

WHY CONTRACTUALISM CANNOT ACCEPT EQUAL TREATMENT FOR EQUAL STATISTICAL LOSS

Jay Zameska

THE CONTRACTUALISM of T. M. Scanlon is a prominent nonconsequentialist approach to the ethics of risk, resulting in a significant and growing body of literature on contractualism and risk.¹ Contractualism holds that an action is morally permissible only if it is justified by principles that no one could reasonably reject. Unlike most consequentialist theories, however, contractualism rejects interpersonal aggregation in moral reasoning. Instead, it requires comparing the strength of individual personal reasons to determine which principles can be reasonably rejected. As such, a very strong reason of one person cannot be outweighed by combining many weaker reasons of other persons.² The goal is to identify “the principle whose implications are most acceptable to the person to whom it is least acceptable.”³

Debates over contractualism and risk typically make a distinction between *ex ante* and *ex post* variants of the theory. *Ex ante* contractualism evaluates principles based on individuals’ prospects under the proposed principle, which allows discounting the value of the expected outcome by its improbability. Here, however, my focus is specifically on *ex post* versions of contractualism, which do not allow discounting benefits and burdens by their improbability. Instead, they require considering the full value of the outcome, regardless of the probability of it occurring. In the rest of the article, unless otherwise specified, by ‘contractualism,’ I mean specifically the *ex post* variant of Scanlon’s contractualism.⁴

1 See Scanlon, *What We Owe to Each Other*. For discussion of contractualism and risk, see, *inter alia*, Hayenhjelm and Wolff, “The Moral Problem of Risk Impositions”; Fried, *Facing Up to Scarcity*; Ashford and Mulgan, “Contractualism”; and John and Curran, “Costa, Cancer and Coronavirus.”

2 Scanlon, *What We Owe to Each Other*.

3 Kumar, “Risking and Wronging,” 35.

4 For the distinction between *ex ante* and *ex post* contractualism, see Frick, “Contractualism and Social Risk”; and Suikkanen, “*Ex Ante* and *Ex Post* Contractualism,” which also

When addressing risky cases, some argue that *ex post* contractualism should comply with the following principle:

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.⁵

From here on, I refer to this as the *equal treatment principle*. Despite being an intuitively plausible principle, in this article, I argue that contractualism cannot comply with the equal treatment principle. This is because contractualism is a fundamentally *relational* moral theory.⁶ In brief, this means that contractualism is premised on the value of standing in a moral relationship with others. As such, contractualist moral deliberation is primarily not about assessing outcomes but about determining whether an action is justifiable to the individuals it affects. To structure this deliberation, contractualism represents those to whom we owe justification through *standpoints*. Standpoints are the generic positions from which individuals assess the burdens imposed by moral principles when determining whether those principles can be reasonably rejected. Importantly, standpoints are merely epistemic heuristics to help represent others and do not have normative importance themselves.⁷ This entails that unoccupied standpoints should not be considered in moral deliberation.

On this basis, I argue that we should draw a distinction between *definite standpoints*—those that we know with certainty pertain to at least one actual

develops a hybrid model. For defenses of *ex post* contractualism in the context of risk, see Reibetanz, “Contractualism and Aggregation,”; and Rüger, “On *Ex Ante* Contractualism.” Finally, for a broad overview of the *ex post* versus *ex ante* debate, see Ashford and Mulgan, “Contractualism.” While the choice between these frameworks is a major debate, often adjudicated by cases like the ones discussed here, the aim of this article is more modest. It seeks to explore a question that is internal to the *ex post* framework: whether the theory’s own foundational commitment to relational justification has unexamined consequences for how it must treat principles governing risk distribution. Understanding these consequences on the theory’s own terms, I suggest, is a necessary step that is logically prior to the larger debate.

- 5 Steuwer, “Contractualism, Complaints, and Risk,” 119. I take Steuwer’s wording and discussion of this as my focus because they are a particularly clear and explicit example of this principle. But this general idea lies behind a significant amount of discussion, especially in opposition to the identified victim bias. I return to discussion of the identified victim bias in section 3.
- 6 Scanlon, *What We Owe to Each Other*.
- 7 See, *inter alia*, Scanlon, *What We Owe to Each Other*, ch. 5; Gibb, “Relational Contractualism and Future Persons”; Katz, “Contractualism, Person-Affecting Wrongness and the Non-Identity Problem”; and Martin, “Navigating Nonidentity.”

person—and *indefinite standpoints*—those about which we are uncertain whether they pertain to any actual individuals. I argue that given contractualism's relational foundation, we should respond differently to definite and indefinite standpoints. We should, in general, prioritize the reasons of definite standpoints over those of indefinite standpoints. This in turn demonstrates that the equal treatment principle cannot always be correct: cases that are probabilistically equal may in fact differ in their justifiability and are thus morally different, from a contractualist perspective.

To develop this argument, in section 1, I introduce the distinction between definite and indefinite standpoints in more detail, and explain why it follows from contractualism's relational foundation. In section 2, I present an example case to demonstrate that even when two choices involve equal expectations of statistical loss, their justifiability may differ depending on whether they involve definite or indefinite standpoints. I argue that, in the case at hand, prioritizing definite standpoints is necessary to avoid acting unjustifiably, and thus, the principle of equal treatment cannot be right. This shows that cases with equal expected statistical losses cannot always be treated equally. In section 3, I address several objections, including concerns about the distinction between objective and epistemic risk, the potential implications of rejecting the equal treatment principle, challenges to the requirement that contractualism requires justifiability actions to actually existing individuals, as opposed to merely hypothetical ones, and the distinction between act and rule contractualism. Finally, in section 4, I consider the broader implications, including how this argument provides a contractualist justification for the identified victim bias and a possible framework for *ex post* discounting.

1. STANDPOINTS, UNCERTAINTY, AND INDIVIDUAL JUSTIFICATION

Consider two possibilities:

Gun: There is a person and a revolver in front of me. The revolver's six-chamber cylinder has already been loaded with a single round, and the cylinder has already been spun. I can pick up the revolver, point it at the person, and pull the trigger one time.

Box: There is an opaque box in front of me. There is a one-sixth chance that someone has been placed inside the box. I can press a button and destroy the box (and anyone inside).

Assuming that I must pick one of these two possibilities, which should I choose? At the moment, it may not seem to matter which option I choose. They

seem to be equal in all relevant respects.⁸ From an *ex ante* perspective, there is a one-sixth chance of killing someone, regardless of which option I choose. From an *ex post* perspective, there is one single death to consider, regardless of which option I choose. As such, Gun versus Box seems to be precisely the kind of dilemma to which the equal treatment principle should apply. In both cases, the expected statistical loss is equal: in one out of six possible worlds, a single person dies regardless of whether I choose Gun or Box. The only thing that changes is the distribution of deaths across designators (the victim of Gun versus the person in the opaque box).

Given this equal expected statistical loss, the equal treatment principle holds that we should treat the two options the same. Perhaps we should flip a coin, perhaps we should simply think either choice is permissible, or perhaps we should employ some sort of tiebreaker.⁹ Here, I assume we should flip a coin, although ultimately it does not matter to this argument how contractualists should respond to equal cases: Gun and Box are probabilistically equivalent, but they are not morally equal for contractualists, and thus, they should not be treated equally.

This is because contractualism is fundamentally a *relational*, or *second-personal*, moral theory.¹⁰ As Rahul Kumar explains, “the contractualist claim is that all persons stand in a particular kind of relationship to one another, the

8 One might worry that from an *ex ante* perspective, the cases as described may not be equal. Gun appears to impose a concentrated risk on a specific person, while Box could be seen as imposing a trivial risk dispersed over a large population. While this article’s focus is on the *ex post* complaints, it is possible to stipulate a background structure that may neutralize this concern. We can imagine that the potential victims for both scenarios are drawn from two separate, equally large populations, such that the initial *ex ante* chance of being the one ultimately harmed is identical in both setups. For Gun, one person is selected from a population and placed in the role of ‘victim of Gun’. For Box, a different person is selected from the other population to be the potential ‘person in the box’, who is then subject to the one-sixth chance of actually being placed inside. Admittedly, from an *ex ante* perspective that employs Frick’s decomposition test, this equalizing setup may not be convincing, as such a view focuses on the unequal risk distributions at the immediate moment of choice, regardless of previous risk distributions. From the strictly *ex post* perspective of this article, however, what matters is that both setups entail an equal expected statistical loss. As such, the moral asymmetry at the moment of decision—the choice between a definite and an indefinite standpoint—remains, and it is this *ex post* feature that my argument addresses. My thanks to an anonymous reviewer for pushing me to clarify this point.

9 See Taurek, “Should the Numbers Count?”; Wasserman, “Let Them Eat Chances”; Saunders, “A Defence of Weighted Lotteries in Life Saving Cases”; Scanlon, *What We Owe to Each Other*; and Reibetanz, “Contractualism and Aggregation.”

10 Scanlon, *What We Owe to Each Other*. For more general discussion of second-personal moral theory, see also Darwall, *The Second-Person Standpoint*.

ideal form of which realizes a distinct value, that of mutual recognition.”¹¹ As such, contractualism is premised on the value of standing in a moral relationship to others, and it is the value of this relationship that requires us to act in ways that others could not reasonably reject.¹² Not only does this relational foundation help define some of the theory’s most controversial and distinctive commitments (for example, its famous ban on aggregation), but I argue that it also means that we must treat Gun and Box differently—and thus, the equal treatment principle cannot always be right. Although Gun and Box may be probabilistically equal, they represent a more fundamental difference in contractualist theory regarding how contractualists should address two different kinds of standpoints, which I call *definite* and *indefinite* standpoints. To motivate the relevance of this distinction, first I explain the concept of standpoints in contractualism in more depth, before explaining why contractualism must treat them differently.

Standpoints are the generic positions from which individuals assess the burdens imposed by a given principle when determining whether that principle can be reasonably rejected. An important but often overlooked feature of such standpoints is that they are intended to serve only as a kind of *epistemic heuristic* that allows us to assess the claims of actually existing but currently unidentified others.¹³ As Corey Katz explains, “the value at the heart of contractualist metaethics is the value of living in *actual second-personal relationships* of mutual justifiability,” and as a result, “what matters for an agent is that her action be justifiable *to the actual people she is in relationship with* on reasons those people could not reasonably reject.”¹⁴ It is the value of these moral relationships, constituted by a shared capacity for rational agency, rather than the reasons available to standpoints per se, that underwrites the contractualist focus on justifying our actions to others on grounds they could not reasonably reject.¹⁵

11 Kumar, “Risking and Wronging,” 258.

12 Scanlon, *What We Owe to Each Other*; and Southwood, “Moral Contractualism” and *Contractualism and the Foundations of Morality*.

13 See Gibb, “Relational Contractualism and Future Persons”; Katz, “Contractualism, Person-Affecting Wrongness and the Non-Identity Problem”; and Martin, “Navigating Non-identity”—all of which provide significant discussion of this epistemic heuristic aspect of standpoints.

14 Katz, “Contractualism, Person-Affecting Wrongness and the Non-Identity Problem,” 112 (emphasis added).

15 It also is this focus on the moral relationships between rational agents rather than on mutual advantage or pure self-interest that leads Scanlon’s contractualism to often be described as a Kantian contract theory (as opposed to, e.g., Hobbesian contract theories). For discussion, see Southwood, “Moral Contractualism.”

In other words, standpoints are simply a way to overcome the limited information regarding the particular identities of the actual individuals our actions will affect.¹⁶ Given that they are limited to an epistemic role, standpoints have normative relevance only insofar as they represent *actual* individuals. This means that the moral relevance of the generic reasons associated with a standpoint is contingent on that standpoint being occupied by at least one actual person. Thus, when we are certain no one actually occupies a given standpoint, we can disregard that standpoint and its generic reasons, as “the reasons of a merely imagined but unoccupied standpoint have no bearing on any of our moral relationships, and it is the value of these relationships, and only these relationships, that drive our moral judgments.”¹⁷

In short, contractualism is fundamentally premised on the moral relationship between actual individuals. Standpoints are an epistemic heuristic to allow us to consider the reasons of actual but unidentified others.¹⁸ As such, they have no independent moral significance beyond the individuals they represent. This entails that when we are certain a standpoint is empty—that it does not correspond to an actual individual—we should not consider reasons from that standpoint. Conversely, if we know a standpoint is occupied—that is, if we know it corresponds to at least one actual individual—then we must consider the reasons from that standpoint.

I take both of the preceding points to be fairly uncontroversial and, at this stage, perhaps rather uninteresting. But in between “certainly empty” and “certainly occupied” are a range of standpoints that are morally interesting. These are standpoints where we do not know if they are in fact occupied. I call these *indefinite standpoints*, in contrast to *definite standpoints*, where we are certain that the standpoint is occupied. Definite standpoints have clear normative force because they represent actual persons to whom justification is owed. Indefinite standpoints, by contrast, represent a special case of uncertainty—where it is unclear whether anyone occupies the standpoint and therefore whether any moral relationship is affected by our decision. The uncertainty involved in each of these is fundamentally different. In a case like Gun, the uncertainty is about the outcome for a person we know exists; the standpoint is definite, but the result is not. In a case like Box, however, the uncertainty is more fundamental: we are uncertain about whether there is an actual person to whom justification could be owed in the first place. This introduces a distinct risk, unique to the

16 For the original discussion of this point, see Scanlon, *What We Owe to Each Other*, 202–6.

17 Gibb, “Relational Contractualism and Future Persons,” 12.

18 *Inter alia*, Scanlon, *What We Owe to Each Other*; and Gibb, “Relational Contractualism and Future Persons.” See also Martin, “Navigating Nonidentity.”

relational commitments of contractualism: the risk of acting on the basis of a merely possible standpoint to the detriment of a certainly existing one. As I argue in the next section, this would be to act in a way that is fundamentally and uniquely unjustifiable.¹⁹ This distinction between definite and indefinite standpoints matters because, as just explained, the normative relevance of a standpoint depends on whether it represents actual individuals. Contractualism's relational foundation requires justifiability to actual persons, and as such, indefinite standpoints introduce uncertainty about whether justification is possible at all. As I argue in the next section, this means that in the cases introduced earlier, the uncertainty in Box is different from the uncertainty in Gun, even if they are probabilistically equal and have equal expected statistical losses.

2. WHY EQUAL STATISTICAL LOSS DOES NOT GUARANTEE EQUAL JUSTIFIABILITY

With this discussion of standpoints in mind, let us return to the choice from earlier in this article.

Gun: There is a person and a revolver in front of me. The revolver's six-chamber cylinder has already been loaded with a single round, and the cylinder has already been spun. I can pick up the revolver, point it at the person, and pull the trigger one time.

Box: There is an opaque box in front of me. There is a one-sixth chance that someone has been placed inside the box. I can press a button and destroy the box (and anyone inside).

The idea that we should treat Gun and Box equally stems at least in part from the idea that both potential victims can present equally strong reasons against shooting the gun or destroying the box. In both cases, they can each offer the reason "I will die."²⁰ However, these reasons should not in fact be treated equally.

19 To make this distinction clearer, we can label these as two different kinds of risk. The uncertainty in Gun can be called *outcome risk*—uncertainty about the consequences of an action. The more fundamental uncertainty in Box is what I call *justificatory risk*—the risk that our action will be foundationally unjustifiable because a necessary condition for justification is absent. As I argue in the next section, contractualism should aim to avoid justificatory risk.

20 Note again, this is because I am focusing only on *ex post* contractualism, which does not allow discounting based on probability. Consequently, each person's complaint should represent the full disvalue of the possible negative outcome, regardless of its likelihood—in this case, death. Contrast this with the *ex ante* version of the complaint—namely, "I face a one-sixth chance of death."

To see why, it is helpful to consider the possible states of the world regarding Box:

S_1 : The box is full.

S_{2-6} : The box is empty.

Imagine that we are in S_1 . If the box is full, either choice (Gun or Box) would be justifiable. This is because presumably, if you can spare only one of two people from an equal harm but not both, and everything else is equal, it is justifiable to spare either one, as the reasons on either side are equally strong. It is a tragic choice, but it is not an unjustifiable one, and justification is what matters to contractualism. If we flip a coin, regardless of whether it is heads or tails, our resulting choice will be justifiable to both parties, and neither will be able to reasonably reject our decision.

Consider how things would change if we are instead in any of S_{2-6} . If we decide to destroy the box, our action would clearly be justifiable, for exactly the same reasons as in S_1 . But now imagine that we instead decide to shoot the gun. For this to be justifiable would introduce a dilemma: given that there is no one in the box, to justify our choice on the basis of the reason "I will die" from the person-in-the-box standpoint would require either that either (1) it is justifiable for reasons that have nothing to do with our moral relationship, as there is no one in the box with whom we have such a relationship or (2) we do in fact have a moral relationship with the empty person-in-the-box standpoint. This raises a dilemma for contractualism, either horn of which is unacceptable.²¹ The first horn is incompatible with the basic structure of the theory, and the second horn stretches the notion of the "moral relationship" to an implausible degree.

This dilemma arises because if the box is empty, there can be no individual, personal reasons of the kind that contractualism demands. The justification for choosing Gun in S_{2-6} (or choosing to flip a coin) would fail to be based on the right kind of reason, as there is no actual individual to hold such a reason. Alternatively, if we were to claim that there are such reasons, this entails that, in Gibb's terms, "we would have to claim that we share a moral relationship with something like a metaphysical entity known as a 'standpoint.'"²² In short, if the box is empty, the contractualist mode of justification breaks down. We cannot properly consider the reasons for rejection from the person-in-the-box standpoint because if the box is empty, they would not be the right kind of reasons

21 For the original discussion of this dilemma, see Gibb, "Relational Contractualism and Future Persons."

22 Gibb, "Relational Contractualism and Future Persons," 11.

required by contractualist moral deliberation, and as a result, we would act in a way that the victim in Gun could reasonably reject.

Thus, if we want to avoid the risk of acting unjustifiably, we must choose Box. Choosing Box does not necessarily minimize the risk of harm, but it minimizes the risk of acting unjustifiably. This is because there is a meaningful distinction between the risk of (justifiably creating) a bad outcome and the risk of acting unjustifiably in the first place; for contractualism, the latter is a more serious failure.²³ Given that contractualism requires that you act justifiably, it also requires that you act in ways that are, *ceteris paribus*, guaranteed to be justifiable over acting in ways that entail a risk of being unjustifiable. This means that the equal treatment principle cannot be right. In this case, we have equal statistical loss, but due to differences in standpoints, we cannot treat them equally. More than just highlighting a problem for the equal treatment principle, this argument also generalizes to any competitive choice between definite and indefinite standpoints with equal harms (or benefits). In competitive cases where both definite and indefinite standpoints face an equal burden, contractualists should not treat them equally. To do so risks acting unjustifiably, in an analogous way to choosing Gun over Box.

As such, in any case where we either prioritize the reasons of indefinite standpoints or treat them equally to definite standpoints (by, e.g., flipping a coin or engaging in some similar procedure), we run the risk of acting unjustifiably. In contrast, if we prioritize the reasons of definite standpoints, we never run the risk of acting unjustifiably. In other words, reasons from definite standpoints carry greater weight than those from indefinite standpoints. Here, I have limited my discussion to cases with equal harms in order to mirror the structure of the case above (and to address the equal treatment principle specifically). Given that in the case at hand, the only difference between the reasons under consideration is whether they stem from definite or indefinite standpoints, this greater weight is a decisive consideration. However, in the next section, I briefly return to the question of how to deal with this priority to definite standpoints when harms and benefits are not equal.

To summarize, in cases involving competitive choice between definite and indefinite standpoints, as in Gun versus Box, we must prioritize the interests of definite standpoints. Prioritizing the reasons of definite standpoints is the only way to avoid the risk of acting unjustifiably. If the indefinite standpoint turns out to be unoccupied, it cannot provide reasons of the kind demanded by

23 This maps onto the distinction drawn earlier between what I call outcome risk and justificatory risk. The risk of a bad outcome (harm) is an outcome risk, present in both Gun and Box. The risk of acting unjustifiably, however, is justificatory risk, introduced by the indefinite standpoint in Box.

contractualism as a fundamentally relational moral theory. As such, acting on the basis of such reasons against the interests of another, actual person is unjustifiable. Given that contractualism requires us to act justifiably, it also requires us to choose actions that are certainly justifiable over actions that are “riskily” justifiable, all else equal. Consequently, in competitive cases with equal burdens, contractualists must always prioritize the interests of definite standpoints over indefinite standpoints, and thus, the equal treatment principle should be rejected. Next, I consider four major objections before turning to the broader implications of this argument.

3. OBJECTIONS: OBJECTIVE VERSUS EPISTEMIC RISK, IMPLAUSIBLE CONSEQUENCES, THE ROLE OF HYPOTHETICAL JUSTIFIABILITY, AND ACT VERSUS RULE CONTRACTUALISM

In this section I address four objections—namely, (1) that this argument confuses or otherwise neglects the distinction between objective and epistemic risk, risk concentration/dispersal, or the distinction between the *luckless* and the *doomed*; (2) that this argument against the principle of equal treatment cannot be right because it entails implausible consequences; (3) that the requirement that contractualism requires justifiability to actually existing people is incorrect; and (4) that the argument is limited to “act contractualism” and fails when applied to the traditional rule-based formulation of the theory. I address each of these in turn.

First, some may object that I am drawing a distinction that already has a more common name: *objective risk* versus *epistemic risk*. In other words, the objection holds that there is already an established difference between Gun and Box, and it has nothing to do with standpoints but rather has to do with whether the risk is objective or merely epistemic. In one common way to cast this distinction, objective risk refers to “truly” risky cases, whereas epistemic risk refers to cases that are risky simply because of our lack of knowledge.²⁴ Some have sought to support the idea that these different kinds of risk matter morally and thus that perhaps we ought to respond to them differently, even if they have the same probabilities.²⁵

24 For general discussion of this objective-versus-epistemic split in interpretations of probability, see Gillies, *Philosophical Theories of Probability*.

25 For discussion see, *inter alia*, Tadros, “Controlling Risk”; Oberdiek, *Imposing Risk*; Fried, *Facing Up to Scarcity*; and Spiekermann, “Good Reasons for Losers.” For discussion of this distinction specifically in the context of contractualist approaches to risk, see Otsuka, “Risking Life and Limb”; Fried, *Facing Up to Scarcity*; and Steuwer, “Contractualism, Complaints, and Risk.”

Along these lines, some may object that Gun is objectively risky while Box is merely epistemically risky. The objection would hold that at the time of decision, it is already decided whether there is someone in the box, but we simply do not know. In contrast, spinning the cylinder in the revolver in Gun is “truly” risky. As such, it may be that these cases are not in fact equal. This is not quite right, however, as assuming there is some mechanism similar to spinning the revolver’s cylinder to decide whether someone is in the box, in both Gun and Box, we are dealing with objectively risky situations—at least insofar as there are such things as objective risks.

Some may still insist that the objective risk in Box occurs *before* you make your decision, and so it truly is a case of objective risk (Gun) versus epistemic risk (Box). At the time of decision, it is either the case that there is someone in the box, or it is the case that there is not—so it cannot be objectively risky in the way that Gun is. Instead, it is merely epistemically risky because it is based on a lack of knowledge of the contents of the box. However, this way of casting the objection does not succeed either. Gun is at the time of trigger pull *also* purely epistemic. It has already been decided which chamber the bullet is in, I simply do not know which. So either both Gun and Box are epistemic, or both are objective, depending on the precise slice of time we select. As such, the difference between epistemic and objective risks cannot be what makes a moral difference here. Instead, as discussed earlier, what matters is whether the standpoint is definite or indefinite.

In other words, either both Gun and Box are objectively risky, or both are epistemically risky. At the time of decision making, it is already decided whether there is a bullet in the firing chamber, just as it is already decided whether there is someone in the box. So the case is merely epistemically risky either way. Or we could partition things differently and instead insist that it is objectively risky, as in the preceding stage of each case, there is an objective one-sixth chance of the bullet ending up in the firing chamber and an objective one-sixth chance of placing a person in the box. The underlying point is that we can shift whether it is objectively or epistemically risky based on how we partition time, but in either case, it applies equally to both Gun and Box. What makes the difference is whether our actions are justifiable to everyone involved at the time of decision, and as argued above, this turns in part on whether we are considering the reasons of definite or indefinite standpoints.

To put the point more directly, the definite/indefinite distinction cuts across the epistemic/objective one by focusing on a different feature of the situation. The moral difference stems not from the nature of the risk (epistemic/objective) but from what it is we are uncertain about. In both cases, we are uncertain whether our action will result in someone’s death; this is a risk

concerning the outcome for a person. But exclusively in Box, we are also uncertain whether our action can be justified at all because we are uncertain whether the person-in-the-box standpoint is actually occupied.²⁶ This second form of uncertainty is unique to a relational moral theory like contractualism, which is grounded in the idea of justifiability to actual, existing individuals. It is this latter uncertainty about the grounds of justification itself, not the epistemic nature of the risk, that makes the cases morally different.

A related objection holds that the definite/indefinite distinction is simply a proxy for the more familiar distinction of risk concentration versus dispersal. One might argue that a preference for Box over Gun is easily explained by the fact that the former seems to distribute risk broadly, while the latter concentrates it on a single individual, thus making the definite/indefinite distinction redundant.²⁷ However, to show that the definite/indefinite distinction is conceptually separate from concerns about risk distribution, it is useful to consider a hypothetical case where the two principles come into direct conflict. Imagine a choice between (1) imposing a large, concentrated risk on a single indefinite standpoint (e.g., a high probability of harm to a person whose existence is uncertain) and (2) imposing a small, dispersed risk across many definite standpoints (e.g., a low probability of harm to many actual, identified people). A principle focused only on avoiding concentrated risk may favor option 2. The argument of this article, however, provides a strong reason to prefer option 1. This is because choosing 2 involves burdening actual people for the sake of a merely possible one, risking a failure of justifiability, whereas choosing 1 avoids this particular kind of relational failure. That the two principles can yield conflicting recommendations proves that they are conceptually distinct.

A final version of this objection holds that the distinction I draw between definite and indefinite standpoints is reducible to another existing distinction: rather than a relabeling of objective versus subjective, it may also appear to be a relabeling of “doomed” and “luckless” victims.²⁸ On this view, a doomed victim is one whose fate is already determined, even if it is epistemically unknown,

26 This corresponds to the two kinds of risk discussed earlier. The uncertainty about death can be described as outcome risk. The second, more fundamental uncertainty about whether the standpoint is occupied—and thus whether our action is justifiable—is what I call justificatory risk. It is the presence of this second kind of risk in Box that makes the case morally different for contractualism.

27 My thanks to an anonymous referee for suggesting this objection. It is worth noting, however, that the moral significance of risk concentration is itself a matter of debate. For skepticism on this point, see Eyal, “Concentrated Risk, the Coventry Blitz, Chamberlain’s Cancer.”

28 I thank an anonymous referee for raising this objection and pushing me to clarify the relationship between my account and this distinction. For discussion of this distinction,

while a luckless victim faces a genuinely probabilistic risk. The objection is that my argument for prioritizing definite standpoints is simply another way of arguing that we should prioritize doomed victims.²⁹

While the two may often align, they are grounded in different moral considerations. The doomed/luckless distinction is about the causal nature of the risk that an individual faces. My account, in contrast, is grounded in the relational requirements of contractualism; the definite/indefinite distinction turns on our certainty about a standpoint's occupation and the corresponding risk of justificatory failure. To see how these justifications come apart, consider the following variation on Gun.

Definite but Luckless: A revolver is pointed at a person, but the trigger is now connected to a truly indeterministic quantum randomizer with a one-sixth chance of firing one time.

The victim in this scenario is luckless, not doomed. A theory based on prioritizing the doomed would not grant this person special priority. Yet on my account, the victim's standpoint is still definite—we are certain an actual person occupies that role—and thus their reasons take priority over the indefinite standpoint in Box to avoid justificatory risk.

Furthermore, my distinction is better equipped to explain the moral difference in the original Gun versus Box decision. Arguably, both victims could be classified as doomed, since the chancy portions of the relevant causal chains (the spin of the revolver's cylinder, the sorting into the box) have already been completed. The doomed/luckless distinction therefore struggles to differentiate them. The definite/indefinite distinction, however, clearly separates the cases based on the certainty of standpoint occupation. Thus, the justification for prioritizing definite standpoints is a distinct, relational one, not based on the nature of the risk that an individual faces.

Perhaps the strongest objection is that this argument may seem to force contractualism into a position where it must endorse implausible conclusions. For example, consider a new choice that must be made.

Gun: There is a person and a revolver in front of me. The revolver's six-chamber cylinder has already been loaded with a single round, and the cylinder has already been spun. I can pick up the revolver, point it at the person, and pull the trigger one time.

see Otsuka, "Risking Life and Limb"; and Steuwer, "Contractualism, Complaints, and Risk."

29 For discussion of prioritizing between doomed and luckless victims, see Steuwer, "Contractualism, Complaints, and Risk."

Box (Massive): There is an opaque box in front of me. There is a one-sixth chance that one hundred people have been placed inside the box. I can press a button and destroy the box (and anyone inside).

The view I defend seems to hold that contractualists should choose to potentially sacrifice all one hundred people in *Box (Massive)* rather than risk a single fatal shot in *Gun*. After all, the first is a definite standpoint, whereas each of the one hundred others is indefinite, and I argue that contractualism should prioritize definite over indefinite standpoints. In other words, my argument seems to require potentially sacrificing one hundred lives to spare a single individual a one-sixth risk of death.³⁰ And this seems extremely implausible.

Notably, this objection does not defend the equal treatment principle or undercut the argument against it. It just shows contractualism to sometimes yield putatively implausible conclusions. However, the equal treatment principle does not prevent this kind of implausible conclusion. In fact, the equal treatment principle does not apply because the expected statistical loss is not equal. As such, this objection does not undermine this argument against the equal treatment principle. Instead, this objection represents a problem inherent to contractualism's anti-aggregative commitments rather than a specific issue with how it handles expected statistical losses. However, this is a more general problem for contractualism and is not new.³¹ There are responses in the contractualist literature, ranging from Scanlon's tiebreaking argument to weighted lotteries.³² Since this is beyond the scope of this article, I do not address this substantial body of literature here.

As such, the argument presented here is not shown wrong by this objection, as this implausible conclusion is a consequence of the structure of contractualism, and similar implausible conclusions can be constructed without reference to risk, uncertainty, or expected statistical loss.³³ Consequently, my argument against the principle of equal treatment may in fact also work as an argument

30 This phrasing sets aside the fact that I am looking at this in terms of *ex post* contractualism, which requires viewing this as one death versus one hundred deaths rather than as a one-sixth chance of one death versus one hundred deaths. But the objection is sharper and stronger when considering it in *ex ante* terms.

31 For an overview and critical discussion of contractualism's difficulties with aggregation, see Fried, *Facing Up to Scarcity*.

32 Scanlon, *What We Owe to Each Other*; and Saunders, "A Defence of Weighted Lotteries in Life Saving Cases."

33 See, for example, discussions over whether contractualists should save a single person or some larger number of people from death: *inter alia*, Taurek, "Should the Numbers Count?"; Otsuka, "Scanlon and the Claims of the Many Versus the One"; Fried, *Facing Up to Scarcity*; and Muñoz, "Each Counts for One."

against contractualism itself, if opposition to this conclusion is sufficiently strong. The fact that contractualism cannot endorse the equal treatment principle and will at least sometimes license sacrificing one hundred lives to avoid a one-sixth chance of death for a single individual may be recast as an objection to contractualism itself.³⁴ Thus, this implausibility objection does not actually challenge the argument against the equal treatment principle, although it may change its strategic use.

Finally, my argument is premised on a strict interpretation of contractualism's relational foundation, which holds that justification is owed to *actual* individuals even if the process of justification itself is idealized and hypothetical. This view, however, is not the only possible understanding, as there is ongoing debate in contractualism over exactly how idealized the process of justification should be. A contrasting, more strongly idealized view holds that the process of justification can be detached from actual persons. On this view, what matters for contractualist moral deliberation are the reasons associated with a standpoint itself, not whether that standpoint is currently occupied.

Against this highly idealized position, my argument defends the stricter, actualist interpretation. While contractualism is about justifiability rather than actual justification, I contend that what matters is justifiability to *actual individuals* rather than justifiability in a purely hypothetical sense.³⁵ As such, we need not justify ourselves to people who we know do not and will not exist. To see this, consider a case where there is equal expected loss between a standpoint we know is certainly occupied and one we know is certainly empty. It would be wrong to treat them equally on the grounds that contractualism is concerned with justifiability rather than justification. Even though the concern is justifiability, it is justifiability to *actual people* that matters. To deny this would be to reject the fundamentally relational basis of the theory.

Furthermore, contractualism's oft discussed *impersonalist restriction* gives another reason to doubt the idea that contractualism is concerned with justifiability to merely hypothetical people. The impersonalist restriction holds that "impersonal values are not themselves grounds for reasonable rejection," and as such, only personal reasons may be employed for contractualist justification.³⁶

34 It may alternatively be read as an objection specifically to the various features of contractualism that permit this rather than to contractualism as a whole. Some have already advocated for removing or revising the individualist restriction on similar grounds (e.g., Parfit, *On What Matters*), but the argument here could also be used to argue for revising the relational character of contractualism.

35 Although for an argument that contractualism should focus on actual justification, not justifiability, see Kim, "On the Need for Real Dialogue."

36 Scanlon, *What We Owe to Each Other*, 222.

However, the only kind of reasons that merely hypothetical and nonexistent individuals could have for reasonable rejection are impersonal reasons. They cannot be personal reasons because they are not the reasons of any (actual) individual. As such, if they are reasons at all, they must be impersonal reasons and are thus inadmissible in a contractualist framework. So the basic commitments of the theory speak against the idea that contractualism should consider merely hypothetical individuals. Nonexistent people cannot provide reasons of the right type.³⁷

The fourth and final objection I consider draws on recent discussions of the distinction between *act contractualism* and the more traditional *principle contractualism*, or *rule contractualism*.³⁸ In this framework, principle contractualism holds that an act is wrong if it violates a principle that no one could reasonably reject for the general regulation of behavior, whereas act contractualism assesses the wrongness of an act based on its own rejectability, independent of general principles.³⁹ The objection holds that the argument presented earlier may apply to the act-based version of the theory but not to the two-stage, rule-based contractualism initially put forward by Scanlon. In the context of the Gun versus Box cases introduced earlier, under the standard rule-based version of contractualism, we would assess not the single choice between Gun and Box but rather competing general principles or rules for governing such situations. For example, we might choose between a principle of equal treatment, which would mandate indifference (e.g., perhaps by flipping a coin in every Gun versus Box decision), and a principle of definite priority, which would require prioritizing the definite standpoint (i.e., always destroying the box).

From this perspective, if we imagine many worlds with Gun versus Box scenarios, a principle of equal treatment would result, *ex post*, in an identical number of bad outcomes as a principle of definite priority. If the number of deaths is the same in the long run, the objection goes, then neither rule is more reasonably rejectable than the other on these grounds. If this is correct, then

37 Further, under most views of reasons, nonexistent people cannot provide any reasons at all. See Martin, "Navigating Nonidentity," 97, especially n. 27.

38 As far as I am aware, Sheinman is the first to systematically put forward this distinction, arguing that principle contractualism is vulnerable to a charge of "principle worship" that is analogous to the "rule worship" objection in consequentialism ("Act and Principle Contractualism"). The possibility and nature of act contractualism has been the subject of recent discussion. For a defense of the possibility of act contractualism, for example, see Bourguignon, "On the Possibility of Act Contractualism." And for further discussion, see Salein, "Leaving Principle Contractualism Behind?" I return to Sheinman's discussion shortly.

39 Sheinman, "Act and Principle Contractualism"; and Bourguignon, "On the Possibility of Act Contractualism."

the equal treatment principle appears to be a natural consequence of rule-based contractualism in this context. As such, the earlier argument for rejecting the equal treatment principle is applicable only to act contractualism but not to rule contractualism.⁴⁰ This is a serious objection, as it threatens to undercut the applicability of this argument to the canonical and most widely discussed formulation of contractualism. However, I believe that the main argument put forward in section 2 works in both a rule-based and act-based understanding of contractualism.

To see why, consider the choice between the two competing principles, equal treatment and definite priority, across two levels of justification: the justifiability of the act relative to the principle (*justifiable*_{RULE}) and the justifiability of the act itself (*justifiable*_{ACT}). Even granting the premise of the objection—namely, that since both the equal treatment and definite priority principles produce the same number of deaths, they are both equally justifiable_{RULE}—there still is a significant difference in overall justifiability between equal treatment and definite priority. The difference is visible when we examine the justifiability of the acts licensed by each rule: the definite priority principle is unique because in Gun versus Box, it *always* licenses actions that are both justifiable_{RULE} and justifiable_{ACT}. Even in a tragic scenario where the box is full, the choice is still justifiable (both in act and rule terms) because it is a choice between two equally strong complaints. The principle of equal treatment, however, does not share this feature. Even granting that equal treatment may be justifiable_{RULE} in approximately one-sixth of cases, it licenses an action that is unjustifiable_{ACT}—namely, shooting the gun when the box is empty. I suggest that this asymmetry creates a clear reason to prefer the definite priority principle. A principle whose licensed actions are always justifiable at both the act level and the rule level is superior to a principle that, while licensing actions that are justifiable at the rule level, sometimes demands that agents perform actions that are unjustifiable at the act level. Put another way, to choose the equal treatment principle over the definite priority principle is to at least sometimes create a needless conflict between what the rule permits and what is justifiable to the actual person in front of us. As such, the definite priority principle dominates the equal treatment principle, given our uncertainty, as it avoids the risk of acting unjustifiably.⁴¹ In other words, the definite priority principle is sometimes better than

40 My thanks to a helpful anonymous reviewer for suggesting this objection and discussion.

41 Here, 'dominance' refers to a principle of choice where one option is considered superior if it is better in at least one possible state of affairs and not worse in any other. For a classic discussion, see Luce and Raiffa, *Games and Decisions*. For a contemporary accessible overview of the dominance principle, see Peterson, *An Introduction to Decision Theory*.

the equal treatment principle (when the box is empty) and never worse than the equal treatment principle (when the box is full).

That justifiable_{RULE} and justifiable_{ACT} can come apart and that sometimes a justifiable rule can license an unjustifiable action lead to the problem of “principle worship” for rule-based versions of contractualism.⁴² While Hanoch Sheinman develops this “principle worship” objection on the basis that principle contractualism sometimes implausibly “forego[es] unrejectability in the actual world for the sake of [principle] conformity,” this objection is also relevant to the relational aspect of contractualism that is central to my own argument.⁴³ The reason we care about finding unrejectable principles is because we value standing in a relationship of mutual recognition with other individuals. Given this, a contractualist rule or principle is only as good as the justifiability of the actions it prescribes. If a rule directs me to act in a way that is not justifiable to the person before me, the fact that the rule produces justifiable outcomes on average or in the abstract does not negate the justificatory failure with respect to that individual.

This points to a more fundamental feature of any plausible version of contractualism: the normative status of contractualist principles is contingent on the standpoints they affect actually being occupied. As such, the distinction between definite standpoints (which we know are occupied) and indefinite standpoints (which may be empty) remains relevant even when formulating and applying general principles, as in the standard version of the theory. As the earlier Gun versus Box choice demonstrates, the complaint from the definite standpoint in Gun has a different justificatory status than the complaint from the indefinite standpoint in Box. Any proposed rule (e.g., the equal treatment principle) must itself be justifiable to all. A rule that ignores the difference between definite and indefinite standpoints is reasonably rejectable by those in definite standpoints, as it requires them to bear risks for the sake of standpoints that may not be occupied. As such, any justifiable rule must take account of the distinction between definite and indefinite standpoints, making this distinction relevant even within the two-stage, rule-based version of the theory.

42 This is analogous to a well-known problem for rule utilitarianism. See Smart and Williams, *Utilitarianism*. For a clear recent discussion of this objection to rule-based ethical theories, see also Copp, “The Rule Worship and Idealization Objections Revisited and Resisted.” I leave further discussion of this aside for now, as I aim not to weigh in on the emerging debate of act versus rule versions of contractualism but only to suggest that we have reason to think that the problem raised here for the equal treatment principle also applies to rule-based versions of contractualism.

43 Sheinman, “Act and Principle Contractualism,” 302.

Finally, even if one were not convinced that this argument applies to the standard rule-based form of contractualism, this objection, at worst, limits the scope of my conclusion. Showing that *ex post* act contractualism cannot accept the equal treatment principle is itself still significant, even if it does not also hold for the standard rule-based version. It still clarifies the commitments of a major branch of contractualist thought and reveals a tension between the theory's relational foundations and an intuitively plausible principle of fairness. This contribution remains valuable even if it does not extend to the classic rule-based formulations of the theory and is relevant for comparing the overall plausibility of act versus rule formulations of the theory more generally.

4. WHY REJECTING EQUAL TREATMENT MATTERS: IDENTIFIED VICTIMS AND EX POST DISCOUNTING

I argue that contractualists should not endorse the equal treatment principle. Instead, in competitive cases with equal harms, contractualism must prioritize definite standpoints over indefinite standpoints. This has significant implications for other areas of contractualist theory, particularly as it relates to how contractualism addresses the identified victim bias. Further, it may provide a new and plausible basis for *ex post* forms of discounting.

As noted, this argument is also relevant to the debate over the normative status of the identified victim bias and questions over prioritizing identified versus statistical lives.⁴⁴ The debate over identified versus statistical lives centers on whether we should give priority to saving *identified* lives over *statistical* lives. To borrow a well-known example, consider whether we should rescue a single trapped miner or spend the same amount to implement safety measures that would save some number of unidentified future miners.⁴⁵ The first is an identified life, whereas the latter are statistical lives. People generally demonstrate a durable "bias" in favor of identified lives, and this bias sometimes appears in the context of ethics and public policy, e.g., in the form of "the rule of rescue."⁴⁶

44 For an excellent book covering various interdisciplinary aspects of the identified victim bias, see Cohen et al., *Identified Versus Statistical Lives*.

45 For discussion of the ethical aspects of this case, see Cohen et al., "Statistical Versus Identified Persons." For a recent grounding of the identified victim bias in the notion of decreasing marginal moral value of survival chances, see also Steffánsson, "Identified Person 'Bias' as Decreasing Marginal Value of Chances."

46 Here, I do not intend 'bias' to be used pejoratively. 'Identified victim effect' is sometimes used to avoid the negative connotations of the term 'bias', but I use 'bias' as it is still the most common way to refer to the phenomenon at hand. Regarding the identified victim bias, see Jenni and Loewenstein, "Explaining the Identifiable Victim Effect." For critical

While many argue that the identified victim bias is morally suspect, the contractualist rejection of the equal treatment principle provides a limited justification for prioritizing identified lives.⁴⁷ The distinction between identified and statistical lives often tracks a more fundamental distinction: the certainty we have about whether a given standpoint is or will be filled. When we face an identified victim, we know with certainty that there is an actual individual whose reasons we must consider. With statistical lives, however, we face uncertainty not just about outcomes but about the very existence of individuals who could provide personal reasons. This concern maps directly onto the distinction between Gun and Box. Just as the victim in Gun is a definite standpoint representing an actual individual, identified victims are represented by definite standpoints. Conversely, just as the person in the opaque box in Box may or may not exist to support reasons of the right type, statistical lives represent indefinite standpoints—we do not know if there actually are or will be individuals capable of grounding reasons.⁴⁸

This in turn provides a novel grounding for the identified victim bias. Rather than arguing that identifiability itself has moral significance, my argument here suggests that what matters is our certainty about the standpoints used to ground reasons for rejection. Identifiability serves as a proxy for this more fundamental moral consideration—namely, our certainty that there is an actual individual whose reasons we must consider.⁴⁹ When discussed this way, we can offer a similar argument to the previous section: prioritizing identified lives helps us to avoid the risk of unjustifiably sacrificing an actual person for an empty standpoint.

This reframing has important implications. First, it suggests the preference for identified lives is not necessarily irrational—at least, not for committed

discussion of the rule of rescue, see Brock and Wikler, “Ethical Challenges in Long-Term Funding for HIV/AIDS.”

47 E.g., Hope, “Rationing and Life-Saving Treatments.”

48 This aspect is often overlooked in discussion of the identified victim bias, as statistical outcomes are typically stipulated in advance. For extended criticism of this aspect of the literature and the way it may distort our ethical reasoning, see Fried, *Facing Up to Scarcity*.

49 In a previous paper, I attempt to give an *ex ante* contractualist argument for the identified victim bias (“An Uncertainty Argument for the Identified Victim Bias”). For criticism of my argument, see Gilbertson, “Indifference, Indeterminacy, and the Uncertainty Argument for Saving Identified Lives.” The justification presented in this article is not exactly comparable to the one I offered previously, as the current article focuses on an *ex post* version of contractualism. However, the justification for the bias is not vulnerable to the same objection that Gilbertson raises, as Gilbertson’s criticism focuses on the use of the *principle of indifference* and the challenges it introduces to *ex ante* contractualism. This current *ex post* justification does not turn on the principle of indifference in the same way.

contractualists—as it reflects a meaningful moral distinction about the certainty of standpoints (and thus, reasons for rejection). Second, it provides guidance about when this preference is justifiable: specifically, it is justifiable when identification tracks (the probability of) existence rather than mere salience. This aspect is sometimes obscured in discussions of the ethical aspects of the identified victim's bias, because the loss of statistical lives is usually stipulated to occur with certainty, thus guaranteeing that the standpoint is or will be filled. In other words, the framing of these cases often obscures the underlying features that make the divide between identified and statistical lives morally relevant in a contractualist framework.⁵⁰

This more nuanced understanding helps explain both the intuitive appeal (for those who have such an intuition) and the limits of the identified victim bias. It suggests we should in fact give special weight to identified lives but only insofar as identification helps us track the existence of actual individuals. Contractualism holds that what ultimately matters is actuality—the existence of real individuals whose complaints we must take account of—rather than identifiability per se. This provides a principled basis for when and why we should prioritize identified over statistical lives.

Beyond implications for the identified victim bias, and perhaps more significantly for contractualism, this argument may provide a basis for an account of *ex post* discounting.⁵¹ The argument here focuses on cases where expected losses are equal, as these are the types of cases that the equal treatment principle seeks to address. However, additional questions arise in cases where the burdens differ significantly between definite and indefinite standpoints (e.g., weighing a small burden for a definite standpoint versus a large burden for an indefinite one). Resolving these cases may require developing a form of *ex post* discounting based on standpoint uncertainty. Such discounting could operate on the likelihood of the standpoint being filled. For example, in Box, this means discounting the complaint of death by one-sixth, whereas in Gun, there would be no discounting, as we know that the potential victim of Gun exists, and we know with certainty that the victim-of-gun standpoint is occupied. So we would compare death for the victim in Gun against one-sixth chance of death for the person in the opaque box. This is not a matter of discounting the strength of the complaint itself as is the case in a traditional *ex ante* approach; it instead pertains to the likelihood of justification being possible. There are challenging questions about this approach, particularly regarding how theoretically

50 Again, for discussion regarding this stipulation of statistical outcomes, see Fried, *Facing Up to Scarcity*.

51 For discussion of an alternative approach to *ex post* discounting, see Steuwer, "Contractualism, Complaints, and Risk."

justified such a move may be. However, here I simply aim to highlight this possibility, as developing a full account of discounting based on standpoints rather than complaints requires a more extensive treatment than I am able to provide here.

This form of discounting could be a plausible way of developing *ex post* contractualism and, importantly, could also be employed with *ex ante* contractualism. While *ex post* discounting makes the strength of the complaint conditional on the likelihood of the relevant standpoint being filled, *ex ante* discounting would discount *both* the complaint and the standpoint. In other words, all *ex ante* probabilities would need to be conditional on the probability of the relevant standpoint being filled. However, this requires a fairly complicated redevelopment of contractualism, and meaningful discussion of this possibility is far beyond the scope of this article. In short, the distinction between definite and indefinite standpoints developed in this article may be able to support the development of a plausible form of discounting for *ex post* contractualism, based on how certain we are that a given standpoint is occupied.⁵² This is an important avenue for future research for the development of contractualist approaches to risk and uncertainty.

5. CONCLUSION

I have argued that contractualism cannot comply with the principle of equal treatment for equal statistical loss. The basic problem with the principle of equal treatment is that contractualism is a relational moral theory and, as such, is concerned primarily with *justifiability to actual individuals* rather than with outcomes themselves. In other words, what matters is not just expected statistical loss but whether the reasons for rejecting a given principle or action are grounded in actual moral relationships. This focus on justifiability to actual others requires contractualism to recognize a distinction between definite and indefinite standpoints. When we know a standpoint is occupied—when we are certain that at least one actual individual is affected—contractualism requires us to take that standpoint's reasons into account. I have called these standpoints definite standpoints. In contrast, indefinite standpoints involve uncertainty and represent instances when we are unsure if there is anyone to whom the standpoint corresponds (as in the person in the opaque box in the Box case). This matters because if no one actually occupies the standpoint, then there is no individual to whom justification is owed, and any reasons arising

52 However, it is unlikely to be applicable in many (perhaps even most) cases, as it is likely rather rare that we have explicit probabilities for standpoints being full, at least outside of philosophy thought experiments.

from such a standpoint cannot be considered in contractualist deliberation. If we were to act on the reasons of an empty standpoint to the detriment of an actual person, we would be acting in an unjustifiable manner. To hold otherwise would require that contractualism either recognizes reasons that have to do with something other than the moral relationship or holds that we have a moral relationship with standpoints themselves.⁵³

This distinction means that in some cases, we can have equal expected statistical losses that are not morally equal because the probability in one pertains to a certainly existing person facing a probabilistic harm, and the probability in the other pertains to the probability of whether there is a person who stands to be harmed at all. Although these two kinds of cases may be probabilistically equal and may entail equal expectations of statistical loss, they are not morally equal. This is because prioritizing the second risks acting unjustifiably in a way that prioritizing the first never does. If the indefinite standpoint turns out to be unoccupied, it cannot provide reasons of the kind demanded by contractualism. As such, acting on the basis of such reasons against the interests of an actual person is unjustifiable. Given that contractualism requires acting justifiably, it also requires choosing actions that are certainly justifiable over actions that are possibly or riskily justifiable. As a result, I have argued that we should prioritize definite over indefinite standpoints in competitive cases with equal expected burdens or benefits. Thus, to avoid acting unjustifiably, contractualist must at least sometimes prioritize definite over indefinite standpoints, which means that in at least some cases, contractualists must reject the equal treatment principle.

This argument has broader implications beyond the need to reject or revise the equal treatment principle. First, it provides a new, *ex post* contractualist justification for the identified victim bias. The preference for identified over statistical lives is frequently dismissed as irrational or otherwise without ethical justification, but the distinction between definite and indefinite standpoints and the contractualist requirement to treat them differently may provide a moral justification for the bias. It is not identifiability per se that matters but rather that identified lives involve reasons that we are certain we must consider. Second, this argument lays groundwork for a possible form of *ex post* discounting. If the normative status of a standpoint depends on whether it is occupied, then one way to approach certain cases of risk and uncertainty is to discount the reasons of indefinite standpoints relative to definite standpoints. This approach to discounting could apply to both *ex post* and *ex ante* versions of contractualism. Further work is needed to explore how exactly

53 Gibb, "Relational Contractualism and Future Persons."

such a discounting approach might be developed. Finally, while this article has focused exclusively on *ex post* contractualism, a full exploration of how the definite/indefinite standpoint distinction operates within an *ex ante* framework remains an important direction for future research.

Ultimately, what this argument shows is that equal expected statistical losses are not always morally equal. When we determine what is justifiable, we must look beyond expected statistical losses and consider the underlying moral relationships and how they influence the justifiability of inflicting such expected statistical losses. If contractualism aims to remain a theory that takes the moral relationship between individuals seriously, then in competitive cases, it must prioritize the interests of definite standpoints over indefinite standpoints, and thus, it must reject the equal treatment principle.⁵⁴

Jagiellonian University
jay@sdu.dk

REFERENCES

- Ashford, Elizabeth, and Tim Mulgan. "Contractualism." *Stanford Encyclopedia of Philosophy* (Summer 2018). <https://plato.stanford.edu/archives/sum2018/entries/contractualism/>.
- Bourguignon, Léa. "On the Possibility of Act Contractualism." *Australasian Journal of Philosophy* 102, no. 2 (2024): 1–19.
- Brock, Dan W., and Daniel Wikler. "Ethical Challenges in Long-Term Funding for HIV/AIDS." *Health Affairs* 28, no. 6 (2009): 1666–76.
- Cohen, I. Glenn, Norman Daniels, and Nir M. Eyal, eds. *Identified Versus Statistical Lives: An Interdisciplinary Perspective*. Population-Level Bioethics Series. Oxford University Press, 2015.
- . "Statistical Versus Identified Persons: An Introduction." In Cohen, Daniels, and Eyal, *Identified Versus Statistical Lives*.
- Copp, David. "The Rule Worship and Idealization Objections Revisited and Resisted." *Oxford Studies in Normative Ethics* 10 (2020): 131–55.

54 I would like to thank Callum MacRae and Christoph Merdes for valuable discussion of some of the ideas in this article, and Anastasiia Babash for helpful comments on an earlier draft of this article. This research is part of project no. 2022/47/P/HS1/02511, co-funded by the National Science Centre and the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement no. 945,339. For the purpose of open access, the author has applied a CC-BY public copyright license to any author accepted manuscript (AAM) version arising from this submission.

- Darwall, Stephen. *The Second-Person Standpoint: Morality, Respect, and Accountability*. Harvard University Press, 2006.
- Eyal, Nir. "Concentrated Risk, the Coventry Blitz, Chamberlain's Cancer." In Cohen, Daniels, and Eyal, *Identified Versus Statistical Lives*.
- Frick, Johann. "Contractualism and Social Risk." *Philosophy and Public Affairs* 43, no. 3 (2015): 175–223.
- Fried, Barbara H. *Facing Up to Scarcity: The Logic and Limits of Nonconsequentialist Thought*. Oxford University Press, 2020.
- Gibb, Michael. "Relational Contractualism and Future Persons." *Journal of Moral Philosophy* 13, no. 2 (2016): 135–60.
- Gilbertson, Eric. "Indifference, Indeterminacy, and the Uncertainty Argument for Saving Identified Lives." *Journal of Applied Philosophy* 41, no. 3 (2024): 480–97.
- Gillies, Donald. *Philosophical Theories of Probability*. Routledge, 2006.
- Hayenhjelm, Madeleine, and Jonathan Wolff. "The Moral Problem of Risk Impositions: A Survey of the Literature." *European Journal of Philosophy* 20, no. 1 (2012): 26–51.
- Hope, Tony. "Rationing and Life-Saving Treatments: Should Identifiable Patients Have Higher Priority?" *Journal of Medical Ethics* 27, no. 3 (2001): 179–85.
- Jenni, Karen, and George Loewenstein. "Explaining the Identifiable Victim Effect." *Journal of Risk and Uncertainty* 14, no. 3 (1997): 235–57.
- John, Stephen David, and Emma J. Curran. "Costa, Cancer and Coronavirus: Contractualism as a Guide to the Ethics of Lockdown." *Journal of Medical Ethics* 48, no. 9 (2022): 643–50.
- Katz, Corey. "Contractualism, Person-Affecting Wrongness and the Non-Identity Problem." *Ethical Theory and Moral Practice* 21, no. 1 (2018): 103–19.
- Kim, Suzie. "On the Need for Real Dialogue: What's Wrong with Monological Contractualism?" *European Journal of Philosophy* 27, no. 4 (2019): 939–56.
- Kumar, Rahul. "Risking and Wronging." *Philosophy and Public Affairs* 43, no. 1 (2015): 27–51.
- Luce, R. Duncan, and Howard Raiffa. *Games and Decisions: Introductions and Critical Survey*. Wiley, 1957.
- Martin, Desa Valeska. "Navigating Nonidentity: Scanlonian Contractualism and Types of Persons." *Journal of Ethics and Social Philosophy* 29, no. 1 (2024): 86–106.
- Muñoz, Daniel. "Each Counts for One." *Philosophical Studies* 181, no. 10 (2024): 2737–54.
- Oberdiek, John. *Imposing Risk: A Normative Framework*. Oxford University Press, 2017.

- Otsuka, Michael. "Risking Life and Limb." In Cohen, Daniels, and Eyal, *Identified Versus Statistical Lives*.
- . "Scanlon and the Claims of the Many Versus the One." *Analysis* 60, no. 3 (2000): 288–93.
- Parfit, Derek. *On What Matters*, vol. 2. Oxford University Press, 2011.
- Peterson, Martin. *An Introduction to Decision Theory*. Cambridge University Press, 2017.
- Reibetanz, Sophia. "Contractualism and Aggregation." *Ethics* 108, no. 2 (1998): 296–311.
- Rüger, Korbinian. "On *Ex Ante* Contractualism." *Journal of Ethics and Social Philosophy* 13, no. 3 (2018): 240–59.
- Salein, Valentin. "Leaving Principle Contractualism Behind? A Response to Salomon." *Journal of Ethics and Social Philosophy* 30, no. 1 (2025): 146–54.
- Saunders, Ben. "A Defence of Weighted Lotteries in Life Saving Cases." *Ethical Theory and Moral Practice* 12, no. 3 (2009): 279–90.
- Scanlon, T. M. *What We Owe to Each Other*. Belknap Press, 1998.
- Sheinman, Hanoch. "Act and Principle Contractualism." *Utilitas* 23, no. 3 (2011): 288–315.
- Smart, John Jamieson Carswell, and Bernard Williams. *Utilitarianism: For and Against*. Cambridge University Press, 1973.
- Southwood, Nicholas. *Contractualism and the Foundations of Morality*. Oxford University Press, 2013.
- . "Moral Contractualism." *Philosophy Compass* 4, no. 6 (2009): 926–37.
- Spiekermann, Kai. "Good Reasons for Losers: Lottery Justification and Social Risk." *Economics and Philosophy* 38, no. 1 (2022): 108–31.
- Stefánsson, H. Orri. "Identified Person 'Bias' as Decreasing Marginal Value of Chances." *Noûs* 58, no. 2 (2024): 536–61.
- Steuwer, Bastian. "Contractualism, Complaints, and Risk." *Journal of Ethics and Social Philosophy* 19, no. 2 (2021): 111–39.
- Suikkanen, Jussi. "*Ex Ante* and *Ex Post* Contractualism: A Synthesis." *Journal of Ethics*, 23, nos. 3–4 (2019): 391–410.
- Tadros, Victor. "Controlling Risk." In *Prevention and the Limits of the Criminal Law*, edited by Andrew Ashworth, Lucia Zedner, and Patrick Tomlin. Oxford University Press, 2013.
- Taurek, John M. "Should the Numbers Count?" *Philosophy and Public Affairs* 6, no. 4 (1977): 293–316.
- Wasserman, David. "Let Them Eat Chances: Probability and Distributive Justice." *Economics and Philosophy* 12, no. 1 (1996): 29–49.
- Zameska, Jay. "An Uncertainty Argument for the Identified Victim Bias." *Journal of Applied Philosophy* 39, no. 3 (2022): 504–18.