

Prevention, Rescue, and Tiny Risks¹

J. Paul Kelleher
University of Wisconsin-Madison

Public Health Ethics 6(3) (2013): 252-261

ABSTRACT: Contrary to popular belief, population-wide preventive measures are rarely cost-reducing. Yet they can still be cost-effective, and indeed more cost-effective than treatment. This is often true of preventive measures that work by slightly reducing the already low risks of death faced by many people. This raises a difficult moral question: when we must choose between life-saving treatment, on the one hand, and preventive measures that avert even more deaths, on the other, is the case for prevention weakened when it works by reducing many healthy people's already low risks by a further tiny amount? I argue the answer is no.

Introduction

Contrary to popular belief, population-wide preventive measures are rarely *cost-reducing*. There are at least two reasons for this. First, such measures must be applied to a very large number of people, and this can be expensive in the aggregate even when the cost of each application is fairly low. Second, much successful prevention simply delays death—compare dying of heart disease at 55 to dying of cancer at 85. In the meantime, other health conditions may be found and treated, adding costs that would not have arisen if one had died earlier rather than later.

Yet even if most population-wide preventive measures are not cost-reducing, they can certainly still be *cost-effective*. That is, preventive measures can have greater benefits than costs, and the cost-per-benefit can still be low when compared to treatment. This raises the following

¹ I would like to thank Paul Menzel, members of the Philosophy Department at the University of Wisconsin-Madison, and two anonymous reviewers for very helpful comments and discussion on previous drafts.

question: should we sometimes let people die of a treatable condition in order to use scarce funds to prevent a greater number of qualitatively similar deaths in the future? This is a difficult moral issue for at least the simple reason that treatment helps existing victims—people who need help *now*—whereas prevention affords a benefit to those who are, by contrast, relatively healthy.

Adding to this potential strike against prevention is the fact that many effective preventive measures work by reducing already low risks by an additional tiny amount. This point was first stressed by the British epidemiologist Geoffrey Rose, who observed that many deaths can be prevented simply by providing tiny reductions in the already low risks faced by a large number of people (Rose 1985; Rose 2008). Consider that a figure displaying the distribution of a risk factor throughout a population will often take the form of a bell-shaped curve. Suppose that the risk factor becomes more serious as one moves from left to right along the curve, so that individuals whose levels fall all the way to the right are those with the highest risks. The curve's hump gets its shape because the vast majority of people fall within the middle part of the distribution. These people have lower risks and thus a lower chance of death than those whose risk falls all the way to the right. Rose's observation was that even if the risks that fall in the middle of the distribution are already rather low—and thus even if the *individuals'* risks are quite low—still this middle part of the distribution can generate many deaths. Indeed, many more deaths can and often will come from the middle of the distribution than will come from the rightmost part. This is because a low individual probability that nevertheless affects *many* people will still give rise to many deaths. In light of this, Rose noted that a good portion of the deaths stemming from a given risk factor can be eliminated simply by shifting the entire bell-shaped distribution just a little bit in the “favorable” direction (in this case to the left). Rose knew that individuals, *as* individuals, may not view this as a very significant benefit. But again the slightest

Please cite only final published version

reduction in risk, applied to a very large number of people, can significantly reduce the aggregate number of deaths that result. One implication of this “population approach” to prevention is that it may not prevent many deaths among those who were and still are at highest risk. But if it is death one wants to prevent, then slightly shifting the entire population’s distribution of risk is a good way to go. I will refer to this mechanism of prevention as *Rose Prevention*. Examples of Rose Prevention can include certain forms of medical screening, health information campaigns, requiring individuals to wear seatbelts, mitigating air pollution, and providing cholesterol-lowering drugs at (some) public expense. In this paper I address the following question: when we must choose between life-saving treatment and more cost-effective Rose Prevention, is the case for Rose Prevention weakened because it works by reducing many people’s relatively low risks by a further tiny amount? I will argue that the answer is no.

I will begin the moral assessment of Rose Prevention by considering in detail a compelling argument in support of it. Tony Hope’s argument is ingenious, but I will argue that it is incomplete in several respects. I then examine what I take to be the two most promising lines of argument for life-saving treatment over Rose Prevention, each of which seeks to exploit a gap in Hope’s argument. One line of argument focuses on the varying strength of different people’s claims to be helped, while the other focuses on individuals’ responsibility for their own health. My conclusion in each case is that neither argument poses a genuine threat to the moral case for Rose Prevention.

Rose Prevention and the Rule of Rescue

The moral case for Rose Prevention faces an immediate hurdle in what is often called “the rule of rescue.” This can be interpreted as a psychological fact about human beings, as a moral rule, or both. As a fact about our psychology, the rule of rescue is related to the strong impulse we feel when someone near us is in dire peril. In such situations, it can feel intolerable to “just stand around and let her die” when there is something that can be done. This makes it is easy to see why the rule of rescue has also been construed as a moral rule: in at least some cases, it not only feels intolerable to just stand around and let someone die, it *is* intolerable. Tony Hope nicely connects such moral convictions to the observations about risk that I have recounted from Rose:

The most powerful reason in support of paying more to save the identified life [with treatment], I believe, is that in the typical cases of prevention the intervention makes only a small difference to the probability of death of any one individual whereas in the typical case of rescue (but by no means in all examples) the intervention makes a large difference to one or more individuals. (Hope 2001, p. 183)

Hope points out that since the preventive benefit that Rose Prevention bestows upon each beneficiary is so small, each potential beneficiary would be “rational to trade that small extra risk in order to make a small but definite contribution towards saving someone else’s life” (2001, p. 183). To illustrate, suppose everyone in a population of 300 million has a 2-in-10 million chance of acquiring and dying from a certain disease. By giving everyone a reduction of 1-in-10 million, 30 expected deaths are prevented—by all accounts a significant population benefit. Note, however, that the 1-in-10 million risk of dying is the same risk that the average American

Please cite only final published version

accepts on the average car trip (Fischhoff and Kadvany 2011, p. 12). But then if it can be rational for an individual to accept a tiny increase in risk by driving to the store for, say, a cupcake, then surely it can be rational for an individual—indeed, for many individuals—to accept a tiny increase in risk in order to provide a life-saving benefit to someone who is currently very sick. This line of reasoning suggests that it might be reasonable for a population to collectively decide to prioritize treatment even when treatment is less cost-effective than prevention. Surely (the argument goes), it cannot be morally wrong for a large number of people to each make such a small sacrifice on behalf of an obviously much worse-off minority.

Hope, however, thinks it *is* morally wrong, or, at least, that it is morally wrong for public policy to encourage and foster this collective sacrifice. He offers an ingenious thought experiment to prove the point. Suppose that a trapped miner can be rescued (i.e. “treated”) by any of several rescue parties. And suppose that the larger the rescue party, the lower the risk of dying faced by any given rescuer in the party. More specifically:

If there were 100 rescuers there would be a 1:1,000 [i.e. a 1-in-1,000] chance for each rescuer of death. If there were 1,000 rescuers each would face a 1:2,000 chance of death. If 10,000 rescuers then each would face a 1:5,000 chance of death. If 100,000 rescuers (an extraordinarily large rescue party—but this is a ‘thought experiment’ to test a theoretical point) then each would face a 1:10,000 risk. (Hope 2004, p. 37)

Thus, as the search party gets larger and larger, the risk to each rescuer decreases. In this way, Hope connects his trapped miner thought experiment to the previous claim that it could be perfectly rational to altruistically forgo a small preventive reduction in risk so that another can

have a large reduction in risk via rescue or treatment. It is here that Hope makes his ingenious argumentative move: for while it could be perfectly rational to want to be part of a search party aimed at saving a trapped miner, it can also be rational to want one's own risk of dying during the rescue to be as low as possible. Assuming one is already willing to help with the rescue, it is in one's self-interest to be part of as large a search party as possible. But then this leads to the implication that Hope finds morally troubling, namely that the rescue party that is in each rescuer's self-interest is also the rescue party that leads (with some statistical certainty) to the deaths of roughly ten rescuers. Such a rescue party is, Hope concludes, "highly problematic from a moral point of view" (Hope 2004, p. 39). Hope's claim here is that it is utterly counterproductive to trade ten lives for one, and that we lose sight of this fact when we focus instead on each rescuer's seemingly negligible change in individual risk. The analogy with the real world is clear: when it is likewise counterproductive to fund rescue treatments instead of Rose Prevention, one cannot defend treatment by citing the seemingly negligible impact that Rose Prevention makes in the lives of those it helps. The tininess of that impact, Hope would say, is a moral distraction.

As I have said, this is an ingenious argument-by-analogy against the rule of rescue. There are, however, at least three problems with it. The first is tied to Hope's claim that "A health care system that spends more per year of life gained on rescue treatments (such as renal dialysis) than on 'statistical' treatments [such as statin drugs for elevated cholesterol] is effectively volunteering those who would benefit from the preventive treatment to take part in a 'rescue party' for those requiring the rescue treatment" (Hope 2004, p. 39). Hope goes on to compare such a health care system to a morally questionable army officer who seeks to organize a rescue party in which ten will die in order to save just one trapped miner. Yet Hope ignores a clear

Please cite only final published version

disanalogy here: in Hope's case of the trapped miner, the rescuers *increase* their risks by joining the rescue party. This is quite unlike the real-world case in which health policy makers are deciding whose pre-existing risks to *reduce* and whose risks to leave *untouched*. This is a subtle difference, but one cannot assume it is morally irrelevant. After all, commonsense morality often recognizes a moral distinction between harming someone, on the one hand, and failing to help someone, on the other. But then if the morality of worsening is not the same as the morality of helping, it is possible that Hope's argument in favor of Rose Prevention is helped by eliding the two. It is possible, for example, that some readers will be swayed by Hope's argument because they are led to believe that in funding renal dialysis instead of statin drugs (which reduce relatively healthy people's cardiovascular risks by a small amount), the health system is thereby *subjecting* those who might benefit from statins to tiny increases in risk. But it is not subjecting them to increased risk; rather, it is failing to reduce the level of risk they started with.

Admittedly, it is not clear how large a problem this is for Hope. After all, the main point of the trapped miner case was to display the sheer folly of trading ten lives for one. As I have been putting it, a very large rescue party in which ten rescuers will die is straightforwardly counterproductive. I imagine that this is what many of Hope's readers will take away from his analogy, and displeasure with this counterproductivity may remain once we see that we are working within the morality of helping and not the morality of worsening. Yet this focus on counterproductivity raises a second problem for Hope. In Hope's presentation, the only relevant reason to save the one miner is simply that death is bad. And of course if *that* is the only reason we have to save a life, then clearly it is problematic to trade ten lives for one. There is, however, a flip-side to this observation: if the badness of death is *not* the only relevant reason, then treatments that are less cost-effective than Rose Prevention might be morally desirable *even if*

they are counterproductive by the lights of a morality that focuses solely on reducing badness. To see this, consider an example that Hope does not discuss, but which is depicted visually by an illustration accompanying his book's discussion. The example is the rescue mission portrayed in the movie *Saving Private Ryan*. In *Saving Private Ryan*, the point of rescuing Private Ryan was not simply that a soldier's death is bad. Rather, General Marshall ordered the rescue in order to spare Ryan's mother the anguish of losing a fourth son to the fighting in World War II. As portrayed in the movie, Marshall was moved by Abraham Lincoln's letter to a woman who had lost five sons in the American Civil War. Now, Lincoln's letter is obviously not concerned with this paper's topic—viz. differential ex ante risks—but rather with the horrific ex post outcome this woman had already ensured. By contrast, Marshall was clearly concerned to eliminate the appreciable ex ante risk that Ryan's mother faced of losing her fourth son. Suppose, then, that Marshall had put his rationale this way: "In the context of a just war, mothers of soldiers must accept the anguishing risk of losing a child in battle; but it is nevertheless wrong to expect mothers to run a significant risk of losing more than three children." If one accepts that rationale, then one may well accept that some rescue missions—like the one launched for Private Ryan—are defensible *even if* they fail to minimize the overall number of deaths.

Now, I am not here defending General Marshall's order. At this stage, I am merely noting that Hope's trapped miner example is free of the complicating moral considerations that Marshall might well have cited to defend a rescue mission that (in the movie) his subordinates criticized as straightforwardly counterproductive. If any such moral consideration is ever relevant in the real-world health policy context that is Hope's ultimate concern, then this is a further reason to resist the temptation to apply the moral conclusion from the trapped miner case

Please cite only final published version

to the trade-offs we face in the real world. I will discuss this issue at greater length in the next section.

Finally, there is a third problem with Hope's trapped miner case, a problem tied to yet another difference between it and the real world. In Hope's thought experiment, it is *stipulated* that if a very large rescue party is needed to save the trapped miner, then roughly ten rescuers will die in any successful rescue mission. Thus for his analogy to the real world to go through, Hope needs it to be the case that we cannot fund, say, long-term renal dialysis for one patient without thereby leading to several other deaths because a great many people will have to go without cholesterol-lowering statins. Should we accept this assumption? At first blush it is hard to see why we should. After all, many of those who would benefit from statins can give themselves the same tiny risk-reduction on their own through improved diet and increased exercise. Unlike in the thought experiment Hope has constructed, it is not a law of nature that at least some deaths will occur among those who are "volunteered" in the real world to help save a patient with dialysis. Since we are considering preventive measures that reduce risks by a tiny amount, many of these can be self-conferred by people who are willing to make small efforts to reduce their own risks by the same degree. I will return at the end of this paper to this particular disanalogy between Hope's thought experiment and the real world.

I have noted three ways in which Hope's trapped miner case might differ from the situation we face in the real world. But I have not argued that Hope's preference for Rose Prevention over rescue treatment is mistaken. Whether he is right depends on what arguments can be made in favor of rescue treatments. The next two sections discuss what I think are the two most promising such arguments.

Rose Prevention and the Morality of Individual Claims

I introduced General Marshall's rationale behind the Private Ryan rescue to show that Hope's Trapped Miner case is free of certain moral factors that might be present in real-world health policy contexts. I asked us to imagine that Marshall was moved by the belief that it is impermissible to impose upon any mother a significant risk of losing a fourth son to war. Note now that Marshall's belief might in turn be grounded in the view that mothers possess strong moral *claims* against their government to protect them from such risks. In that case, Marshall's view would be that Ryan's mother has a stronger claim in favor of the rescue mission than anyone has against it. Claims, then, would appear to be just the sort of moral factor that one might invoke to justify a policy that fails to minimize the overall number of deaths.

As it happens, Norman Daniels has recently argued for the importance of individual claims to assistance in assessing the relative moral priority between helping "identified" and "statistical" victims. For our purposes here, the distinction between helping identified victims and helping statistical victims can be viewed as coextensive with the distinction between rescue treatment and Rose Prevention. The relevant question is the same: conceding that we could never know whose death would be prevented with Rose Prevention, should we nevertheless prevent more statistical deaths with Rose Prevention when we could instead rescue fewer identifiable individuals who are quite obviously in peril right now? According to Daniels, the morality of individual claims does provide some support for treating the identified victims instead. This is because they are "worse off" than statistical victims (Daniels 2012, p. 41). Daniels offers the following thought experiment. Suppose we have five tablets of medication and face six people. Alice will die if she does not receive all five tablets. Each of the other five individuals (which I

Please cite only final published version

will call the “opposing group”) has a 1-in-5 chance of dying of the same terminal disease; by giving everyone in the opposing group one tablet, we can vaccinate them and eliminate this risk. Daniels writes, “I believe we have a stronger obligation to treat Alice than to vaccinate the five others,” even though “one expected life is saved” in either case (p. 41). Daniels claims that if one does not share his intuition that Alice should be helped in this one-vs.-five case, at some point one will likely agree that the “concentration of risk matters morally”; for example, Alice’s higher risk clearly matters, Daniels claims, in a case where Alice needs the available 1,000 pills, which can instead be used to eliminate a 0.1 percent risk of dying in each of 1,000 others. Daniels suggests (but does not explicitly state) that Alice should be given the medicine even if it is a statistical certainty that one member of the opposing group will in fact die if the group is not vaccinated.

Within moral philosophy, the morality of claims is often invoked as a source of reasons to reject consequentialism, which tells us to add up the good that can be produced through different courses of action, and to choose a course that produces at least as much goodness as any other. Many non-consequentialists are troubled by consequentialism’s implication that if one had to choose between (1) preventing a death and (2) curing mild headaches in some *very* large number of people, there is some finite number of headaches such that curing them is morally preferable to preventing the death. A standard non-consequentialist solution to this “aggregation” problem is to insist that in such cases, the proper decision procedure involves making “pairwise comparisons” between individuals’ claims to assistance. That is, instead of simply comparing the aggregate good that each option would yield, one should compare each individual’s claim to be helped against each other individual’s claim. Since the person whose life is in danger intuitively has a very strong claim to assistance, and since each headache-sufferer intuitively has a rather

weak claim, a morality of claims will conclude that the one death should be prevented. This is because the one person's claim wins out over *each* of the others' claims. So if it is true that Alice has a stronger claim to assistance than any of the others in the opposing group, and if the morality of claims is a defensible moral framework, then it can be invoked to explain and support Daniels' intuition that Alice should be saved.

The decision procedure of pairwise comparison does seem to yield the right answer in the death-vs.-headaches case, and Daniels suggests it gets things right in his Alice-vs.-others case too. In what follows, I will say that a *Type 1* case is a case that pits a certain number of identifiable victims against the *same* number of statistical victims. (Following Daniels, I will assume for ease of presentation that if an opposing group is not helped, the probabilities will run true and precisely that many statistical victims will result.) Thus Daniels's Alice-vs.-others case is a Type 1 case. Consider now a variation that I will call *Type 2* cases. These are cases that pit a certain number of identifiable victims against *more* than that number of statistical victims, but where everyone in the opposing group still has *much* lower risks than the imperiled identifiable individuals. Here is an example of a Type 2 case: we can either treat and cure Alice of her fatal disease or vaccinate an opposing group of one million people and thereby eliminate each member's 5-in-1 million risk of death. Now note that if the strength of one's claim is commensurate with the magnitude of one's antecedent risk, as Daniels suggests, then Alice would still have a much stronger claim than any of the others in this Type 2 case. The procedure of pairwise comparison would then still support saving Alice even if that meant allowing five deaths within the opposing group. This in turn would count in favor of the rule of rescue and against more cost-effective Rose Prevention.

Please cite only final published version

Let us now ask: *should* pairwise comparison be applied in Type 2 cases? Daniels suggests not, although he does not attempt to account for this stance theoretically, presumably because it is grounded directly in intuitions. Daniels simply concludes that, “Risk concentration might be such a weak factor that it matters morally only in breaking ties” between preventing Alice’s one death and preventing the *same* number of statistical death among those in the opposing group (p. 41). Thus Daniels would say that Alice’s claim wins out in Type 1 cases but not in Type 2 cases. But this seems problematic. If the morality of claims should be used to decide Type 1 cases (as Daniels suggests), and if Alice continues to have a *much* stronger claim in Type 2 cases, then why does her claim not continue to trump much smaller claims in cases of Type 2?

Daniels could rightly reply here that virtually all non-consequentialists embrace a self-consciously moderate view on which pairwise comparison is to be relied upon in some cases but not in others. As Dan Brock and Daniel Wikler write:

[A]ccording to many moral theories, individuals should confront others competitors for scarce resources as individuals, and their priority for treatment should be determined by the urgency of their individual claims to treatment. Then again, most people and most moral theories do not reject all aggregation of different sizes and costs of health benefits in setting priorities and allocation, although there is no consensus either on when aggregation should be permitted or for what reasons. (Brock and Wikler 2006, pp. 263-4)

Brock and Wikler go on to state that pairwise comparison is precisely the right approach in cases like death-vs.-headaches. But they, like Daniels, suggest that pairwise comparison should be rejected in other cases. As an example of the sort of case they might have in mind, consider the

choice between saving one person's life and preventing 1 million others from developing paraplegia. As Brock and Wikler note, many people (including many non-consequentialists) believe that the large number tips the ethical scales toward the 1 million, despite the fact that no one among the 1 million has an individual claim to assistance that is as strong as the claim held by the one person who faces death. If preventing the 1 million cases of paraplegia is morally permissible—and I agree it is—then this is at least one example where pairwise comparison gives way to more aggregative concerns. Thus Brock and Wikler, like Daniels, hold a moderate view that endorses the use of pairwise comparison in some contexts while rejecting it in others.

Suppose we agree with Daniels and Brock and Wikler that *some* such moderate theory is superior to a thoroughgoing consequentialist view that allows enough headaches to trump saving a life. Our question now is whether the most plausible moderate view would count against Rose Prevention by permitting the claims of those in dire peril to win out over the claims of those who could benefit from tiny reductions in risk. We have seen that on Daniels' view, those in dire peril do have stronger claims, but those claims do not win out in Type 2 cases because the procedure of pairwise comparison does not apply there. However, Brock and Wikler believe that Daniels has made a crucial mistake. More specifically, they believe that pairwise comparison should be used *both* in Type 1 cases *and* in Type 2 cases. If Brock and Wikler are right, what would this mean for Rose Prevention? The answer turns on whether Daniels is correct that Alice has a stronger claim than any other person in *both* types of case. If she does, and if Brock and Wikler's stance on pairwise comparison is correct, then Alice's claim would win out in both cases, and that would count against Rose Prevention. But Brock and Wikler take a different tack, for they reject Daniels' view that Alice has the stronger claim in both cases; that is, they hold that a

Please cite only final published version

person can have a strong claim to risk-reduction *despite having an exceedingly tiny antecedent risk*. Brock and Wikler write:

Consider the relative moral importance of providing treatment to an acutely ill AIDS patient versus offering a larger number of people who have a slight risk of infection with HIV a modest reduction in that risk. As an object of compassion, a person suffering from AIDS now presents a stronger claim than someone with a slight chance of becoming infected. Would it be wrong, therefore, to give higher priority to the prevention program, even if it saved more lives? Once again, the argument fails in its own terms. In a large population, the numbers add up, and without prevention there will be many people suffering from AIDS. Each is just as real as the current patient... The differences are that this patient's suffering will occur in the future... (Brock and Wikler 2009, p. 1673)

According to Brock and Wikler, then, Daniels is wrong to highlight the fact that in Type 1 cases Alice is worse off than n others who each face a 1-in- n chance of dying. For while Alice is *now* worse off than they are *now*, she is not worse off than one of them *will be* in the future. Thus Brock and Wikler suggest that Daniels's argument "fails in its own terms": its own terms suggest that the strongest claim should win out in Type 1 cases, and in fact (Brock and Wikler contend) Type 1 cases involve a clash between *two equally strong* claims—one in the present and one in the future. Brock and Wikler could therefore say that in Type 1 cases, we should give equal chances to Alice and to the other person—*whoever* it is—who will die in the future if the opposing group is not vaccinated. If the resulting lottery comes out in favor of preventing the

future death, then the opposing group should be vaccinated as a means to preventing that one future death.

Brock and Wikler's view has definite strengths. For one thing, it accounts for the intuition that Daniels reports many others having about Type 1 cases, namely that there is no obligation to prevent the "identified" death rather than the one "statistical" death that will occur in the group (p. 41). Further, Brock and Wikler's view explains Daniels' own intuition that we should vaccinate the group in Type 2 cases. Whereas Daniels accounted for this intuition by saying that pairwise comparison no longer applies in cases of Type 2, Brock and Wikler hold onto pairwise comparison and maintain that Alice's strong claim is here outweighed by a *greater number of equally strong* opposing claims. Brock and Wikler's view therefore supports Rose Prevention while offering a theoretical explanation for why we should vaccinate the opposing group when doing so averts more deaths than can be averted with rescue treatment. By contrast, Daniels lacks a theoretical account for why the opposing group should be vaccinated in Type 2 cases. He therefore lacks an account for why Rose Prevention should be prioritized over less cost-effective rescue treatment.

Despite their view's theoretical strengths, I think Brock and Wikler are wrong to hold that there is someone among the group opposed to Alice who has as strong a claim to assistance as she has. I have two reasons for this. First, on Brock and Wikler's view, the strength of one's claim to assistance depends upon what will *in fact* happen to one. That is why they can claim that there is *someone* among the many whose claim is as strong as Alice's. But now imagine a case in which Alice has a 99 percent chance of death and one other person has a headache, and we must choose between eliminating Alice's risk and curing the other person's headache. On Brock and Wikler's view, the strength of Alice's claim to assistance turns on what will happen to her, not

Please cite only final published version

on her current risk profile. This means that if we choose to cure the headache and Alice miraculously does not die, she cannot claim that she was wronged, since she had *no claim at all* to the risk-reduction that, it turns out, she did not need. This conclusion strikes me as clearly mistaken. Alice *was* wronged: her antecedent risk-profile *did* ground a claim that was *much* stronger than the other person's claim to headache-relief. Likewise, in a Type 1 case, Alice's higher baseline risk gives her a very strong claim, and the fact that *someone* among the others will die if not helped does not show that that person—whoever it will be—has an antecedent claim that is as strong as Alice's.

Here is a second argument against Brock and Wikler's view that strengths of claims do not depend on antecedent risk. Suppose each of a million people has a 1-in-1 million chance of dying in the next hour, and that this risk is truly probabilistic. That is, it might turn out that no one dies in the next hour, that exactly one person dies in the next hour, or even that more than one person will die in the next hour. On Daniels' and my view, the strength of each individual's claim to risk-reduction remains constant across all these possible outcome-scenarios. But Brock and Wikler must say that (1) where no one dies in the next hour, there were no claims to assistance, and that (2) where one person dies, there was one very strong claim to assistance, and that (3) where more than one person dies, there was more than one very strong claim to assistance. That again seems very implausible.

For these reasons, I agree with Daniels as against Brock and Wikler that Alice has a much stronger claim than any of the others, and this is so even if we stipulate that at least one of the others will in fact die if the opposing group is not vaccinated. This does not, however, mean that I agree entirely with Daniels, for his account combines a view about claims (with which I agree) and a view about when to rely upon pairwise comparisons between those claims. And I do

not agree with Daniels that pairwise comparison applies in Type 1 cases (when it is Alice's one life versus just one "statistical" life), but not in Type 2 cases (when there is more than one statistical life on the line). To see my worry, consider the following scenario. Suppose that before disbanding and retreating into the hills, the enemy set into motion two automated firing squads that still threaten innocent lives. Firing Squad 1 (FS1) has one automated gun pointed at one innocent person tied to a post. Firing Squad 2 (FS2) also has one automated gun, but in FS2 there are 1,000 innocent people tied to numbered posts. If FS1 proceeds, the gun will shoot, and the one innocent person will surely be killed. If FS2 proceeds, a random number generator will select a number between 1 and 1,000, and then send a signal to the gun which will then shoot and kill the innocent person tied to the post with that number. Suppose we know all this and there is time to stop one firing squad but not both. Is there a moral reason to stop one rather than the other? I cannot see one. Yet this is a Type 1 case that is structurally identical to the original Alice-vs.-others case, about which Daniels says, "I believe we have a stronger obligation to treat Alice than to vaccinate the...others." Since I cannot see a reason why one must foil FS1 rather than FS2, I am inclined to disagree with Daniels and conclude that it is permissible to vaccinate the opposing group in all Type 1 cases.

I am not sure how much stock I should put in my intuitions about the Firing Squad case. But in light of how compelling pairwise comparison appears to be in death-vs.-headaches cases, and given the trouble pairwise comparison could cause for Rose Prevention, I believe it is helpful to have a case (such as the Firing Squad case) in which one very large claim is pitted against many very small claims, and yet it seems perfectly permissible to neglect the very large claim. If my intuition about this case is reliable, then it can be combined with what I have argued to support a conclusion different from both Daniels' and Brock and Wikler's. My conclusion differs

Please cite only final published version

from Brock and Wikler's insofar as I agree with Daniels about the sizes of respective claims: Alice has a much stronger claim in both Type 1 cases and in Type 2 cases. Further, I disagree with both Daniels and Brock and Wikler about when pairwise comparison applies. Brock and Wikler say it applies in both types of case, whereas Daniels says it applies only in Type 1 cases. By contrast, my view—admittedly based largely on the Firing Squad case—is that pairwise comparison does *not* apply in Type 1 cases. And if pairwise comparison does not apply there, then *a fortiori* it does not apply in Type 2 cases (in which there are even more statistical deaths on the line). Thus despite whatever merit pairwise comparison has in death-vs.-headaches cases (and I must agree it has a lot), I am led to conclude that this merit dissolves once we shift our attention to the risk-related cases that bear on the evaluation of Rose Prevention.

In the end, and despite the differences between them, Daniels' view, Brock and Wikler's view, *and* my view all give Rose Prevention priority over less cost-effective rescue treatments. It is worth noting here that John Broome has offered a still different account of claims (1990, 1994). But I have chosen to focus on Daniels' and Brock and Wikler's views because they, unlike Broome, explicitly address cases involving risk. Further, Broome's discussion suggests that whatever moral importance pairwise comparison has, it can be outweighed both by saving single net life (1998) *and* by a great many prevented headaches (2002). It is therefore clear that Broome would support Rose Prevention that prevents death more cost-effectively than rescue treatment.

Since I cannot claim to have provided a knock-down argument for the view that pairwise comparison is inappropriate in the context of Rose Prevention, let me add that I take the case for Rose Prevention to be strengthened by the fact that it is supported by all of the accounts of claims I have discussed. In the absence of a knock-down argument in favor of one or another

approach to the morality of claims, it is reasonable to provisionally accept a conclusion that is supported by all of them. From the general perspective of a morality of claims, then, there is (to use Rawls's term) an overlapping consensus in support of Rose Prevention over less cost-effective rescue treatments.

The Self-Help Argument Against Rose Prevention

If Rose Prevention is not opposed by any plausible morality of individual claims, is there any other route to a preference for less cost-effective treatment when the two conflict? The only additional route I can see is tied to what I earlier identified as the third problem with Hope's Trapped Miner analogy. That problem had to do with the fact that Hope's thought experiment *stipulates* that a very large rescue party will result in 10 expected deaths among the rescuers. As I noted, it is not clear that the real world is like this. For since Rose Prevention often works by reducing risks by a tiny amount, this opens the possibility that potential beneficiaries of prevention can reduce their own risks by that amount without much effort. To use Hope's stock example of cholesterol-lowering statin drugs, many people who benefit from these drugs can derive the same degree of benefit from an improved diet and increased exercise. Perhaps, then, treatment should be preferred in some cases because those who would benefit significantly from treatment cannot (easily) provide a benefit that large to themselves. If people can reduce their tiny risks on their own, then the real-world case would differ from the Firing Squad 2 case, in which it is implicitly assumed that the many prisoners did not subject themselves to these tiny risks and that they cannot now eliminate them without outside assistance. To what extent, then,

Please cite only final published version

does the availability of self-help support rescue treatment over Rose Prevention in the real world?

Consider a fantastical example involving seatbelts. Suppose all drivers and passengers currently use their seatbelt despite finding them mildly annoying to wear. And suppose that we are now able to give drivers and passengers a “crash-vaccine” that does one thing and one thing only: *if* a driver or passenger fails to buckle-up, the vaccine then provides them the same degree of crash-protection that the seatbelt *would have* provided. (I won’t try to explain how this magical vaccine is supposed to work.) Now imagine that we can choose between (1) providing the vaccine to millions of people (thereby enabling them to travel sans seatbelt without increasing their risk) and (2) providing an expensive cure for a small number of people suffering from a fatal genetic disease. I suspect many would be against paying for the vaccine, and for at least the following reasons: first, drivers and passengers are already easily achieving the same degree of risk-reduction on their own with seatbelts; second, those dying of the fatal disease *cannot* cure themselves. If this is a reasonable way to think about this science-fiction case, it suggests that it might be permissible to expect many people to reduce their own risks in other mildly annoying ways so that public policy can use scarce resources to help those who have a harder time doing so. If those who find it harder to help themselves are, by and large, those who face much higher baseline risks, then that is some reason to think that rescue treatments should be favored over Rose Prevention.

One problem with this self-help argument for treatment over Rose Prevention is that the underlying point cuts both ways in many real-world contexts. For unlike the science-fiction seatbelts example, many of those who now need treatment do so precisely because *they* did not previously reduce their own avoidably small risk of getting to that point. If it is morally

acceptable to expect people with tiny risks to reduce these risks on their own (or otherwise suffer the consequences), then we cannot take the brute fact that someone is in serious danger *now* as a point in their favor. Their present bad situation may simply reflect their not reducing their risk back when it was much smaller and much easier to avoid.

A further problem with the self-help argument for treatment is that it may not in fact be reasonable to expect individuals to reduce their own risks by a tiny amount. Note that the seatbelts case assumed that people are *already* doing the risk-lowering thing that they find mildly annoying to do. Yet if the prospective real-world reduction in risk is genuinely tiny, it will be understandably difficult to convince any particular individual that it is worth his or her making any effort at all to reduce it. Rose had a name for this problem: the “Prevention Paradox.” It is a paradox, he suggested, that a preventive measure yielding such large benefits at the population level can yield such small (expected) benefits to each individual (Rose 1985, p. 38; Rose 2008, p. 47). As I have stressed, it can be perfectly rational to trade tiny increases in risks for quite trivial benefits that are more certain (e.g. cupcakes). The upshot, I believe, is that individuals cannot always be criticized for passing up opportunities to reduce their risks by foregoing more certain benefits that they also care about. Of course, it is still true that if risks *are* reduced by tiny amounts in lots of people, then many premature deaths will be averted. But even if that is a worthy goal for public policy, it will still not always be a goal that policymakers can reasonably expect individuals to promote on their own, *as* individuals. As individuals, they will often need external help, if only to convince them that slightly reducing the population’s risks is a worthy goal from a wider population-perspective. But as soon as external help is required, the self-help argument against Rose Prevention is weakened.

A further problem facing the self-help argument for treatment is that it invites moralism. In order to know if potential beneficiaries of prevention are candidates for external help on this argument, one needs to arrive at justified beliefs about what is or is not in individuals' control. And to know whether the argument sanctions the rescue of someone who now needs treatment, one must assess whether or not she could have easily prevented her unfortunate condition years ago. There is thus a worry that those who are charged with making these determinations will, as humans often do, draw hasty, moralized conclusions about others' degree of responsibility for their own plight. This worry about moralism blends seamlessly into the large philosophical literature on whether and when it is appropriate to hold individuals responsible for their health.² Delving into that discussion is not possible here, but neither is it necessary, in my view. For the two problems already mentioned—that the self-help argument can cut against treatment as well as prevention, and that self-help may not be in any given individual's interests and thus may not be within any individual's psychological reach—together suffice to block the self-help argument in favor of rescue treatment. Most people need an outside incentive or nudge or threat of some kind to make the effort to reduce their risks by a tiny amount. Whether public policy should allocate scarce resources to implement such measures is precisely the question we are asking. Since we cannot simply assume that individuals can nudge themselves along, the real-world situation is, in the end, more similar to the Firing Squad 2 case than it is to the science-fiction case involving seatbelts. Although it is not an iron law of nature that the preventable deaths will occur if treatment is prioritized, it is highly likely that they will, and their occurrence arguably should not be attributed to a failure of individual rationality or of personal responsibility.

² See, *inter alia*, Wikler (2002), Daniels (2001), and Voigt (2013).

Conclusion

I have examined the case for and against using Rose Prevention to avert premature deaths when doing so means neglecting the urgent needs of a smaller number of people needing rescue now. I should stress that nothing I have said suggests that we should invoke Rose Prevention to give each of one-hundred ninety year-olds ten more years of life when we could instead give one twenty year-old sixty more years of life. Even holding the age and potential life-expectancies of beneficiaries fixed, Rose Prevention raises distinctive moral issues on its own, and it is these I have focused on in this paper. I therefore also do not claim to have provided a comprehensive analysis of the myriad moral tensions that fall under the umbrella of “prevention versus treatment.” The moral status of Rose Prevention is but one issue within a much wider web of related issues. Still, I believe we can conclude, on the basis of what I have argued here, that when all else is equal, the moral importance of Rose Prevention should not be discounted simply because it works by providing tiny reductions in already low risks.

REFERENCES

- Brock, D.W., and Wikler, D. (2006). ‘Ethical Issues in Resource Allocation, Research, and New Product Development’, in Jamison, D.T., Breman, J.G., Measham, A.R., Alleyne, G., et al. (eds.) *Disease Control Priorities in Developing Countries*. Washington, D.C.: The International Bank for Reconstruction and Development / The World Bank, pp. 259-270.
- Brock, D.W., and Wikler, D. (2009). ‘Ethical Challenges in Long-term Funding for HIV/AIDS’. *Health Affairs*, 28(6): 1666-1676.

Please cite only final published version

Broome, J. (1990). 'Fairness'. *Proceedings of the Aristotelian Society*, 91: 87-101.

Broome, J. (1998). 'Kamm on Fairness'. *Philosophy and Phenomenological Research*, 58(4): 955-961

Broome, J. (1994). 'Fairness Versus Doing the Most Good'. *Hastings Center Report*, 24(4): 36-39.

Broome, J. (2002). 'All Goods Are Relevant' in Murray, C., et al. (eds.) *Summary Measures of Population Health*. Geneva: World Health Organization, pp. 727-9.

Daniels, N. (2001). 'Social and Individual Responsibility for Health,' in Knight, C., and Stemplowska, Z. (eds.) *Responsibility and Distributive Justice*. Oxford: Oxford University Press, pp. 266-86.

Daniels, N. (2012). 'Reasonable Disagreement about Identified vs. Statistical Victims'. *Hastings Center Report*, 42(1): 35-45.

Fischhoff, B., and Kadvany, J. (2011). *Risk: A Very Short Introduction*. Oxford: Oxford University Press.

Hope, T. (2001). 'Rationing and Life-Saving Treatments: Should Identifiable Patients Have Higher Priority?' *Journal of Medical Ethics*, 27(3): 179-185.

Hope, T. (2004). *Medical Ethics: A Very Short Introduction*. Oxford: Oxford University Press.

Rose, G. (1985). 'Sick Individuals and Sick Populations'. *International Journal of Epidemiology*, 14(1): 32-38.

Rose, G. (2008). *Rose's Strategy of Preventive Medicine*. Oxford: Oxford University Press.

Voigt, K. (2013). 'Appeals to Personal Responsibility for Health: Reconsidering the Luck Egalitarian Perspective'. *Cambridge Quarterly of Healthcare Ethics* 22(2): 146-158.
[Erratum: *Cambridge Quarterly of Healthcare Ethics* 22(3) (2013): 328-9.]

Wikler, D. (2002). 'Personal and Social Responsibility for Health'. *Ethics & International Affairs*, 16(2): 47-55.