 See List and Pivato, "Emergent Chance," 139–42.

is a probability of how many worlds are head-landing worlds.[42] The fact that this probability is not merely epistemic can be seen if we consider the case in which the conditional is a counterfactual conditional. Processes like this coin flip are counterfactually open. No head-landing world is relevantly more similar to our actual world than any tail-landing world. Since the process is counterfactually open, there will not be a fact of the matter about what would have happened had we flipped the coin. There would only be a counterfactual probability. Since there is no fact of the matter about what would have happened, this probability cannot be interpreted as referring to our ignorance about what would have happened.

Now why should we be interested in underspecified conditionals as opposed to fully specified conditionals? After all, in a conditional that is specified at the microphysical level, there are no nontrivial probabilities if we assume determinism at the microphysical level. The reason is the link between contractualism and evidence-based criteria of rightness. Risk impositions are only an issue for contractualism if it is interpreted as an evidence-based criterion of rightness. If contractualism is interpreted as a fact-based criterion of rightness, a risk imposition would be wrong if and only if it leads to eventual harm. But a fact-based criterion is unhelpful in guiding the choices of agents. Evidence-based criteria, on the other hand, link moral permissibility to a choice an agent can make. They capture morality as answering deliberative questions for agents. The actions that contractualism is concerned with are therefore those that are in the choice set of an agent.[43] As agents, we are unable to choose the option "flip the coin with a force between 0.18345 and 0.18348 N." This is simply not an option available to us. The option that is available to us is an option at the agential level, namely "flip the coin." This gives us an argument for specifying conditionals at the agential level. The agential level captures the options that are available, open to the agent, whereas a microphysical level does not.

The argument for the compatibility of lower-level determinism and objective chances has another upshot. A perennial challenge to *ex post* contractualism is that it prohibits many intuitively permissible forms of risk imposition where small risks are imposed on large populations. It would seem that traffic victims have reason to reject principles that allow higher speed limits. Starting major construction works would be impermissible because of the risk of harm to workers. Air traffic may be difficult to justify because it leads to harms to bystanders.

42 See also Hare, "Obligation and Regret When There Is No Fact of the Matter about What Would Have Happened If You Had Not Done What You Did," 190–94; and Hare, "Obligations to Merely Statistical People," 380–82.

43 Scanlon, *Moral Dimensions*, 56–62.

The list goes on.[44] What these divergent risks all have in common is that they appear random in a relevant sense. They contrast with, for example, the risk of a lethal allergic reaction in an individual. Such an individual's death may have been difficult to prevent, but it is not random in the same sense. The aforementioned examples all appear random because none of these events is determined by the previous history of the world at the agential level. The event "person is killed in car accident" is not already determined by the past history of the world. At most, a description of the event in microphysical language is determined. This means that at the agential level, the level that counts, all the familiar examples are objectively risky. Therefore, objective *ex ante* contractualism can appealingly explain why it is permissible to impose such risks.

4.2. Identified Victim Bias

The second objection arises from the discussion concerning identified and statistical lives. *Ex ante* contractualism generally favors a bias toward identified lives and has received criticism for giving too strong an endorsement to saving identified lives over statistical lives. While this observation is broadly correct, the relationship between my version of *ex ante* contractualism and the problem of identified and statistical lives is more complex. Objective *ex ante* contractualism does not place any emphasis on the victim being identified. Rather, what is relevant is whether the victim is already determined. In a case like Dust, we do not have a way to identify the victim, but given that we have a rigid designator for the victim, we should favor her.

Indeed, my proposal can at times account for saving a statistical life rather than an identified life. For this, see a simplified version of a case by Caspar Hare.[45] You have two options: either you head north or you head south. If you head north, you will save one person for certain. If you head south, you can flip an indeterministic coin. If it lands heads, you will save another person. If it lands tails, you will save yet another person. The two potential southern victims can complain that if you head north they will die. You deprived them of a 50 percent chance to live. They can also complain that you would allocate chances to live more unequally if you were to head north. The potential northern victim can complain that by heading south you deprived her of a 100 percent chance to live. The northern victim cannot raise an additional complaint about the unfairness of the unequal distribution of chances. If we accept limited aggregation, then it seems plausible that a complaint against a 50 percent chance of death is close

44 See Norcross, "Comparing Harms," 159–67; Ashford, "The Demandingness of Scanlon's Contractualism," 298–99; James, "Contractualism's (Not So) Slippery Slope," 268–72.

45 Hare, "Obligations to Merely Statistical People," 382, 385.

enough to a complaint against a 100 percent chance of death. If this is correct, and we are permitted to aggregate the claims of the southern victims, then the added complaints against unfairness would tip the balance. It would follow, on my view, that you ought to head south and save the statistical, rather than the identified, life.

Nevertheless, the general observation is correct. *Ex ante* contractualism retains a bias against statistical lives, even though this bias is substantially weakened due to the permissibility of limited aggregation. Take, for example, the following revision of Wheel: the indeterministic roulette wheel does not release one ball but ten balls that will kill ten different persons. To many, it is difficult to accept that we should prioritize Bob Johnson's leg over multiple statistical victims. However, we should note that the individual risk for each person, while higher than in the standard version of Wheel, is still vanishingly low at one in four million.

On reflection we notice that small risks of serious harms are omnipresent. It is inevitable that large-scale policies will lead to serious harms. In many such cases of social risk, we nonetheless believe that the risk imposition is permissible. Indeed, accounting for these cases is a key challenge to *ex post* contractualism. Take, for example, the following stylized case:

> *Vaccine*: In order to protect the entire population of California from an infectious disease, which everyone would come down with in the absence of any intervention, the government is considering a mass vaccination program. The disease is not life threatening but would cause the Californians to limp for two months, similar to the effects of a sprained ankle. While the temporary limp is much less bad than the impairment due to loss of a leg, it is significant enough that the Californians want to avoid it. In extraordinary circumstances, the vaccine can, however, be lethal, although the chance of death for each Californian is only one in four million. The government is able to administer the vaccine without intrusion on the bodies of any Californian.

Even though the policy in Vaccine will also lead to ten expected statistical deaths, we want to account for the permissibility of Vaccine. The risk of death is sufficiently small that it is outweighed by the benefit of avoiding the temporary limp. For example, according to the National Safety Council, the odds of a US resident being struck by lightning in their lifetime are a bit over one in 180,000, more than twenty-two times more likely than the harm due to the vaccine.[46] Rejecting risks of the kind involved in Vaccine would make it difficult to pursue many large-

46 See the overview at National Safety Council, "Odds of Dying."

scale policies or practices. The challenge is now the following. In the case of Vaccine, we prefer saving the population of California from the temporary limp over the loss of ten statistical lives. In the revised Wheel case, we prefer saving the ten statistical lives over Bob Johnson's loss of a limb. Now what if we could choose between saving the population of California from the temporary limp or saving Bob Johnson from the loss of a leg? Since the temporary limp is much less bad than the permanent loss of a leg, it is plausible that a contractualist would reject the aggregation of the complaints against the temporary limp. Hence, we should save Bob Johnson. This leads us to a preference cycle over the three options.

It is not clear how we could justify such a preference cycle. One attempt would be to point out that in Vaccine the gamble is in the *ex ante* interest of all, whereas this is not the case in the revised Wheel case.[47] This may explain why the option of "ten statistical victims when it was in their *ex ante* interest to take the risk" is not the same option as "ten statistical victims." I am not convinced that this explains our intuitions well. While it is true that the gamble is in the *ex ante* interest of all in the stylized Vaccine case, I do not believe that this is necessary to the case. I believe that delivering the vaccine would be permissible even if some small and unidentifiable part of the population were already known to be immunoresistant. The vaccine would, therefore, be neither to the *ex ante* nor the *ex post* benefit of any of them. In fact, it appears that in most cases of intuitively permissible large-scale risks, the benefits are widespread but not universal.

What the response shows, however, is that it is a mistake to frame the problem in the revised Wheel case as either saving ten people from death or saving one person from the loss of a leg. Such a framing already assumes that what matters is the harm that is the result of the risk imposition. In other words, this framing already assumes the *ex post* perspective. If my arguments against the *ex post* perspective are successful, then we should rather phrase this choice as saving the leg of one and reducing the risks of very many by a small amount. So understood, it is more plausible to maintain that it is permissible to impose the risk in the revised Wheel case.

We can give the following justification for our choice. At the time of our decision, there was no person who had as strong of a complaint as Bob Johnson did. We were able to justify our action to each of the forty million persons involved, each of whom faced only a very small risk of death. In fact, none of the forty million would have been permitted to save themselves from such a small risk if doing so had required the loss of Bob Johnson's leg. For example, each would have been required to call an ambulance to save Bob Johnson's leg even if this would have created a one in four million chance of being killed by an ambu-

47 See Walen, "Risks and Weak Aggregation."

lance sliding out of control. We can acknowledge that a better outcome could have been brought about, in which only one person loses a limb rather than ten people losing a life. But that is the sort of thing nonconsequentialists are already willing to acknowledge across a range of familiar cases. Nonconsequentialists accept that oftentimes it is impermissible to do what brings about the best outcome because doing so would violate the claims of a single individual. We can understand deontological constraints in this way.

In line with the analogy to deontological constraints, we can accept a further claim. While nonconsequentialists accept some inefficiency in terms of failing to bring about the best outcome, they typically accept that there are some cases in which deontological constraints can be overridden. Most nonconsequentialists believe that rights may permissibly be violated in cases where doing so is necessary to avoid a moral catastrophe or some other high threshold of weighty moral considerations. In those cases, even deontological constraints such as those that stand in the way of being harmfully used as a mere means can be exceptionally suspended.[48] In such cases, it can be permissible to do what would otherwise be unjustifiable to the rights holder—for example, violating the right not to be harmed as a mere means. If it is plausible that we can override the individual complaint not to be used as a mere means, then it also seems plausible that we can sometimes override the individual complaint of a determined victim against not being saved. If anything, the complaint against being used as a mere means appears to be a stronger complaint than the complaint against failing to be saved in the cases under discussion in this article.

The analogy is strengthened by a deep theoretical connection that contractualism has with a rights-based morality. Contractualism only covers a part of morality, the part that Scanlon identifies with "what we owe to each other." This part is concerned with our relations to other persons. A natural thought is that when we act in ways that are not justifiable to a given person, we thereby wrong this person. Similarly, when we violate the right of a person, we thereby wrong this person. This suggests an important theoretical connection between contractualism and a rights-based morality, given that both are concerned with wrongs done to other persons.[49] Therefore, the idea that there is some threshold of statistical victims at which point we need to depart from contractualist morality is no more problematic than the widely accepted idea that there is some threshold of bad consequences at which point we need to depart from deontological constraints.

48 See, e.g., Nagel, *Mortal Questions*, ch. 5; and Thomson, *The Realm of Rights*, ch. 6.

49 See, e.g., Kamm, *Intricate Ethics*, 461–68.

5. CONCLUSION

In this article, I have argued for a new version of *ex ante* contractualism that focuses on the complaints that rigidly designated individuals can bring forward. Their complaints ought to be discounted by the objective probability that the harm will come about. Unlike other *ex ante* contractualists, I do not believe that we should always discount epistemic risk, nor do I believe that we should be concerned only with individuals that we can identify. Such an objective version of *ex ante* contractualism provides us with a plausible model of justifiability to each. It insists that our actions must be justifiable to everyone at the time that we act. It also insists that justification is owed to separate persons. But it does not require the use of morally superfluous identifying information that would make actual justification to each possible. Objective *ex ante* contractualism is alone in drawing a distinction between cases in which objective risks are at stake and cases in which merely epistemic risks are at stake. But far from being a defect, this is a virtue. We can thereby illuminate the morally relevant difference between luckless and doomed victims. For these reasons, I conclude that objective *ex ante* contractualism is a viable and better alternative that is theoretically superior to both epistemic *ex ante* and *ex post* contractualism.[50]

Rutgers University
steuwer@cplb.rutgers.edu

REFERENCES

Ashford, Elizabeth. "The Demandingness of Scanlon's Contractualism." *Ethics* 113, no. 2 (January 2003): 273–302.

Belnap, Nuel, and Mitchell Green. "Indeterminism and the Thin Red Line." *Philosophical Perspectives* 8 (1994): 365–88.

50 Versions of this article were presented at the LSE Choice Group, the LSE/UCL Risk and Aggregation Workshop, the TIF Conference at the University of Barcelona, the London Graduate Moral and Political Philosophy Workshop at King's College London, and a research seminar at the London School of Economics. I am grateful for helpful comments and suggestions by members of the audiences. I also thank Susanne Burri, David Coombs, Nicolas Côté, Christina Easton, Johann Frick, Joe Horton, Christian List, Anna Mahtani, Chris Marshall, Todd Karhu, Thulasi K. Raj, Tom Rowe, Korbinian Rüger, and multiple anonymous reviewers for very helpful discussions and comments on various drafts. I am especially grateful to Michael Otsuka and Alex Voorhoeve for extensive discussion and comments on countless drafts.

Broome, John. "Fairness." *Proceedings of the Aristotelian Society* 91 (1991): 87–101.
Carnap, Rudolf. "The Two Concepts of Probability." *Philosophy and Phenomenological Research* 5, no. 4 (June 1945): 513–32.
Cohen, I. Glenn, Norman Daniels, and Nir Eyal, eds. *Identified versus Statistical Lives: An Interdisciplinary Perspective.* Oxford: Oxford University Press, 2015.
Daniels, Norman. "Can There Be Moral Force to Favoring an Identified over a Statistical Life?" In Cohen, Daniels, and Eyal, *Identified versus Statistical Lives*, 110–23.
Dougherty, Tom. "Aggregation, Beneficence, and Chance." *Journal of Ethics and Social Philosophy* 7, no. 2 (May 2013): 1–19.
Eagle, Anthony. "Chance versus Randomness." *Stanford Encyclopedia of Philosophy* (Spring 2019). https://plato.stanford.edu/archives/spr2019/entries/chance-randomness/.
———. "Deterministic Chance." *Noûs* 45, no. 2 (June 2011): 269–99.
Emery, Nina. "Chance, Possibility, and Explanation." *British Journal for the Philosophy of Science* 66, no. 1 (March 2015): 95–120.
Eyal, Nir. "Concentrated Risk, the Coventry Blitz, Chamberlain's Cancer." In Cohen, Daniels, and Eyal, *Identified versus Statistical Lives*, 94–109.
Fleurbaey, Marc, and Alex Voorhoeve. "Decide As You Would with Full Information! An Argument against *Ex Ante* Pareto." In *Inequalities in Health: Concepts, Measures, and Ethics*, edited by Nir Eyal, Samia A. Hurst, Ole F. Norheim, and Dan Wikler, 113–28. Oxford: Oxford University Press, 2013.
Frick, Johann. "Contractualism and Social Risk." *Philosophy and Public Affairs* 43, no. 3 (Summer 2015): 175–223.
———. "Uncertainty and Justifiability to Each Person: Response to Fleurbaey and Voorhoeve." In *Inequalities in Health: Concepts, Measures, and Ethics*, edited by Nir Eyal, Samia A. Hurst, Ole F. Norheim, and Dan Wikler, 129–46. Oxford: Oxford University Press, 2013.
Fried, Barbara H. "Can Contractualism Save Us from Aggregation?" *Journal of Ethics* 16, no. 1 (March 2012): 39–66.
Frigg, Roman, and Carl Hoefer. "The Best Humean System for Statistical Mechanics." *Erkenntnis* 80, no. 3 (December 2015): 551–74.
Gauthier, David. *Practical Reasoning*. Oxford: Clarendon Press, 1963.
Glover, Jonathan. *Causing Death and Saving Lives*. London: Penguin Books, 1977.
Glynn, Luke. "Deterministic Chance." *British Journal for the Philosophy of Science* 61, no. 1 (March 2010): 51–80.
Hájek, Alan. "Interpretations of Probability." *Stanford Encyclopedia of Philosophy* (Fall 2019). https://plato.stanford.edu/archives/fall2019/entries/probability-interpret/.

Hare, Caspar. "Obligation and Regret When There Is No Fact of the Matter about What Would Have Happened If You Had Not Done What You Did." *Noûs* 45, no. 1 (March 2011): 190–206.

———. "Obligations to Merely Statistical People." *Journal of Philosophy* 109, no. 5/6 (May/June 2012): 378–90.

———. "Should We Wish Well to All?" *Philosophical Review* 125, no. 4 (October: 2016): 451–72.

Hoefer, Carl. "The Third Way on Objective Probability: A Sceptic's Guide to Objective Chance." *Mind* 116, no. 463 (July 2007): 549–96.

Horton, Joe. "Aggregation, Complaints, and Risk." *Philosophy and Public Affairs* 45, no. 1 (Winter 2017): 54–81.

James, Aaron. "Contractualism's (Not So) Slippery Slope." *Legal Theory* 18, no. 3 (September 2012): 263–92.

John, S. D. "Risk, Contractualism, and Rose's 'Prevention Paradox.'" *Social Theory and Practice* 40, no. 1 (January 2014): 28–50.

Kamm, F. M. *Intricate Ethics*. Oxford: Oxford University Press, 2007.

———. *Morality, Mortality*. Vol. 1, *Death and Whom to Save from It*. New York: Oxford University Press, 1993.

Kumar, Rahul. "Risking and Wronging." *Philosophy and Public Affairs* 43, no. 1 (Winter 2015): 27–51.

Lenman, James. "Contractualism and Risk Imposition." *Politics, Philosophy and Economics* 7, no. 1 (February 2008): 99–122.

Lewis, David. *On the Plurality of Worlds*. Oxford: Basil Blackwell, 1986.

———. "A Subjectivist's Guide to Objective Chance." In *Studies in Inductive Logic and Probability*, vol. 2, edited by Richard C. Jeffrey, 263–93. Berkeley: University of California Press, 1980.

List, Christian, and Marcus Pivato. "Emergent Chance." *Philosophical Review* 124, no. 1 (January 2015): 119–52.

Loewer, Barry. "Determinism and Chance." *Studies in History and Philosophy of Modern Physics* 32, no. 4 (December 2001): 609–20.

Lyon, Aidan. "Deterministic Probability: Neither Chance nor Credence." *Synthese* 182, no. 3 (October 2011): 413–32.

Mahtani, Anna. "The *Ex Ante* Pareto Principle." *Journal of Philosophy* 114, no. 6 (June 2017): 303–23.

Munoz-Dardé, Véronique. "The Distribution of Numbers and the Comprehensiveness of Reasons." *Proceedings of the Aristotelian Society* 105, no. 1 (June 2005): 191–217.

Nagel, Thomas. *Equality and Partiality*. Oxford: Oxford University Press, 1991.

———. *Mortal Questions*. Cambridge: Cambridge University Press, 1979.

———. *The Possibility of Altruism*. Princeton: Princeton University Press, 1970.
National Safety Council. "Odds of Dying." Accessed November 15, 2020. https://injuryfacts.nsc.org/all-injuries/preventable-death-overview/odds-of-dying/.
Norcross, Alastair. "Comparing Harms: Headaches and Human Lives." *Philosophy and Public Affairs* 26, no. 2 (April 1997): 135–67.
Nozick, Robert. *Anarchy, State, and Utopia*. New York: Basic Books, 1974.
Oberdiek, John. *Imposing Risk*. Oxford: Oxford University Press, 2017.
Otsuka, Michael. "Determinism and the Value and Fairness of Equal Chances." Unpublished manuscript.
———. "Risking Life and Limb: How to Discount Harms by Their Improbability." In Cohen, Daniels, and Eyal, *Identified versus Statistical Lives*, 77–93.
Rapoport, Anatol. *Strategy and Conscience*. New York: Harper and Row, 1964.
Rawls, John. *A Theory of Justice*. Rev. ed. Oxford: Oxford University Press, 1999.
Reibetanz, Sophia. "Contractualism and Aggregation." *Ethics* 108, no. 2 (January 1998): 296–311.
Rüger, Korbinian. "On *Ex Ante* Contractualism." *Journal of Ethics and Social Philosophy* 13, no. 3 (June 2018): 240–58.
Scanlon, T. M. "Contractualism and Justification." Unpublished manuscript.
———. "Contractualism and Utilitarianism." In *Utilitarianism and Beyond*, edited by Amartya Sen and Bernard Williams, 103–28. Cambridge: Cambridge University Press, 1982.
———. *Moral Dimensions*. Cambridge, MA: Belknap Press of Harvard University Press, 2008.
———. "Reply to Zofia Stemplowska." *Journal of Moral Philosophy* 10, no. 4 (January 2013): 508–14.
———. *What We Owe to Each Other*. Cambridge, MA: Belknap Press of Harvard University Press, 1998.
Sher, George. "What Makes a Lottery Fair?" *Noûs* 14, no. 2 (May 1980): 203–16.
Sober, Elliott. "Evolutionary Theory and the Reality of Macro-Probabilities." In *The Place of Probability in Science*, edited by Ellery Ells and J. H. Fetzer, 133–61. Dordrecht: Springer, 2010.
Spiekermann, Kai. "Good Reasons for Losers: Lottery Fairness and Social Risk." Unpublished manuscript.
Steuwer, Bastian. "Aggregation, Balancing, and Respect for the Claims of Individuals." *Utilitas* (forthcoming). Published ahead of print, August 25, 2020. https://doi.org/10.1017/S0953820820000217.
Strevens, Michael. "Probability out of Determinism." In *Probabilities in Physics*,

edited by Claus Beisbart and Stephan Hartmann, 339–64. Oxford: Oxford University Press, 2011.

Tadros, Victor. "Controlling Risk." In *Prevention and the Limits of the Criminal Law*, edited by Andrew Ashworth, Lucia Zedner, and Patrick Tomlin, 133–55. Oxford: Oxford University Press, 2013.

Temkin, Larry. *Rethinking the Good*. Oxford: Oxford University Press, 2012.

Thomson, Judith Jarvis. *The Realm of Rights*. Cambridge, MA: Harvard University Press, 1990.

Voorhoeve, Alex. "How Should We Aggregate Competing Claims?" *Ethics* 125, no. 1 (October 2014): 64–87.

Walen, Alec. "Risks and Weak Aggregation: Why Different Models of Risk Suit Different Types of Cases." *Ethics* 131, no. 1 (October 2020): 62–86.