



AGGREGATION, BENEFICENCE AND CHANCE

BY TOM DOUGHERTY

JOURNAL OF ETHICS & SOCIAL PHILOSOPHY

VOL. 7, No. 2 | MAY 2013

URL: WWW.JESP.ORG

COPYRIGHT © TOM DOUGHERTY 2013

Aggregation, Beneficence and Chance

Tom Dougherty

IT IS PLAUSIBLE TO THINK THAT IT IS WRONG to cure many people's headaches rather than save someone else's life. On the other hand, it is plausible to think that it is not wrong to expose someone to a tiny risk of death when curing this person's headache. I will argue that these claims are inconsistent. For if we keep taking this tiny risk, then it is likely that one person dies, while many others' headaches are cured. In light of this inconsistency, there is a conflict in our intuitions about beneficence and chance. This conflict is perplexing. And I have not been able to find a satisfactory way of resolving it. Perhaps you can do better?

1. Intuitive Opposition to Aggregation

I will begin by fleshing out the two claims that I will argue are inconsistent. The first claim concerns decisions about whom to help.¹ Consider the following case:

ISLANDS. If you go to the north island, then you can save Jones's life. If you go to the south island, then you can cure the headaches of a billion people. You only have enough fuel to go to one island.²

You have two options. First, you can provide a large benefit to Jones by saving Jones's life. Second, you can provide a small benefit to a billion people by curing their headaches. Most of us have the intuition that you ought to save Jones's life. Moreover, you ought to do so, no matter how many people's headaches you could otherwise cure.

What principle would explain the fact that you must save Jones? A natural candidate is the following:

Many-Few. You may not provide small benefits to many people rather than save the life of someone else. (All else being equal.)

When I say you may not provide these benefits, "all else being equal," I am holding fixed features of these people, such as their relationship to you. I am also assuming that your behavior affects only the people who could receive

¹ I am assuming that the potential beneficiaries do not have claim-rights against you, and that you do not act in any institutional capacity.

² This case is inspired by Scanlon's classic case:

WORLD CUP. A worker has fallen onto electric cables in the television studio and is in agony. If you turn off the power, billions of people miss the climax of the World Cup final.

the small benefits and the person whose life you could save. I am also assuming that the saved person would continue to live for a reasonably long period of time.³

If Many-Few is correct, then we should reject explanations that endorse a utilitarian approach to beneficence. According to this approach, it is always possible that small benefits to sufficiently many people “aggregate” to form a total amount of welfare that is more important than Jones’s survival. If you ought to save Jones in ISLANDS, then the utilitarian approach to beneficence is misguided.⁴

2. Intuitive Tolerance of Risk-Taking

The second view concerns risk-taking. Sometimes, we take minuscule risks that someone will die in order to provide this person with a small benefit. To spare Amanda the inconvenience of taking the subway, I might drive her to the airport. I do so, even though this exposes her to a slightly increased chance of death. (I will ignore the risks to pedestrians and other motorists.) For Ben’s culinary pleasure, I might serve him raw fish. I do so, even though cooking the fish would slightly reduce his risk of dying from a fatal bacterium. Do we sometimes kiss people for their sake? If so, we should pay heed to James Lenman’s observation that:

it seems morally pretty unobjectionable to kiss one’s lover in circumstances where there is (and I guess there *always* is) a fantastically small chance that, for example, one has contracted some deadly transmissible disease which one thereby passes on to them.⁵

It seems a kiss is worth the coyest flirtation with death.

What would explain the fact that we may take these risks? A natural candidate is the following principle:

Risk Tolerance. You may expose someone to a negligible risk of death in order to otherwise provide this person with a small benefit. (All else being equal.)

Again, the “all else being equal” clause is in place to restrict the scope of the principle – the clause ensures that the principle has bite only in a case where the only morally significant feature of your action is that it has these chances of harms and benefits. I briefly postpone for now discussion of what difference it makes whether this risk is metaphysical or doxastic in nature. (The quick answer? It makes little difference.)

³ Given the inevitability of death, all we can do is to extend our brief stays on this earth. If the length of the extension is incredibly short, we may prefer to provide many small benefits to extending one person’s life by this amount.

⁴ This is the point that Scanlon makes about his WORLD CUP example, described in footnote 2. He offers this case as a counterexample to aggregative theories.

⁵ James Lenman (2008) “Contractualism and Risk Imposition,” *Politics Philosophy Economics* 7: 99-122, at 112.

3. The Repetition Argument

We have seen two claims, Many-Few and Risk Tolerance, which are initially plausible. Now I will offer an argument that concludes that they are inconsistent. This argument turns on the likely effects of repeating risky actions. It is helpful to consider the repetition of risks when these risks are small. This allows us to put these small risks “under the moral microscope,” to use Derek Parfit’s apt phrase.⁶ Parfit’s preferred microscope is universalization: He magnifies a small effect of an action by considering *everyone* performing the action. In addition, we can magnify a small effect of an action by considering *you* performing the action many times: Another way of putting a risk under the moral microscope would be to consider your repeating the risk.

The main argument of this paper turns on the likely effects of repeating the risk. In this respect, my argument will follow the same broad strategy employed by Alastair Norcross.⁷ Norcross argues as follows:

P1 Other things being equal, it is better that one person die than that five million each incur a one in a million risk of dying (153).

P2 Other things being equal, it is better that five million people each incur a one in a million risk of dying than that each suffers a moderate headache for twenty-four hours (156).

C Therefore, it is better that one person die a premature death than that five million people each suffers a moderate headache for twenty-four hours.

And if we accept the conclusion of this argument, then we must reject the thesis that Norcross is targeting:

Worse: Other things being equal, it is worse that one person die a premature death than that *any number* of people suffer moderate headaches for twenty-four hours (152).

Norcross’s target, Worse, is a close cousin of the claim Many-Few: While Many-Few is a claim about permissibility, Worse is a parallel claim about which outcomes are better than others. Norcross argues against his target as follows. Norcross defends P1 by appealing to our intuition that one ought to direct a deadly gas so that it certainly kills one person, instead of exposing 5 million people to a one-in-a-million risk of death in the following way: Although the “expected utility” of the action is equivalent to a situation in which five people among the 5 million certainly die, “there is a finite, but small (less than one percent), chance that no one will die.” (153). Meanwhile, Norcross defends P2 by appealing to our intuition that it is permissible for a single

⁶ Derek Parfit (2011) *On What Matters*, Oxford: Oxford University Press.

⁷ Alastair Norcross (1998) “Great Harms from Small Benefits Grow: How Death Can Be Outweighed by Headaches,” *Analysis* 58: 152-58.

individual to take a tiny risk of death by driving to the pharmacy for a pain-killer. The intuition to which Norcross is appealing here is the first-personal analogue to Risk Tolerance: While Risk Tolerance is the thesis that it is permissible to expose a stranger to a tiny fatal risk for a small benefit to her, Norcross's approach appeals to the claim that a better outcome results when these individuals take these risks for themselves. Norcross then "scales up" this risk by imagining 5 million individuals separately driving to the pharmacy. On the grounds that it is better for them to do so, he argues that we should accept P2.

The background idea that animates Norcross's argument is, I suggest, an important insight for our topic. However, the central inference in Norcross's argument presupposes a consequentialist approach to the ethics of risk and aggregation. Norcross's premises and conclusion are claims about which outcomes are better or worse than other outcomes, and he justifies his central inference by appealing to the "transitivity of 'better than'" (157). This inference will be happily accepted by nearly all consequentialists. But it is an inference that will be rejected by many non-consequentialists. They will respond that attempts to translate all ethical features of situations into a single metric of "goodness" is an illegitimate strategy, in general. Consequently, they are unlikely to be persuaded, by means of this strategy, that claims like Many-Few and Risk Tolerance are inconsistent. This potential opposition is dialectically significant because many of the friends of Many-Few are non-consequentialists. Indeed, as we noted earlier, it is claims such as Many-Few that some non-consequentialists see as the key grounds for rejecting an aggregative teleological approach to ethics.⁸

My aim is to show that the broader strategy of considering the repetition of risks can be used to formulate an argument against which non-consequentialists should have no complaint. This argument will not appeal to the transitivity of "better than"; indeed, it will not even suppose that there is a property of a "better or worse" outcome – a property that some hard-nosed non-consequentialists deny.⁹ The argument will nonetheless follow the same insight that lies behind Norcross's argument by considering the repetition of a tiny risk. As such, I will call it the "Repetition Argument." It starts with a premise entailed by Risk Tolerance, and reaches a conclusion that contradicts Many-Few. I should stress at the outset that I intend this argument to show only that Risk Tolerance and Many-Few are *inconsistent*. I take no stance on which thesis we should reject. In this respect, my argument is less ambitious than that of Norcross, who *affirms* an analogue of Risk Tolerance in order to persuade us to *reject* an analogue of Many-Few.

⁸ See, for example, Scanlon (1998).

⁹ See, for example, Judith Jarvis Thomson (1993) "Goodness and Utilitarianism," *Proceedings and Addresses of the American Philosophical Association* 67: 145-59.

To run my argument, I will need an example of a risky action that you would be permitted to perform, according to Risk Tolerance. I will choose the following toy case:

POISON. In a nearby room, Victim has swallowed poison. If you do nothing, then Victim will have a headache and then die. You can bring Victim only one of two antidotes:

The Reliable Antidote is certain to save Victim, but will do nothing for Victim's headache.

The Chancy Antidote is most likely to cure Victim's headache and save Victim's life, apart from a one-in-a-billion chance that it does nothing.

If you choose the Chancy Antidote, then you are most likely to relieve Victim's headache, but take a tiny risk that you fail to save Victim's life. I assume that if the situation described in POISON occurs in isolation, then you may bring the Chancy Antidote to Victim, according to Risk Tolerance. I invite you to alter the size of the risk as you see fit so that you are satisfied that Risk Tolerance entails that you may bring the Chancy Antidote. For example, if you feel that a one-in-a-billion chance is too large, then make it smaller. The numbers that I use will have no substantial effect on the argument that I offer. So I will assume that if we accept Risk Tolerance, we should accept the following:

1. You may bring the Chancy Antidote in POISON once.

This will be the first premise in the argument.

Now, suppose the situation described in POISON occurs again and again. On the first day, Smith is in the position of Victim, then the next day Roberts is in this position, and so on. If you may bring the Chancy Antidote to Smith, then you may also bring the Chancy Antidote to Smith and then bring it to Roberts on the following day. If the risk of bringing the Chancy Antidote is such that you may bring it to one person, then the risk is such that you may bring it to one person and then another. Indeed, we could suppose that you face a very long sequence of people in exactly the same situation. Assuming you may bring the Chancy Antidote to each person alone, you may bring it to each of them in turn:

2. If you may bring the Chancy Antidote in POISON once, then you may bring the Chancy Antidote to any sequence of people in POISON.

From these two premises, we can infer:

3. *Therefore*, you may bring the Chancy Antidote to any sequence of people in POISON. (From 1 and 2.)

Now we should ask what will happen if you do bring the antidote to many people. As you increase the number of times that you repeat a risk that

someone might die, it becomes increasingly likely that one of these repetitions results in a death. Let us bring this point to bear on POISON. Suppose this situation repeats with different people playing the role of Victim. Indeed, suppose it occurs 1 billion times, and you bring the Chancy Antidote each time. Among these 1 billion repetitions, we would expect that one time you will fail to save someone's life. Moreover, the "Law of Large Numbers" entails that, as the number of times that you bring the Chancy Antidote increases, it becomes increasingly likely that you will fail to save someone's life one in every billion times.¹⁰ The law entails that it is near certain that this distribution will result in a sufficiently large, finite sequence. Note that I really do mean *near certain*. If we make this sequence of people large enough, then we can bring the chance as close to certainty as we like.¹¹ If we made the sequence long enough, then the probability could be 0.99999999, for example. Thus:

4. If you bring the Chancy Antidote to a sufficiently large sequence of people, it is near certain that you will fail to save the life of one in every billion people to whom you bring the Chancy Antidote and cure the headaches of the rest. (From the Law of Large Numbers.)

Notice that these repetitions do not increase the chances of your failing to save any particular person – this chance remains one in a billion. To think otherwise would be to commit the gambler's fallacy.

Continuing the argument, we can infer:

5. *Therefore*, you may make it near certain that you fail to save the life of one in every billion people to whom you bring the Chancy Antidote in POISON and cure the headaches of the rest. (From 3 and 4.)

That is, repeating these risks would be likely to result in a particular distribution of benefits. The distribution is that for every billion people to whom you bring the Chancy Antidote, one person dies, while all the rest have their lives saved and their headaches cured. By always bringing the Reliable Antidote, another distribution would result. This distribution is that, for every billion people to whom you bring the Chancy Antidote, all have their lives saved but suffer headaches. The comparative difference between these distributions is that in the former, 999,999,999 people enjoy a small benefit of having their headaches cured, while one person suffers the large loss of death:

¹⁰ Roughly, the law says that as the number of repetitions of the risky action tends to infinity, the probability that the resulting distribution is one death in every billion tends to one (certainty). For our purposes, the differences between the weak and strong versions of the law do not matter.

¹¹ Indeed, the Law of Large Numbers entails that if you repeated this action *infinitely* many times, then statistically you would be guaranteed to fail to save someone's life one in every billion times – the probability of this pattern emerging is 1.

6. *Therefore*, you may make it near certain that you cure the headaches of 999,999,999 people rather than save the life of someone else. (From 5.)

This is the consequence of your being permitted to repeat the risky action.

This brings us to another substantive premise in the Repetition Argument:

7. If it is wrong to bring about a particular distribution, then it is wrong to make it near certain that this distribution obtains.

By contraposition, if it is permissible to make it near certain that the distribution obtains, then it is permissible to bring about the distribution. This premise is an important step in the argument as I offer it here.¹² But it strikes me as very plausible, and I cannot see any reason for denying it. Any reason not to bring about a distribution would equally seem to be a reason not to make the distribution near certain. Recall that I really mean *near certain* here. The likelihood of this distribution arising can be as close to certainty as we like. I suggest that we accept this premise unless we are given some reason not to do so. I cannot see any such reason. So I conclude that we ought to accept it.

Accepting this premise puts us in a position to infer:

¹² In addition, there are two ways of running a similar argument that appeal to weaker premises. First, Richard Holton has pointed out to me that an analogous argument could be run that relies on this premise instead:

7*. If it is wrong to bring about a particular distribution, then there is a probability, p , such that it is wrong to make the probability of this distribution obtaining p .

This would require some tweaks to the other premises in the argument that strike me as benign, but would be complicated to write out. For example, we would have to replace references to “near certainty,” with references to this probability p . Stating this would require some complexity in specifying the quantifiers that bind each premise, which is why I avoid doing so here.

The second way of altering the argument begins by strengthening premise 2 to 2*:

2*. If you may bring the Chancy Antidote in POISON once, then you may bring the Chancy Antidote to any sequence of people in POISON.

By the Law of Large Numbers:

4*. If you bring the Chancy Antidote to an infinite sequence of people, there is a probability of 1 that you will fail to save the life of one in every billion people to whom you bring the Chancy Antidote and cure the headaches of the rest.

Then we would only need to appeal to the weaker premise 7**:

7**. If it is wrong to bring about a particular distribution, then it is wrong to make it the case that there is a probability of 1 that this distribution obtains.

This premise strikes me as undeniable.

8. *Therefore*, you may cure the headaches of 999,999,999 people rather than save the life of someone else. (From 6 and 7.)

But this conclusion in line 8 contradicts the thesis that we encountered earlier:

Many-Few. You may never provide small benefits to many people rather than save the life of someone else. (All else being equal.)

So by starting with an assumption entailed by Risk Tolerance, we have ended with a conclusion that contradicts Many-Few.

Let me summarize the Repetition Argument:

1. You may bring the Chancy Antidote in POISON once. (From Risk Tolerance.)
2. If you may bring the Chancy Antidote in POISON once, then you may bring the Chancy Antidote to any sequence of people in POISON.
3. *Therefore*, you may bring the Chancy Antidote to any sequence of people in POISON. (From 1 and 2.)
4. If you bring the Chancy Antidote to a sufficiently large sequence of people, it is near certain that you will fail to save the life of one in every billion people to whom you bring the Chancy Antidote, and cure the headaches of the rest. (From the Law of Large Numbers.)
5. *Therefore*, you may make it near certain that you will fail to save the life of one in every billion people to whom you bring the Chancy Antidote in POISON and cure the headaches of the rest. (From 3 and 4.)
6. *Therefore*, you may make it near certain that you will cure the headaches of 999,999,999 people rather than save the life of someone else. (From 5.)
7. If it is wrong to bring about a particular distribution, then it is wrong to make it near certain that this distribution obtains.
8. *Therefore*, you may cure the headaches of 999,999,999 people rather than save the life of someone else. (From 6 and 7.)

Since line 1 is entailed by Risk Tolerance, and Many-Few contradicts line 8, the argument shows that we cannot hold both principles.

4. Defending the Iterative Premise

To resist the argument, one would have to drive a wedge between the permissibility of taking a risk and the permissibility of bringing about a distribution that is near certain to obtain as a result of repeating the risk. A natural way of doing so would be to deny the iterative premise:

2. If you may bring the Chancy Antidote in POISON, then you may bring the Chancy Antidote to any sequence of people in POISON.

One might object that you may take a risk occasionally, but that you may not repeat the risk so many times that there is a high chance that someone denies. For example, Derek Parfit claims:

We can usually ignore a very small chance. But we should not do so when we may affect a very large number of people, or when the chance will be taken a very large number of times. These large numbers roughly cancel out the smallness of the chance.¹³

Suppose we applied Parfit's reasoning to POISON. If you faced the case only once, then you may bring the Chancy Antidote since you can "ignore a very small chance." But you may not bring the Chancy Antidote to many people since then "the chance will be taken a very large number of times." As a result, I would expect that Parfit may be inclined to deny the iterative premise.

This line of thought may be tempting, but it springs from a mistake that we often make when thinking about risks. In evaluating a set of risky actions, we sometimes make a mistake of not appreciating the importance of the *independence* of the risks in question. Here is an example of the same mistake.¹⁴ Suppose you leave your house and worry that you have left the gas on. You judge that there is a very small risk that you have. It is tempting to hold two thoughts. First, it is fine for you to take this risk *just this once*. The convenience of not going back to check is worth the risk. Second, it would not be fine for you to take this risk very many times. Were you to do so, it would become likely that one time you would have left the gas on, and this would not be worth many small conveniences. Conjoining these two thoughts, we might say that it is fine for you to tempt fate now and again, but that you may not tempt fate too much.

However, this is a mistaken way to think about risks. Suppose it is Tuesday and you are considering whether or not to go back to check the gas. Whether you may continue on your way without checking the gas depends only on the effects associated with this particular action – do the risks of harms and benefits associated with continuing without checking justify your doing so? Whether you may continue does not depend on whether you continued last week, or the week before. Nor does it depend on whether you go on to do so in the future. All that matters when deciding whether to go back to check on the gas on any particular occasion is the risk associated with this particular occasion. So if the sum of the conveniences is not worth the sum of the risks, then each convenience is not worth each risk.¹⁵

This point applies with equal force to POISON. The correct way to think about the risks is to look at the possible or actual effects of each action independently, and assess the permissibility of each action on these grounds. The permissibility of exposing someone to a risk does not depend on whether you have previously exposed others to this risk. The following principle is a plausible way of capturing this point:

¹³ Derek Parfit (1984) *Reasons and Persons* (reprinted with further corrections), New York: Oxford University Press, p. 75.

¹⁴ Thanks to Richard Holton for suggesting this type of example.

¹⁵ I assume here that the conveniences have constant marginal utility for you.

Only Effects Matter. When only one person is affected by a risky action, the permissibility of the action depends on the possible or actual effects of *that action* for *that person*.¹⁶

Now, this principle entails the iterative premise. Consider a sequence of people, Marks, Roberts and Smith. May you bring the Chancy Antidote to each in sequence? This question would depend only on the risks involved for each person, individually. If bringing the Chancy Antidote does not create too large a risk of death for Marks, Roberts and Smith individually, then it must be permissible for you to bring it to all of them in sequence. Conversely, if you are forbidden from bringing the Chancy Antidote to all in sequence, then this could only be because the risk for some individual was too high.

In making these claims, I assume we are dealing only with cases where the magnitudes of the risks, benefits and harms are independent of each other. This assumption has two parts. First, I assume that the size of the benefits and harms are the same, and that these are the only morally relevant features associated with taking these risks. For example, I assume, plausibly, that whether Marks dies does not affect how bad it is for Roberts to die. Moreover, I assume that the only morally relevant feature of their deaths is their well-being. That is, I assume that we are not dealing with a case where there are further undesirable effects that arise from them both dying. For example, I assume that we are not dealing with a case in which they are the last two remaining males in the species. In this case, there would be a further morally relevant consequence of both dying – no one would be able to procreate to continue the human race. Second, I assume that the risks are statistically independent. For example, whether you give Marks the Reliable or Chancy Antidote makes no difference to the size of the risk that you would expose Roberts to by bringing Roberts either antidote. Things would be different if the risks were cumulative:

SINKING SHIP. A ship is sinking with many children on board. There are two lifeboats that you can direct them to.

The Reliable Lifeboat is a large and sturdy boat. It is very safe and certain to reach the shore. It will take a day to reach land, by which time any children on board will be tired, hungry and cranky.

The Chancy Lifeboat is a small and light boat. It will reach the shore in an hour. However, there is a one-in-one-hundred-million chance that it will sink if it has one

¹⁶ Things are more complicated when there are causal relations between the risks to different individuals. For example, there are cases where each individual faces a small risk of suffering a harm, but where it is guaranteed that *some* unspecified individual will suffer this harm. The risks to each individual are interconnected because it is impossible that every individual avoids the harm. Consequently, it is impossible that an individual avoids the harm when every other individual has avoided the harm. Lina Eriksson discusses this more complicated issue in her unpublished manuscript “Probabilities and Prohibitions.”

child passenger on board. Furthermore, for each additional child that is on board, the probability of it sinking doubles.

Let us suppose, for the sake of argument, that it is permissible to take a risk of a child's death in order to spare this child several hours' crankiness if and only if the risk of death is below one in 30 million. (Again, my choice of these numbers is arbitrary, so you should feel free to change them. For our purposes, all that matters is that the risk is cumulative, and there is a threshold of risk that you may take in pursuit of this benefit.) Now, whether you have directed toddlers Timmy and Sammy to the Reliable Lifeboat matters a great deal to whether you may direct toddler Vicky to this boat. For if both Timmy and Sammy are the only children on board the Reliable Lifeboat, then the probability of the boat sinking is one in 50 million. Were Vicky to board, this probability would rise to one in 25 million – an unacceptably high risk. But if neither Timmy nor Sammy are on board, nor any other child, then the probability of the boat sinking with Vicky on board is merely one in 100 million – a risk that we have stipulated to be acceptably low in light of the benefit to her.

In the SINKING SHIP case, you may not repeat the risky action of directing a child to the Chancy Boat. This is because the risks are cumulative: Exposing one child to a risk increases the risk that you would expose the next child to. Still, it is the risk of an action to each child that affects the permissibility of performing this action. In the POISON case, you may repeat the risky action (if you may perform it at all), because it does not affect the risks that you expose other individuals to. Again, all that matters is the risk that each person faces.

5. A First Possible Solution: Distinguishing Subjective and Objective Risks

Up until now, I have postponed discussing the issue of whether the risks in question are metaphysical or doxastic in nature. I will argue that it makes little difference, if any, to the significance of the previous discussion.

Before I can make this point, I need to set up some machinery. Let me begin by invoking the orthodox distinction between objective chance and credence.¹⁷ Objective chance is metaphysical in nature. Objective chance is surprisingly difficult to define,¹⁸ but we do have an intuitive grip on the notion: We have an intuitive grip on the notion that a coin toss is *itself*

¹⁷ The locus classicus is David Lewis (1980) "A Subjectivist's Guide to Objective Chance," in Richard C. Jeffrey, ed., *Studies in Inductive Logic and Probability*, Berkeley: University of California Press, pp. 263-93.

¹⁸ Take the definition that an event is objectively chancy when previous events do not determine whether it occurs. Even this definition is controversial, since some people hold that statistical mechanics involves objective chances, even on the assumption that determinism is true. See Barry Loewer (2001) "Determinism and Chance," *Studies in History and Philosophy of Modern Physics* 32B(4): 609-20.

chancy, regardless of our attitudes toward the toss. Metaphysical chance is different from doxastic credence. Your credence in a proposition is your degree of belief that this proposition is true. We can illustrate the difference between objective chance and credence with an example. Suppose you do not know whether there is an even or odd number of words in yesterday's newspaper. You are equally inclined to believe that the number is even or odd. In fact, there was an odd number. The objective chance of there being an even number is zero. But your credence in the proposition that there is an even number is half.

In light of the distinction between objective chance and credence, let us distinguish objective risk from subjective risk. An action has an objective risk of a harm just in case there is an objective chance that the action leads to this harm. By contrast, an action carries a subjective risk of a harm from your point of view just in case on the basis of your evidence you ought to assign some credence to the proposition that the action leads to this harm. We can use the previous example to illustrate the difference between objective risk and subjective risk. Suppose that you are offered a bet: You win \$10 if there was an odd number of words in yesterday's newspaper, and lose \$10 if there was an even number. Again suppose there was an odd number. Taking this bet has no objective risk, as there is no objective chance that there was an even number – there is no objective chance you can lose the bet. But in light of your evidence, you should have a credence of 0.5 in the proposition that there was an even number of words. Therefore, from your point of view, taking the bet exposes you to a subjective risk of losing \$10.

Now consider Risk Tolerance. When we distinguish objective and subjective risk, we can ask which type of risk we tolerate, and make the thesis more specific, accordingly. This principle was motivated by our intuitions about everyday actions. We intuitively judge that we may give people lifts in a car, that we may serve them raw fish and that we may kiss them (although hopefully we do not perform these actions at the same time). Our intuitions about these everyday actions are, I think, based on the assumption that the actions in question are objectively risky. So I take it that these intuitive judgments support this thesis:

Objective Risk Tolerance. You may expose someone to a negligible objective risk of death in order to otherwise provide this person with a small benefit. (All else being equal.)

So our intuitions support Objective Risk Tolerance. The Repetition Argument shows that this thesis is inconsistent with Many-Few. To satisfy yourself of this, suppose that the risks of bringing the Chancy Antidote in POISON are objective risks. If you repeat these objective risks enough times, you will near certainly end up with a distribution of many small benefits and one death. Appealing to this consideration, the Repetition Argument finds that Objective Risk Tolerance is inconsistent with Many-Few.

Therefore we cannot hold both Objective Risk Tolerance and Many-Few. Suppose we give up Objective Risk Tolerance. Our naive view of everyday risky actions is that they are objectively risky. If the naive view is correct, and we reject Objective Risk Tolerance, then we lack a justification for these everyday actions. But whether the naive view is correct will turn on some tricky issues in physics and metaphysics. Consider the following two claims. First, the effects of these everyday actions are determined. Second, if an event is determined, then it is not objectively chancy. Presently, there seems to be little consensus among the experts concerning whether each of these claims is true.¹⁹ For the sake of argument, let us suppose that they both are true. In this case, everyday actions are merely subjectively risky.

This might seem to present a solution to the problem for us. The Repetition Argument shows that we cannot hold Objective Risk Tolerance and Many-Few. The reason why we want to hold onto Objective Risk Tolerance is to justify everyday risky actions like giving people lifts in a car. We could reject Objective Risk Tolerance and still hold that these everyday actions are permissible because the actions are merely subjectively risky.²⁰ Call this the Subjective Risk Solution.

Alas, this solution does not work. In addition to the permissibility of an action, we should also consider whether we would be blameworthy for performing the action. The solution is no defense against the charge that we are blameworthy for performing these everyday risky actions. Let me say why.

When people correctly apportion their credences on the basis of their available evidence, the blameworthiness of their actions depends on their evidence.²¹ To see this, suppose Alf buys a packet of sugar, which he adds to someone's tea. The packet he buys contains an untraceable poison. Still, he acts blamelessly. This is because his evidence was that the sugar was harmless, and he based his beliefs on the evidence.

Because the blameworthiness of an action turns on an agent's *evidence*, an agent can be blameworthy for performing a subjectively risky action, even if the action is not objectively risky. Now, consider the following principle:

¹⁹ As things stand, it seems that the experts disagree about (i) whether the correct interpretation of quantum mechanics is indeterministic; (ii) whether indeterminism could affect macroscopic events like human actions; (iii) whether actions could be objectively chancy even on the assumption that determinism is true.

²⁰ This raises the question of whether we may perform merely subjectively risky actions. There is a controversy among moral philosophers about this issue. Some hold that the permissibility of these actions depends on your available evidence. See Thomas M. Scanlon, (2009) *Moral Dimensions*, Cambridge Mass.: Belknap Press of Harvard University Press. Others deny that the permissibility of merely subjectively risky action depends on the agent's evidence. Instead, they say that the permissibility of such an action depends on the action's actual effects. See Judith Jarvis Thomson (2008) *Normativity*, Chicago and La Salle, IL: Open Court Press.

²¹ I pass over complications that arise concerning the blameworthiness of an agent's conduct on the basis of a belief that is not backed by the available evidence.

Blame. You are blameworthy for taking a subjective risk of a certain magnitude just in case you are not permitted to take an objective risk of the same magnitude.

This principle strikes me as very attractive. If this principle is correct, then Objective Risk Tolerance is true just in case the following thesis is true:

Subjective Risk Tolerance. You are not blameworthy for exposing someone to a negligible subjective risk of death in order to otherwise provide this person with a small benefit. (All else being equal.)

So if we give up Objective Risk Tolerance, then we must also give up Subjective Risk Tolerance. Therefore, in all events, we are forced to concede that our everyday risky actions are blameworthy. In light of this concession, it would be a pyrrhic victory to claim that these actions are permissible. So, quite apart from the fact that the Subjective Risk Solution is dubiously motivated, insofar as the jury is out over whether everyday actions are objectively risky, additionally this solution does not adequately reconcile all our moral intuitions about aggregation and chance, since these intuitions also pertain to blameworthiness.

6. A Second Possible Solution: Do *Ex Ante* Interests Matter?

Assuming the Repetition Argument is sound, it shows that we cannot hold both of these claims:

Risk Tolerance. You may expose someone to a negligible risk of death in order to otherwise provide this person with a small benefit. (All else being equal.)

Many-Few. You may not let someone die in order to provide small benefits to many. (All else being equal.)

These claims seem attractive because they could explain intuitions of ours. Let us focus on Many-Few. The reason why we are drawn to Many-Few is that it is a way of capturing an intuition about the case we encountered earlier:

ISLANDS. If you go to the north island, then you can save Jones's life. If you go to the south island, then you can cure the headaches of a billion people. You only have enough fuel to go to one island.

Can we explain the intuition that you ought to save Jones in ISLANDS by appealing to a different principle that also tolerates bringing the Chancy Antidote to many people in POISON?

To do so, we should look for a morally important difference between the two cases. There is one difference that looks as if it may be important. In ISLANDS, there is a single person, Jones, whom you would fail to save, were you to cure others' headaches. On the other hand, if you keep bringing the Chancy Antidote to sufficiently many people in POISON, then you would be

near certain to fail to save someone, were you to cure others' headaches. But there is no single person whom you would be near certain to fail to save. This contrast is sometimes put as follows: In ISLANDS there is an "identifiable" victim, but in POISON there is a merely "statistical" victim.²²

In light of this observation, let us coin some terminology. Let us say that a risky action is in someone's "*ex ante* interest" if this action produces a positive amount of expected utility for him or her.²³ In these terms, we can formulate this principle:

Ex Ante View. You may let someone die and provide small benefits to many when and only when your doing so is in everyone's *ex ante* interest at the time that you are acting. (All else being equal.)²⁴

As I have defined the view, the *Ex Ante* View evaluates actions' permissibility at the time at which they are performed. Hence it determines the permissibility of a risky action before any particular outcome has resulted from it. According to this principle, you may not fail to save Jones in ISLANDS – doing so is not in Jones's *ex ante* interest. More generally, the principle would not permit you to sacrifice an identifiable life for the sake of many small benefits.

²² Thanks to an anonymous reviewer for the helpful suggestion that I distinguish clearly here and throughout whether the lost life is that of an identifiable victim or a merely statistical life. The terminology derives from Schelling:

There is a distinction between an individual life and a statistical life. Let a 6-year-old girl with brown hair need thousands of dollars for an operation that will prolong her life until Christmas, and the post office will be swamped with nickels and dimes to save her. But let it be reported that without a sales tax the hospital facilities of Massachusetts will deteriorate and cause a barely perceptible increase in preventable deaths – not many will drop a tear or reach for their checkbooks.

Thomas C. Schelling (1968) "The Life You Save May Be Your Own," in Samuel Chase, ed., *Problems in Public Expenditure Analysis*, Washington, D.C.: The Brookings Institute.

²³ We should calculate expected utility in the traditional method – we multiply the utility of each harm or benefit by the probability that it results, and then sum these products.

²⁴ Michael Otsuka considers a view similar to this one in an appendix that was formerly a part of his unpublished essay, "Risking Life and Limb: How to Discount Harms by Their Improbability." Within the context of his view that the justification of actions depends on the complaints that people might have against them, he considers the view that:

for any individual who ends up suffering harm as the result of risky activity, we might allow for the discounting of his complaint against suffering harm by the improbability that *he in particular* would suffer harm in the following restricted set of circumstances: those in which it is in his *ex ante* self-interest to be exposed to such risks when one takes into account the expected benefit to him of such risky activity.

Since Otsuka is interested in the justification of everyday actions that expose people to risk, I take him to be considering the possibility that we should "discount" this person's "complaint" so much that the relevant actions are permissible.

On the other hand, this principle would permit you to repetitively bring the Chancy Antidote to people in POISON – doing so is in each person’s *ex ante* interest, including the *ex ante* interest of anyone who turns out to die. In this case, it permits the loss of a merely statistical life in these circumstances because, *at the time of action*, each individual had a positive expectation in terms of his or her own utility. When the lost life is a merely statistical one, the *Ex Ante* View directs us to think in terms of intrapersonal aggregation. But when the lost life is one that is identifiable at the point of action, the *Ex Ante* View maintains that the issue is no longer one of intrapersonal aggregation.

Why would it be morally significant that an action is in everyone’s *ex ante* interest? A possible answer is the contractualist claim that our actions must be capable of being “justified to *each* affected person,” as James Lenman puts it.²⁵ Some people argue that you could justify your action to someone if the action was in this person’s *ex ante* interest.²⁶ In a similar vein, one could motivate the *Ex Ante* View by stressing the importance of *hypothetical consent*. Someone would be able to hypothetically consent to an action if this action was in his or her *ex ante* interest.

The *Ex Ante* View might seem like the way to reconcile our intuitions about POISON and ISLANDS. But it faces an initial problem in that it is questionably motivated. There are well-known controversies about contractualism and hypothetical consent that I will not enter into. In addition, there is a further worry as to why *ex ante* interests should matter. Suppose a risky action ends up harming somebody. Why would it make a moral difference that at *an earlier time* this action was likely to be in their interest? Why should we be concerned with the likely effects of an action rather than its actual effects? These are questions for the contractualist as much as the rest of us: Why would actions be justified to someone on the basis of the actions’ likely effects rather than their actual effects?

We can set these worries about motivation to one side, for there is a separate and serious problem with the *Ex Ante* View: It is committed to unacceptable claims. Consider:

²⁵ Lenman (2008) Lenman denies that an action’s being in someone’s *ex ante* interest is sufficient for being justifiable to that person. He thinks that in addition the action must not preclude taking reasonable precautions, and must be capable of being guided by others’ safety. Thus, he suggests that a contractualist adopt the “Aim Consistency Principle,” which:

would then say something like this: my act of imposing a risk on you may be rendered permissible by furthering some purpose that does not in principle preclude my taking reasonable precautions against your coming to harm; other acts I might perform are rendered impermissible when the purpose from which they might seek a warrant is inconsistent with your safety’s being one of my guiding aims (111).

²⁶ See Johann Frick, “Health Resource Allocation Behind a Natural Veil of Ignorance,” manuscript.

ROULETTE WHEEL.²⁷ There is a giant roulette wheel in the sky. It has a billion slots that are connected to chutes. The chutes are directed at individuals who have swallowed poison. The poison will make each individual have a headache and die, if untreated. To make the wheel work, you have to press a button. This will give the wheel an indeterministic spin, and then release the contents of each slot into a chute. Before you do so, you have to make an antidote to put in the slots. You only have enough resources to choose between these options:

Vials For Everyone. You make a billion vials of an antidote that will certainly save each recipient's life, but do nothing for each recipient's headache.

Vials That Cure Headaches. You make a billion minus one vials of an antidote that will certainly save each recipient's life and cure this person's headache.²⁸

The net effect of choosing Vials That Cure Headaches rather than Vials For Everyone is that 999,999,999 people have their headaches cured and one person dies. Intuitively, you may not choose Vials That Cure Headaches. However, the *Ex Ante* View predicts that you may. This is because this option would be in the *ex ante* interest of all. (I assume that a one-in-a-billion chance of death is small enough for headache relief to compensate for it. Feel free to make the probability smaller until you agree.) Since the *Ex Ante* View reaches an unacceptable result about cases like this, we have reason to reject it.²⁹

²⁷ I borrow the whimsical idea of a giant roulette wheel in the sky from Otsuka's "Risking Life and Limb."

²⁸ James Lenman considers a similar government policy consisting of performing two actions. First a lottery is held. Then the winners are given a benefit and the losers are killed by the police. Lenman (2008: 102). He proposes this as a counterexample to the claim that *policies* are justified when they are in the *ex ante* interest of all. Later, he considers "Russian Roulette Cases," which involve a single *action*. Some of these seem to be a variant of the objectionable cases where:

the push of a button sets the lottery in train and some complex, entirely automatic mechanism then brings it about that the losers are killed without the involvement of the police or any other human agency after the button is pushed (112).

His remarks imply that he would consider it impermissible to push the button in such a case, and that this would refute the claim that an action is permissible when it is in the *ex ante* interest of all. Lenman claims that his "Aim Consistency Principle" (see previous footnote) can explain why this is so, on the grounds that:

it is hard to see how [the] redeeming feature [of not precluding the taking of precautions] could be present at least in any remotely realistic example involving Russian roulette or elaborate, automated, survival-lottery machines (112).

²⁹ One might worry that this conclusion is too hasty, on the following grounds: Many people, especially policy makers in the environment and health, on reflection, change their intuition about cases like ROULETTE WHEEL and come to think that there is no morally relevant difference between an identified life and a statistical life. (I owe the substance and much of the formulation of this objection to an anonymous reviewer.) But I would respond to this objection by maintaining that it is wrong to choose the option Vials That Cure Headaches,

7. Conclusion

I have presented an argument that shows that the principles Many-Few and Risk Tolerance are inconsistent: This turns on the fact that, by iterating a risk of a harm, we make it near certain that this harm occurs in the long run. The significance of this inconsistency is not merely limited to these specific principles. It points to a problem in our intuitions about risks and aggregation that makes these specific claims seem appealing. Ideally, we would like to retain these intuitions by finding alternative principles that accommodate these intuitions but do not conflict with each other. Unfortunately, I have been unsuccessful in discovering any such principles. I like to think that others will be more successful in doing so. So my essay aims to present a challenge for others, which can be put as the following question: What alternative way is there of reconciling our intuitions about cases like ISLANDS and about everyday actions that expose potential beneficiaries to minuscule risks of death?

Although I remain optimistic that someone will be able to answer this question by finding a satisfactory way of reconciling these intuitions, it is worth considering the worst-case scenario in which no answer is forthcoming. In this situation, we must revise some of our intuitions. We have two options, which I will discuss in turn.

The first option would be to revise our intuition about ISLANDS. We might motivate doing so by noting that it is an intuition about an incredibly large number of people. Our ability to compute such large numbers may be limited. For example, I know that my attempts to imagine a million people are suspiciously like my attempts to imagine a billion. Maybe we are unable to appreciate the moral significance of helping so many people. Revising the intuition about ISLANDS would be important for our moral theory. It would affect what view we should take of a utilitarian approach to beneficence. The utilitarian approach endorses “aggregation” – it aggregates benefits to individuals into total amounts of welfare, and recommends that you bring about

whether or not the lost life is identified or statistical. As stated in the main text, the lost life is a merely statistical life – it is a matter of chance who it is that misses out on the antidote. In this original formulation, it is intuitively wrong to choose the option Vials That Cure Headaches. But we could change the case so that it is an identifiable individual, Smith, who would not receive the vial. It remains the case that intuitively it is wrong to choose the option Vials That Cure Headaches. Therefore, I conclude that our intuition that it is wrong to choose the option Vials That Cure Headaches is independent of our views about the relative value of identifiable and merely statistical lives. Perhaps a more general worry arises from the fact that some people revise their intuitions about whether there is a moral difference between an identified or merely statistical life. If our intuitions are that there is a moral difference here, but on reflection we revise these intuitions, then we may worry that more generally our intuitions about cases involving risks and large numbers of harm are unreliable. I am not able to respond adequately to an objection along these lines that proposes a general skepticism about our intuitions about these sorts of cases. I can say only that I have sufficient confidence in my own intuitions and I am relying on the reader having a similar confidence in his or hers. To the extent that this turns out not to be the case, my argument is ineffective.

the greatest total amount of welfare possible. Some of us thought that cases like ISLANDS were the death knell for the utilitarian approach to beneficence. If our intuition about ISLANDS is misguided, then our rejection of the utilitarian approach would be called into question.

The second option would be to revise our intuitions about everyday risky actions. We could motivate this option by noting that these intuitions concern incredibly small probabilities. Our ability to compute such small numbers may be limited. For example, I know that my deliberations about a one-in-a-billion risk are suspiciously like my deliberations about no risk at all. Moreover, some people's intuitions about risks depend on whether these risks are described quantitatively or qualitatively. For example, they may fear a risk of "one in a million," but be comfortable with a risk "equivalent to that of being hit by lightning." Indeed, empirical work suggests that in general we reason poorly about probabilities: We adopt heuristics, and commit probabilistic fallacies.³⁰ Lastly, our intuitions about the permissibility of taking certain risks may be influenced psychologically by the fact that we *do* take these risks. The psychological literature on "cognitive dissonance" finds that we are under psychological pressure to avoid the unpleasant feeling of believing both that a behavior is wrong and that we engage in this behavior.³¹ This puts pressure on us to rationalize our actual behavior: We are under pressure to judge that our behavior is permissible. So it may be that these intuitions about everyday risks are less trustworthy than we might hope. However, if we revised our intuitions about everyday actions then we would have to change the way we live parts of our everyday lives, on pain of acting impermissibly. We would have to be much more cautious than we are now when helping people. We would have to avoid all beneficial actions that expose the potential beneficiaries to any risk of death at all: no lifts to the airport, sushi, kisses and so on. This is not an exciting prospect.³²

Tom Dougherty
University of Sydney
Department of Philosophy
tom.dou@gmail.com

³⁰ See, for example, Amos Tversky and Daniel Kahneman (1974) "Judgment Under Uncertainty: Heuristics and Biases," *Science* 185: 1124-31.

³¹ See, for example, Leon Festinger (1957) *A Theory of Cognitive Dissonance*, Stanford: Stanford University Press.

³² For helpful comments, I would like to thank Johann Frick, Caspar Hare, Richard Holton, Rae Langton, Bradford Skow, Judith Jarvis Thomson, the M.A.T.T.I. reading group at MIT, the audience of my doctoral public defense at MIT and two anonymous reviewers for the *Journal of Ethics & Social Philosophy*.