

PROBABILITY, NORMALCY AND THE RIGHT AGAINST RISK IMPOSITION

Martin Smith

Many philosophers accept that, as well as having a right that others not harm us, we also have a right that others not subject us to a *risk* of harm. And yet, when we attempt to spell out precisely what this “right against risk imposition” involves, we encounter a series of notorious puzzles. Existing attempts to deal with these puzzles have tended to focus on the nature of *rights* – but I propose an approach that focusses instead on the nature of *risk*. The key move is to distinguish two different ways in which to conceptualise the risk that a given action presents – one of which is linked to the notion of *probability* and the other to the notion of *normalcy*.

1. The Risk Thesis and the High Risk Thesis

Consider the following case of “pure” risk imposition.¹ Suppose A plays Russian roulette on B. That is, suppose A takes a revolver, inserts a bullet into one chamber, spins the cylinder, aims at B’s head and pulls the trigger. Suppose that, as it happens, the chamber that rotates into the firing position when the trigger is pulled is empty, and the gun doesn’t discharge. Suppose, further, that B is asleep or otherwise unaware of what is happening and, as a result, experiences no fear or distress.

We can all agree that, other things being equal, A’s action is morally impermissible. But surely B wouldn’t merely regard this as an “impermissible” action – an action that A “ought not to have performed” – he would see this as a *violation of his rights*. And A’s action has many of the telltale signs of a rights infringement: A’s action could never be justified on purely hedonistic grounds – it could never be justified on the grounds that it would bring *pleasure* to A or to others, irrespective of the amount of pleasure that might be derived. Given the opportunity B, or a third party, would have been morally permitted to use force – even extreme force – against A in order to prevent him from undertaking the action. Finally, in the absence of a strong reason or excuse, it would be legitimate for A to be punished – perhaps even severely punished – for his conduct.²

¹ Thomson “Imposing risks”, p126.

² On the permissibility of preventive force and punishment in this kind of case, see Bergelson, “Self-defense and risks”, §3, Thomson “Some questions about government regulation of behavior”, §4. Many theorists agree that one of the crucial roles of a right is to legitimise defensive and punitive actions when the right is threatened or infringed (see, for instance, McKerlie, “Rights and risk”, pp241-242, Thomson, “Some questions about government regulation of behavior”, §2, *The Realm of Rights*, p2, ch.14, §5). If it is legitimate for an action to be punished by the state, or for a person to use force to prevent it, this is treated here as defeasible evidence that

But what right of B's could have been infringed by A's action? It's natural to think that we each have a right not to be harmed by others³ – but, by stipulation, B has not suffered any actual harm at the hands of A (that's why the risk imposition is described as "pure").⁴ Perhaps the most obvious suggestion is that, in addition to the right that others not harm us, we also have a right that others not subject us to a *risk* of harm – and this is the right that A infringes. Call this the *Risk Thesis*.⁵

Although it provides a straightforward treatment of this example, the Risk Thesis faces an immediate problem – many of the ordinary activities we engage in every day will impose *some* risk of harm on others. If, for instance, A drops a piece of bread into his toaster and presses down the lever, there is some risk that this could cause a fire in which his neighbour B dies.⁶ But presumably B has no right that A refrain from making toast. The risk that B would die as a result of A making toast is of course very low – at least in the order of one in billions – and this suggests an obvious fix: perhaps our right against risk imposition only applies to risks that are relatively *high* or significant, like the risk imposed by Russian roulette. The idea, more precisely, is that we have a right that others not subject us to a high risk of harm – that others not act in such a way that the risk of our being harmed, as a

the action infringes the rights of another. Furthermore, many theorists have claimed that rights serve to *trump* certain justifications for action – including justifications that cite personal pleasure or preference (Dworkin, *Taking Rights Seriously* ch.4, §3, ch.7, Nozick, *Anarchy, State and Utopia*, ch.3, Railton, "Locke, stock and peril", p189). If an action could never be justified on these grounds then this will also be taken as a defeasible indication of a rights infringement.

³ If "harm" is construed broadly then there may be certain harms that one can inflict on others without infringing their rights (see, for instance, Thomson, "Some questions about government regulation of behavior", §2). I put this issue to one side here – "harm" in the main text can be read as restricted to physical injury and death.

⁴ Some have argued that even pure risk impositions constitute harms, on the grounds that they frustrate one's interests or diminish one's autonomy (see Finkelstein, "Is risk a harm?", Oberdiek, "The moral significance of risking", for related discussion see Rowe, "Can a risk of harm itself be a harm?", Thomson, "Some questions about government regulation of behavior", §3, *The Realm of Rights*, p244). This opens up a different way of thinking about a right against risk imposition – and offers the potential of subsuming such a right within a broader right not to be harmed. While this view would require us to reformulate the problems that I will consider, it does not, as far as I can tell, offer any immediate solutions.

⁵ See Holm, "A right against risk imposition and the problem of paralysis", p918, McCarthy, "Rights, explanation and risks", p208, Song, "Rights against high-level risk impositions", p765, Thomson, *The Realm of Rights*, p243.

⁶ B will, of course, be subject to a base level risk of dying in a fire even if A does not make toast. The risk that A imposes is that of B dying in a fire as a result of, or in a way that is caused by, A making toast. The general point that many day-to-day activities impose risks of harm on others is familiar in discussions of the ethics of risk imposition. See, for instance, Fried, *An Anatomy of Values* pp192-193, Hayenhjelm and Wolff "The moral problem of risk impositions", Holm, "A right against risk imposition and the problem of paralysis", McCarthy, "Rights, explanation and risks", Railton "Locke, stock and peril", p207, Song, "Rights against high-level risk impositions", Thomson, *The Realm of Rights*, ch.9, §6.

result of their action, exceeds a threshold t . This is sometimes referred to as the Threshold Risk Thesis or *High Risk Thesis*.⁷

While it may appear more promising than the original Risk Thesis, the High Risk Thesis faces at least two problems of its own. The first problem concerns what we might call cases of “low-risk” Russian roulette.⁸ Suppose a bullet is placed in a single chamber of one out of a set of otherwise empty revolvers. Suppose A chooses a revolver at random, spins the cylinder, aims at B’s head and pulls the trigger. The larger the set of revolvers, the lower the risk of harm that A imposes on B. If the set were sufficiently large, the risk could be lower than any positive threshold, and could even be lower than the risk imposed by making toast. Suppose, for argument’s sake, that we set the threshold at one-in-five-hundred-thousand. In this case, if A chooses from, say, one hundred thousand revolvers when playing Russian roulette on B, his actions would not infringe the right posited by the High Risk Thesis. And yet, A’s action still has all of the hallmarks of a rights infringement. In spite of the number of revolvers, A’s behaviour could never be justified on the grounds that A, or others, find it enjoyable. In spite of the number of revolvers, it would be permissible to use force to prevent A’s action, and legitimate for A to be punished if he has no strong reason or excuse (I’m inclined to think that even extreme force and severe punishment could still be warranted).

The second problem for the High Risk Thesis concerns cases of “distributed risk” – cases in which there is a high risk that some member of a group will be harmed, even though the risk to each individual member is low. Consider the following example due to McCarthy.⁹ Suppose A is considering two options for disposing of a large quantity of a toxic chemical. First, he could surreptitiously dump the chemical into a pond that he shares with his neighbour B. Second, he could surreptitiously dump the chemical into the river that flows through his property, even though there are a million people who live downstream. The former option involves a high risk – say a one-in-a-thousand chance – that B will be exposed to a harmful quantity of the chemical. The latter option involves a *very* high risk that at least one of the people living downstream will be exposed to a harmful quantity of the chemical, even though the risk to each *individual* person is low – say a one-in-a-million chance.

⁷ See Holm, “A right against risk imposition and the problem of paralysis”, p920, McCarthy, “Rights, explanation and risks”, p212, Song, “Rights against high-level risk impositions”, p767, Thomson, *The Realm of Rights*, p245.

⁸ Nozick, *Anarchy State and Utopia*, p73, Thomson, “Some questions about government regulation of behavior”, §4.

⁹ McCarthy, “Rights, explanation and risks”, pp213-214. The example is adapted from McKerlie, “Rights and risk”, pp247-248. The fact that cases of distributed risk pose a potential problem for the High Risk Thesis is observed by Railton, “Locke, stock and peril”, pp209-210.

If a one-in-a-thousand chance is above the threshold then, according to the High Risk Thesis, B's rights would be infringed if A dumped the chemical in the pond. If a one-in-a-million chance is below the threshold then, as far as the High Risk thesis is concerned, there is no person whose rights would be infringed if A dumped the chemical in the river. All else equal, then, dumping the chemical in the river would be the morally preferable option, as this involves no rights infringements. But this prediction seems incorrect – after all, dumping the chemical in the river involves a much higher *overall* risk of harm. If there are a million people living downstream who are each subjected to a one-in-a-million risk of harm then, assuming these risks are independent, the risk of at least one person being harmed works out to approximately *sixty-four percent*. The High Risk Thesis appears, then, to create a dubious moral preference for cases in which risk is distributed amongst a group of individuals (the river-option) over cases in which risk is imposed upon a single individual (the pond-option).¹⁰

2. Revisiting the Risk Thesis

In response to these problems, McCarthy abandons the High Risk Thesis and advocates a return to the original Risk Thesis, on which *any* risk imposition, no matter how slight, constitutes a rights infringement¹¹. As McCarthy observes, the Risk Thesis offers a more satisfactory treatment of the toxic chemical case. The Risk Thesis predicts that dumping the chemical in the river would involve a million rights infringements, while dumping it in the pond would involve only one, leading to an immediate reversal of the above verdict; all else equal it is the river-option that would now be reckoned to be morally worse. More generally, the Risk Thesis predicts that a risk imposition cannot be made more morally acceptable by distributing the risk amongst a number of individuals – on the contrary, this will simply introduce further rights infringements.

The Risk Thesis still faces a basic problem of course. As discussed above, it predicts that many of our day-to-day actions will infringe others' rights. According to McCarthy, though, this result is only problematic if it leads to the conclusion that many of our day-to-day actions are *morally impermissible*. We can infer the latter from the former if we assume that rights are *absolute*, and can never be

¹⁰ While my focus here is on the problem of spelling out a right against risk imposition, it's worth noting that *any* approach to the ethics of risk imposition, whether or not it assigns a prominent role to such a right, will still need to thread its way through cases of Russian roulette vs low-risk Russian roulette vs day-to-day risk imposition and through cases of distributed vs non-distributed risk. While other approaches are beyond the scope of the discussion here, I will mention some in passing (n21 and n32).

¹¹ McCarthy, "Rights, explanation and risks".

permissibly infringed – but, according to McCarthy, this view is untenable.¹² Suppose I suddenly fall ill and you possess a large quantity of a drug that I need to save my life. While I have the opportunity to take the required amount from your stockpile, the situation is so urgent that I don't have time to seek your permission or to procure the drug elsewhere. If I take the drug then I infringe your rights, and do so knowingly – after all, I know that the drugs belong to you and you have a right that others not take them without your permission. Nevertheless, if this is the only way in which I can save my life, it seems that my action is morally permitted.

In McCarthy's view, an action that infringes the rights of another will be morally permissible if the reasons in favour of performing it sufficiently outweigh the burden to the bearer of the right.¹³ So even if many of our day-to-day actions infringe the rights of others, as the Risk Thesis implies, these actions may yet be permissible, provided they are backed by sufficiently strong reasons. But *are* they backed by sufficiently strong reasons? Think again of the toast example. By making toast A imposes a low risk of injury or death upon his neighbours and, according to the Risk Thesis, he thereby infringes their rights. And yet, the reasons in favour of making toast are, by and large, pretty trivial in which case, if we are to have the desired result that the activity is morally permissible, then these rights infringements would have to be *more trivial still*. But there is something jarring about the idea that another person's rights could count for so little. This is sometimes referred to as the "cheapening of rights" problem.¹⁴ One could perfectly well reject absolutism about rights – perhaps on the strength of examples like McCarthy's drug case – and still insist that a rights infringement could never be made permissible by something like a *desire for toast*. That is, one could insist that a desire for a piece of toast (rather than a slice of bread) is not the kind of thing that could ever sufficiently outweigh the burden of having a right infringed.

¹² McCarthy, "Rights, explanation and risks", §3. Nozick comes close to endorsing an absolutism about rights, arguing that it is impermissible to infringe a person's rights even if one could reduce the total number of rights infringements thereby (Nozick, *Anarchy, State and Utopia*, pp28-30). Even Nozick allows, however, that it may be permissible to infringe another's rights in order to avert "catastrophic moral horror". For discussion, and criticism, of Nozick's near-absolutism, see Thomson "Some ruminations on rights". The existence of some absolute rights is defended by Gerwith, "Are there any absolute rights?".

¹³ See McCarthy, "Rights, explanation and risks", pp209-210, see also Thomson, "Some ruminations on rights", *The Realm of Rights*, ch.6. On one terminology (Gerwith, "Are there any absolute rights?", Thomson, "Some ruminations on rights") a right is *violated* just in case it is impermissibly infringed – infringed without sufficient justification. We might say, then, that my taking the drug from your stockpile would, under the circumstances, constitute an infringement, but not a violation, of your rights. Absolutism, in this terminology, can be expressed by saying that all infringements are violations. For discussion of the violation/infringement distinction see Oberdiek "Lost in moral space".

¹⁴ See Song, "Rights against high-level risk impositions".

Here is another way to put the worry: while absolutism about rights will take us from the premise that many of our day-to-day actions infringe the rights of others to the conclusion that many of our day-to-day actions are morally impermissible, it is not the *only* way to bridge this gap. One supposition I have been taking for granted so far is that an act which infringes another person's rights can never be justified solely on the grounds that it will bring *pleasure* to oneself or to others.¹⁵ This is clearly much weaker than absolutism about rights – and is perfectly consistent with McCarthy's preferred verdict about the drugs case – but it is *inconsistent* with the idea that many of our day-to-day actions permissibly infringe others' rights. Many of our day-to-day actions (like making toast) have no discernible benefit other than to bring some small pleasure.¹⁶ In any case, I won't pursue this further here – for McCarthy's defence of the Risk Thesis faces another, perhaps even more serious, objection.

3. The Role of Intentions?

Although the Risk Thesis appears to give the right verdict in the toxic chemical case, the low-risk Russian roulette case continues to pose a problem – though somewhat subtler than the problem it poses for the High Risk Thesis. If A plays low-risk Russian roulette on B then the Risk Thesis, unlike the High Risk Thesis, will straightforwardly predict that A infringes B's rights. However, given that A *also* infringes B's rights when he makes toast, we are still in need of some explanation of the blatant moral difference between the two actions. As mentioned above, even if the only benefit of making toast is to bring some small pleasure to A, the action is clearly permissible. But it is clearly *impermissible* for A to play low-risk Russian roulette on B, no matter the pleasure he might derive by doing so. As discussed, it would be permissible for one to use force to prevent A from playing low-risk Russian roulette on B and, absent a strong reason or excuse, legitimate for A to be punished for such an action.

¹⁵ There are a number of different views as to what rights are and what kind of moral significance they carry – but many would agree that rights must be something *more* than just one further ingredient in the balance of considerations that bear upon the moral permissibility of an action. As observed in n2, many theorists endorse the metaphor of rights as “trumps” which would seem to require, at a minimum, that there be some considerations that can count in favour of an action, but which could never outweigh or counterbalance a rights infringement. This, at the very least, illustrates that there is a broad conceptual space between absolutism about rights, and the view that rights infringements can in principle be justified by any action-favouring considerations whatsoever.

¹⁶ Even those who grant that a rights infringement can be weighed against goods such as pleasure may still deny that it could ever be justified by a small or trifling pleasure. Thomson (*The Realm of Rights*, p153, n2) resists the metaphor of rights as trumps, but suggests that they might still be considered “high cards”.

Obviously, one cannot legitimately punish A for making toast or permissibly use force to prevent him from doing so.

Picking up on a suggestion from Thomson, McCarthy proposes that the moral difference between these two actions lies in A's *intentions*.¹⁷ To play Russian roulette on an innocent person is to *intend* to impose a risk of death – this seems to be the very point of the action. In contrast, imposing a risk of death is not the point of making toast – while one may be *aware* of this risk, it is not intended. The reason it is impermissible for A to play low-risk Russian roulette on B, according to McCarthy, is that this action involves an intentional imposition of risk and, thus, intentionally infringes B's rights.¹⁸ In contrast, when A makes toast, although B's rights may be infringed, the infringement is not intended (but merely foreseen).

On closer inspection, though, this suggestion comes nowhere close to capturing the moral difference between these two actions. Suppose A's toaster is broken, and a third party offers to make him a slice of toast if only he plays low-risk Russian roulette on B. In this example, A has no particular wish to impose a risk of death on B – his only aim is to procure toast. Nevertheless, if A agrees to this, then his actions are hardly better than if he played low-risk Russian roulette on B with the express aim of imposing risk. It is still the case that A could be punished for such an action, and it is still the case that B, or a third party, could permissibly use force to stop him. The wrongness of playing Russian roulette on an innocent person has little to do with one's intentions – a willingness to impose this risk for a trivial payoff is little better than a direct desire to impose it.¹⁹

¹⁷ See McCarthy "Rights, explanation and risks", pp211-212 and Thomson, "Some questions about government regulation of behavior", §4. As far as I'm aware, it was Nozick who first offered a hypothesis about the moral difference between low-risk Russian roulette and an activity like making toast (Nozick's examples are mining, running trains and driving). According to Nozick, what distinguishes the former action is that it has no value for society and is not a normal and/or important part of people's lives (Nozick, *Anarchy, State and Utopia*, pp73, 82). Nozick doesn't elaborate, but the suggestion does not appear promising. Even if making toast were a rare practice, and few people owned toasters, it's dubious that this would make any moral difference to the activity – and it certainly wouldn't make it into the moral equivalent of playing low-risk Russian roulette on innocent people. We could also imagine a situation in which something like low-risk Russian roulette *was* part of an entrenched social practice to which people assigned importance (as in Shirley Jackson's *The Lottery*). Once again, it's doubtful that this would make much, if any, difference to the moral status of the activity.

¹⁸ Holm ("A right against risk imposition and the problem of paralysis", pp921-922) defends a variant on this: what makes it impermissible for A to play low-risk Russian roulette on B are A's *reasons for action*, which include the fact that the action will impose a risk upon B. When A makes toast, the fact that this imposes a risk upon B is not one of the reasons for which A acts. Holm's proposal is equally subject to the objections in the main text.

¹⁹ One might suggest that, although the imposition of risk is not A's ultimate aim in the new example, it is still being used as a *means*, and this is enough for it to count as *intended* – an idea reflected in certain formulations of the doctrine of double effect (see, for instance, Nelkin and Rickless, "Three cheers for double effect", §1

As well as imagining a case in which A playing low-risk Russian roulette on B does not involve an intentional imposition of risk, we could also imagine a case in which A making toast *does* involve an intentional imposition of risk. Suppose A doesn't want toast at all and uses his toaster solely to impose some small risk of death upon his neighbour B. While A's motives are certainly criticisable, and would seem to reflect poorly on his moral character, it is plausible that his *action* is nevertheless a permissible one.²⁰ And, even if we do insist that A has acted impermissibly, the action is hardly the moral equivalent of playing low-risk Russian roulette on B. If B became aware of A's reasons for putting on his toaster, he may be perturbed by A's apparent maliciousness towards him, but it wouldn't be proportionate for him to use *force* in order to prevent the action. (If B *were* to use force against A – knock him out, break his fingers, even just smash his toaster – our sympathies, in this story, would quickly switch from B to A). Similarly, it would seem cruel and vindictive to *punish* A for putting on his toaster. The overwhelming sense is that, while there is clearly something morally amiss about A's state of mind, the action itself is essentially *harmless* and doesn't warrant any strong response.

The moral difference between A playing low-risk Russian roulette on B and A making toast in the house next to B's cannot be captured in the way that McCarthy proposes. What *would* capture this moral difference is the verdict that the former action involves an infringement of B's rights, while the latter does not. But this, of course, is the very prediction that neither the Risk Thesis nor the High Risk Thesis seems able to deliver. The right posited by the Risk Thesis is infringed by both of these actions, while the right posited by the High Risk Thesis is not infringed by either.

Scanlon, *Moral Dimensions*, p14, Thomson, "Physician assisted suicide", pp512-513). I'm unsure if this suggestion is right – while low-risk Russian roulette is obviously being treated as a means by A, it is less clear whether the imposition of risk *per se* is playing this role. Whatever the truth, the motivational structure of this example is intended to mirror that of the original toast case. In both cases, A's only aim is to get toast. In both cases A employs a means to this end – be it using the toaster or playing low-risk Russian roulette – which he foresees will impose a risk of death on B. In neither case is A motivated by this imposition of risk – he may even regard it as regrettable (though obviously not so much as to make him reconsider). If the risk imposition counts as intended in the new example, it must also count as intended in the toast example – and whatever makes for the moral difference between the cases is not to be found in A's intentions.

²⁰ Some philosophers insist on a sharp divide between the moral evaluation of an *action* and the moral evaluation of an *agent*, with intentions and reasons bearing upon the latter but not the former (see, for instance, Oberdiek, "The moral significance of risking", section 3A, Scanlon, *Moral Dimensions*, partic. ch.1, Thomson, "Physician assisted suicide", §4, §5). To undermine McCarthy's strategy, it is not necessary that we endorse this general view – it is enough to maintain that one's intentions are not relevant to the moral permissibility of making toast or of playing low-risk Russian roulette. In fact, even *this* is not strictly necessary, so long as we maintain that whatever difference one's intentions make here is not enough to bridge the moral gulf between these actions.

4. Revisiting the High Risk Thesis

I will now outline a way of delivering the desired verdict. The first step is to highlight a tacit assumption that has guided the discussion so far. Return to the case of low-risk Russian roulette: a single bullet is placed into a single chamber of one out of a large set of revolvers, before A chooses a revolver at random, spins the cylinder, points at B and pulls the trigger. So far, this has been classified as a 'low-risk' scenario on the grounds that the *probability* of B being killed is very small. And yet, if we were actually watching these events unfold, the thing that would surely shock us – and move us to intervene if we could – is the perceived *riskiness* of what A is doing. It would be natural to have something like the following thought; the bullet has to be located in some chamber, and it would be just as *normal* for it to be in any one chamber as any other – including the chamber that slides into alignment with the barrel of A's revolver when he pulls the trigger. If B were shot and killed then, given the nature of the set-up, we wouldn't need any special explanation as to how this could have happened. One who is struck by this thought wouldn't be altogether reassured by learning how *many* revolvers were in the initial set. The more revolvers there are, the more places the bullet could end up – but there is still nothing *preventing* it from being in the one chamber that would result in B's death.

Here is another way to put the point: the most *normal* possible worlds in which A plays Russian roulette on B will include worlds in which the bullet is in each of the available chambers. As a result, some of the most normal possible worlds in which A plays Russian roulette on B will be worlds in which B is killed – this represents one normal outcome of the action. Clearly, the notion of normalcy that is being invoked here is distinct from the idea of statistical frequency – B's death is not an outcome that would *frequently* arise from this action, were it repeated over and over. Rather, this outcome is normal in the same sense that it would be normal for, say, "10, 7, 13, 8, 25, 19" to be the winning lottery numbers – some sequence of numbers has to come up, and this sequence would require no more explanation than any other.

When it comes to A putting on his toaster, however, the situation seems altogether different. While it is *possible* that this action could cause a fire which leads to the death of his neighbour B, there is no sense in which this would count as a *normal* outcome of the action. On the contrary, there would have to be some explanation as to how the fire started (was there an electrical fault in the toaster or in the wiring of A's house?), how it took hold (was there inflammable material around the toaster, was there a gas leak?), how A failed to extinguish the fire or raise the alarm (was he asleep, did he leave the house?) and so on. If we were told that B had died in a fire as a result of A putting on his toaster our immediate reaction would be to ask *how* this could have possibly happened. If we were told that B had died as a result of A subjecting him to Russian roulette, our reaction would be quite different –

no matter how many revolvers A was choosing from. I think that this contrast in our reactions is tracking a genuine difference between these two hypothetical events. Of all the outcomes which could result from A putting on his toaster, B's death is a highly abnormal one. In possible worlds talk, the most normal worlds in which A puts on his toaster are worlds in which B suffers no harm as a result, and any worlds in which this action leads to B's death are highly abnormal.

It has been taken for granted, in the discussion so far, that the risk of a given outcome is determined by its probability – the greater the probability, the greater the risk, and the lower the probability, the lower the risk. It has also been assumed, accordingly, that any risk threshold we use in spelling out the High Risk Thesis must take the form of a probability value (such as one in-five-hundred-thousand). This *probabilistic* conception of risk is entrenched across a range of areas, and has been largely assumed, unquestioned, in discussions of the right against risk imposition.²¹ But this conception of risk is not inevitable – and should be seen as another potential moving part in the puzzles we have been considering. The dominance of the probabilistic account of risk has recently been challenged by several authors, who have put forward alternatives such as the *modal* account, the *relevant alternatives* account, and the *normic* account – which will be my focus here.²²

According to the probabilistic account, the risk that a particular outcome would result from a given action depends on how probable it is that the outcome would result from the action.²³ According

²¹ And, indeed, throughout the literature on the ethics of risk imposition. According to one well-known family of views, we are morally required to act in a way that, roughly speaking, minimises the strongest individual complaint against our action. On the “ex ante” version of this view, when an action imposes a risk of harm upon an individual they have a complaint against it, the strength of which is discounted according to the level of risk involved (see for instance Frick, “Contractualism and social risk”, Kumar, “Risking and wronging”, for critical discussion see Horton, “Aggregation, complaints and risk”. But it is assumed, when determining the strength of such a complaint, that the level of risk is to be measured *probabilistically* (see for instance Frick, “Contractualism and social risk”, p188, Horton, “Aggregation, complaints and risk”, §2, Kumar, “Risking and wronging”, §4A). This will make the view ill-equipped to handle cases of low-risk Russian roulette vs making toast – and also leads to a problem with cases of distributed risk. I won't explore this further here.

²² For the modal account, see Pritchard, “Risk”. For the relevant alternatives account see Gardiner, “Relevance and risk”. For the normic account see Ebert, Smith and Durbach, “Varieties of risk”, Smith “Decision theory and de minimis risk”.

²³ The probability in question could be understood as “objective” – determined perhaps by the frequency with which the outcome would accompany the act or some such. My own view is that the probability is best understood as epistemic or *evidential* – the probability that a given outcome will eventuate, given the evidence that the action has been performed. There are important questions about how much should be included in the relevant description of an action – and different answers may, of course, give rise to different assessments of the risk that the action poses. In the examples I consider here, the relevant descriptions seem relatively clear – but there will undoubtedly be more difficult cases and, arguably, a more principled approach to this issue would be needed for any complete ethics of risk imposition.

to the normic account, the risk that a particular outcome would result from a given action depends on how *abnormal* it would be for the outcome to result from the action. As above, the notion of normalcy at work here is linked with the need for explanation – an outcome is abnormal to the extent that it requires special explanation, in terms of factors that are additional to the action. On the probabilistic account, when A plays Russian roulette on B, the risk to B depends upon the number of revolvers involved and can, as a result, be made arbitrarily close to zero. On the normic account, things look altogether different; given the set-up, no special explanation would be needed if B were shot and killed – no matter how many revolvers are involved, this will represent one of the normal outcomes of the action.

Suppose, as hinted above, that possible worlds can be ranked according to their normalcy – the most normal worlds are assigned a rank of zero, the next most normal worlds are assigned a rank of one and so on.²⁴ Suppose an action could, in principle, result in harm to a given individual. If the most normal worlds in which the action is performed include worlds in which the individual is harmed, then this outcome will have an abnormality of zero, given the action – it will, in short, represent one of the normal outcomes of the action. If the individual does not come to harm in any of the most normal worlds in which the action is performed, then this will not be a normal outcome of the action – and its abnormality may be gauged by the difference in rank between the most normal worlds in which harm results from the action and the most normal worlds in which the action is performed. If the former worlds are one rank more abnormal than the latter then the abnormality of the individual suffering harm, given the action, will be equal to one. If the former worlds are seven ranks more abnormal than the latter, then the abnormality of the individual suffering harm, given the action, will be equal to seven, and so on.

Amongst the most normal worlds in which A plays Russian roulette on B are worlds in which B is shot and killed. On the normic account of risk, when A plays Russian roulette on B, the risk to B is *maximal*, as B's death has an abnormality of zero, given A's action. While I have spoken of cases of "low-risk" Russian roulette (and will, for ease, continue to use that term), on a normic interpretation there is, in effect, *no such thing* as low-risk Russian roulette – the normic risk is maximal, no matter how many revolvers are involved. In contrast, when A makes toast, this will count as a low-risk activity in both the probabilistic and normic senses. While there are possible worlds in which A's putting on his toaster results in a fire in which B dies, these worlds are highly abnormal.

²⁴ See Smith, *Between Probability and Certainty*, ch.8, "The logic of epistemic justification", "The hardest paradox for closure".

It is important to emphasise that I am not proposing the normic account of risk as a *competitor* to the probabilistic account. In my view, both probabilistic risk and normic risk represent legitimate ways of precisifying our ordinary risk concept.²⁵ In fact, this kind of pluralist approach fits well with the example of low-risk Russian roulette – in which our intuitions about the risks involved do, arguably, pull in different directions. On the one hand, it's intuitive that the risk to B diminishes as the number of revolvers is increased – that the risk is halved if two revolvers are used instead of one, and halved again if four revolvers are used etc. This is why low-risk Russian roulette may be morally preferable to standard Russian roulette. On the other hand, it's intuitive that, irrespective of the number of revolvers involved, the risk to B is greater than that imposed by A making toast or engaging in other everyday activities. This is why low-risk Russian roulette will never be the moral equivalent of an everyday activity. While the former intuition is captured by the probabilistic account, the latter is captured by the normic account.²⁶

For the original Risk Thesis, it makes no difference whether the “risk” in question is interpreted probabilistically or normically. If an action involves *some* probabilistic risk of harm, then there must be a possible world in which harm results from the action, in which case the action will also involve some normic risk of harm and vice versa.²⁷ When it comes to the High Risk Thesis, however, the two different ways of disambiguating the notion of risk give rise to two distinct theses: the *Probabilistic* High Risk Thesis (which we have been taking for granted so far) and the *Normic* High Risk Thesis.

Probabilistic High Risk Thesis We have a right that others not subject us to a high probabilistic risk of harm. More precisely, we have a right that others not act in such a way that the probability of our being harmed, as a result of their action, is above a threshold t .

²⁵ Ebert, Smith and Durbach, “Varieties of risk”, §6.

²⁶ More precisely, the following three claims are inconsistent if “risk” is given the same interpretation in each:

- (i) The risk involved in low-risk Russian roulette is halved when the number of revolvers is doubled.
- (ii) Low-risk Russian roulette always involves a greater risk than making toast.
- (iii) Making toast involves some positive level of risk.

²⁷ This assumes that every possible world is assigned some normalcy rank and some non-zero probability. The latter assumption is typically dropped in the case of infinite probability spaces and, without it, the existence of possible worlds in which an action results in harm will be consistent with the probability of harm being zero, conditional upon the action being performed. In this case, the “vice versa” direction of the above will fail – an action could present some normic risk of harm without presenting any probabilistic risk of harm – though it's doubtful that this would make any difference in practice.

Normic High Risk Thesis We have a right that others not subject us to a high normic risk of harm. More precisely, we have a right that others not act in such a way that the abnormality of our being harmed, as a result of their action, is below a threshold t .

Unlike the Risk Thesis and the Probabilistic High Risk Thesis, the Normic High Risk Thesis can separate the act of playing low-risk Russian roulette on an innocent, and the act of making toast. While the right posited by the Risk Thesis is infringed by both acts, and the right posited by the Probabilistic High Risk Thesis is infringed by neither, the right posited by the Normic High Risk Thesis, given an appropriate choice of threshold, will be infringed by the first but not the second.

Having distinguished between probabilistic and normic risk, one might wonder *why* it is the second kind of risk that should figure in our right against risk imposition. Why should we have a right that others not subject us to a high *normic* risk of harm? On first impressions, this claim might seem rather mysterious. My primary aim here is to argue that, by understanding the right against risk imposition in normic terms, we are able to solve a number of problems that arise for rights-based approaches to the ethics of risk imposition. Questions about the foundation or basis of such a right lie, for the most part, beyond the scope of this paper – but I will conclude this section with one speculative line of thought, based on a connection between normic risk and the limits of our responsibility for the consequences of our actions.

Suppose A's decision to make toast really *did* lead to a fire in which his neighbour B died. In this case, it's plausible that this outcome would not be wholly attributable to A's action – for it would owe in part to circumstances (be they faulty wiring, a gas leak etc.) that lie completely outside of A's awareness and control. As a result, A would not be considered *fully responsible* for B's death. Similar remarks apply to any action that presents a low normic risk of harm. If there is a low normic risk that an action will cause harm to an individual, then harm could only result through the intervention of independent, interfering factors, which would serve to mitigate the agent's responsibility.

There is no equivalent connection between responsibility and probabilistic risk. Even if an action presents a low probabilistic risk of harm, one may still bear full responsibility for any harm that ensues. If A plays low-risk Russian roulette on B then, irrespective of the number of revolvers involved, A would be fully responsible in the event that B were shot and killed. These brief remarks don't of course amount to a full explanation of why the Normic High Risk Thesis should be true – but they do perhaps dispel some of the mystery which might otherwise surround it.²⁸

²⁸ Another point to bear in mind is that the Probabilistic and Normic High Risk Theses are not formally inconsistent. While I have been assuming (as seems standard in the literature) that there is at most *one* right

5. Distributed Risk

In section 1 I presented two problems for the High Risk Thesis – one of which concerned cases of low-risk Russian roulette and the other of which concerned cases of distributed risk. In the previous section I argued that, on a normic interpretation, there is no such thing as a case of low-risk Russian roulette, and the problem dissolves. In this section I will argue that the same is true for cases of distributed risk – on a normic interpretation, such cases simply *cannot arise*. A case of distributed risk, recall, was defined as one in which there is a high risk that some member of a group will be harmed, even though the risk to each individual member is low. Consider a group of individuals C, D, E, F ... and suppose there is a high risk that some member of the group will suffer harm. On the normic account, this means that there is a relatively normal possible world in which some member of the group suffers harm. But any world in which some member of the group suffers harm must either be a world in which C suffers harm, or a world in which D suffers harm ... in which case either C must be at a high risk of harm or D must be at a high risk of harm ... That is, if there is a high normic risk that some member of the group will be harmed, then there must be some member of the group who is at a high normic risk of harm. If the risk to each member is equal, then they will *all* face a high normic risk of harm.

More formally, if x is a variable ranging over the members of some group, H is read “... is harmed” and \diamond is read “There is a high risk that ...” then, on a normic reading, we will have an instance of the “Barcan Formula”: $\diamond\exists xHx \rightarrow \exists x\diamond Hx$. That is, if there is a high risk that some individual in the group is harmed then, on the normic reading, there is some individual in the group who is at a high risk of harm. Since the converse clearly holds, the normic account predicts that the two risk attributions are, in fact, equivalent – it makes no difference whether the high risk operator or the existential quantifier is given wide scope: $\diamond\exists xHx \leftrightarrow \exists x\diamond Hx$.²⁹

What, then, will the Normic High Risk Thesis predict in putative cases of distributed risk such as the toxic chemical case? In the toxic chemical case, we are explicitly told the probabilistic risks associated with each possible action – we are told that, if A dumps the chemical in the pond there is

against risk imposition, there is no logical barrier to accepting *both* theses. And the Normic High Risk Thesis would surely seem less mysterious if the Probabilistic High Risk Thesis were accepted alongside it. In this case, any imposition of high risk would infringe a person’s rights, no matter how the notion of risk is interpreted. This “combined” view might also offer a more satisfactory treatment of certain cases – such as those in which an individual is subjected to a high probabilistic, but low normic, risk of harm (see Smith, “Decision theory and de minimis risk”, §5). I won’t explore this further here.

²⁹ If R is read “...’s rights are infringed” then the High Risk Thesis gives us the conditional $\forall x (\diamond Hx \rightarrow Rx)$, from which we can derive $\exists x\diamond Hx \rightarrow \exists xRx$. If the risk is interpreted normically, we can infer $\exists xRx$ from $\diamond\exists xHx$ (via $\exists x\diamond Hx$). If the risk is interpreted probabilistically, however, then the inference is blocked.

a one-in-a-thousand chance that B will be exposed to a harmful amount, and if A dumps the chemical in the river then, for each of the million people living downstream, there is a one-in-a-million chance that the individual will be exposed to a harmful amount. Obviously, there are no normic risks stipulated – and the details that would be needed to assess these risks are also largely missing in existing descriptions of the case.

Here, perhaps, is one natural way of filling in the required details. If A were to dump the chemical in the pond then, given the quantity and potency of the chemical, the volume of the pond, and the way in which B normally uses the pond water, the amount of chemical to which B is exposed will vary throughout a certain range. While most of the values in this range would result in no ill effects, the highest values would cause harm to B. In this case, if B were to suffer harm as a result of A dumping the chemical in the pond, then no special explanation would be needed – this would be like subjecting B to a kind of low(medium?)-risk Russian roulette.

Similarly, if A were to dump the chemical in the river, then each individual living downstream faces a potential exposure range, given facts about the quantity and potency of the chemical, the volume and flow of the river, and the way in which the river water is normally used. If some of the values in this range are above the harmful level then, once again, for a given individual to suffer harm as a result of A dumping the chemical in the river would not demand special explanation – this would be like subjecting each of these individuals to a kind of low-risk Russian roulette. When the details are filled in like this, all of the normic risks are maximal – if A dumps the chemical in the pond then B is at maximal normic risk of harm, and if A dumps the chemical in the river then every individual downstream is at maximal normic risk of harm. As a result, the Normic High Risk Thesis will predict that, all else equal, it would be morally worse to dump the chemical in the river, as this would involve a million rights infringements, while dumping the chemical in the pond would involve only one.

If, however, we were to alter the case, in such a way that the presence of the chemical in the river would present only a *low* normic risk to each of the people living downstream, then the Normic High Risk Thesis could make a different prediction. What might such a case look like? Suppose a series of measures is in place to prevent the people living downstream from ever coming into contact with the river water – perhaps the land around the river is private and trespassers face penalties or prosecution, perhaps the river is protected by a high fence etc. None of this would make it *certain* that a given person living downstream won't be harmed if the chemical is dumped in the river, but it would generate the need for special explanation in the event that they are. How did they get past the fence? Why were they willing to trespass? And so on. If the details are filled in in this way then, for any individual x living downstream, it would be abnormal for x to be harmed as a result of A dumping the

chemical in the river. If the degree of abnormality is greater than the threshold posited by the Normic High Risk Thesis then, as far as this thesis is concerned, dumping the chemical in the river will involve no rights infringements. If the situation with B and the pond is unchanged from the description given above then, all else equal, the Normic High Risk Thesis will predict that it is morally preferable for A to dump the chemical in the river.³⁰

As we have seen, though, since the normic risk to each individual living downstream is low, so too is the normic risk to the group – dumping the chemical in the river presents a low normic risk of *any* individual being harmed. Not only, then, does dumping the chemical in the pond infringe B's rights, it also involves a higher overall normic risk of harm. If A were to dump the chemical in the pond, then B could suffer harm in a way that is consistent with conditions being normal – this would represent one of the normal outcomes of the action. If A were to dump the chemical in the river then, under normal conditions, there is no individual who would be harmed. Furthermore, if an individual living downstream *were* harmed as a result of A dumping the chemical in the river, then A would not be fully responsible for this harm, as it would be due in part to the individual's own actions – trespassing, scaling the fence, etc.

It might still be the case, of course, that dumping the chemical in the river would involve a higher overall *probabilistic* risk of harm. Indeed, there is no reason why the probabilities could not remain as originally stipulated – given the sheer number of people who live downstream, if A dumps the chemical in the river, there is a sixty-four percent chance that at least one of these people will, for some reason, flout the rules, come into contact with the water and be harmed.³¹ In light of this, some would balk at the idea that facts about normic risk could ever make the river-option morally preferable to the pond-option. Some would insist that, with the probabilities as they are, the river-option would *always* be morally worse.

³⁰ We could of course avoid this result if we were willing to set the abnormality threshold very high. Even if the penalties, fences etc. would make it highly abnormal for any given individual to be harmed as a result of A dumping the chemical into the river, if the abnormality threshold that features in the Normic High Risk Thesis were higher *still*, then this action would nevertheless infringe the rights of those living downstream, and the pond option would remain morally preferable. But the higher we push the abnormality threshold, the more of our ordinary everyday activities will turn out to infringe others' rights and, in the limit, we would end up mired in the same problems that beset the original Risk Thesis. Whatever one thinks about this particular case, I don't think that threshold-raising is viable as a general strategy for dealing with cases of this kind.

³¹ To compensate for the probabilistic effect of the fences, penalties etc. we could imagine that there is a higher probability that any person who comes into contact with the water suffers harm. Alternately, we could achieve the same effect by increasing the number of people who live downstream.

I won't attempt to engage this position here, except to say this: the prediction that the river-option *may* be morally worse, depending on how the non-probabilistic details of the case are filled in, is as close as we can come to the above prediction whilst working exclusively within the framework of individual rights. The only right against risk imposition that will yield the result that the river-option is always morally worse than the pond-option is the right posited by the original Risk Thesis – and, as argued, this thesis is untenable. Those who wish to maintain that the river-option is always morally worse should, I suggest, give up on attempting to derive this result purely from a right against risk imposition. One could still accept the existence of such a right – and still perhaps find explanatory work for it – but would need to argue that, when it comes to comparing these two options, rights infringements are not the decisive factor.³²

6. Conclusion

I have argued that the most promising rights-based approach to the ethics of risk imposition comes in the form of the Normic High Risk Thesis – the claim that we each have a right that others not impose a high normic risk of harm upon us. Unlike the Risk Thesis, the Normic High Risk Thesis doesn't make for rampant, trivial rights infringements. Unlike the Probabilistic High Risk Thesis, the Normic High Risk Thesis doesn't generate a moral preference for cases in which a risk is distributed amongst the members of a group. Unlike either of these theses, the Normic High Risk Thesis is able to account for the moral difference between the risks imposed by making toast and the risks imposed by “low-risk” Russian roulette.³³

³² Those who insist that the river-option is always morally worse than the pond-option may find it natural to appeal to considerations of *expected utility* – where the expected utility of an action is equal to the probability-weighted average of the utilities of its possible outcomes. If an individual suffering harm as a result of the chemical is assigned a constant, finite disutility then, with the probabilities as stipulated, the pond-option will always have a significantly higher expected utility than the river-option. Once we assign expected utilities a moral role, however, one might think that there is no longer any *need* for the Normic High Risk Thesis, or for any right against risk imposition – why not let our moral assessment of a risk imposing activity be determined purely by its expected utility? There are reasons to be dissatisfied with this approach however. This view will, for instance, generate the wrong predictions about low-risk Russian roulette which could, given enough revolvers, have a higher expected utility than making toast (for related discussion see Smith, “Decision theory and de minimis risk”, §3). A full evaluation of this view is beyond the scope of this paper.

³³ This paper was presented (online) at the Risk and Recklessness workshop, University College London in April 2021, at the Epistemology and Normality workshop, Dianoia Institute of Philosophy, Australian Catholic University in January 2022, and at the Imposing Risk workshop, University of Manchester in September 2022. Thanks to all of those who participated on these occasions. Thanks also to two anonymous referees for this

References

- Bergelson, V. (2016) "Self-defense and risks" in Coons, C. and Weber, M. eds. *The Ethics of Self-Defense* (Oxford: Oxford University Press)
- Dworkin, R. (1977) *Taking Rights Seriously* (London: Duckworth)
- Ebert, P., Smith, M. and Durbach, I. (2020) "Varieties of risk" *Philosophy and Phenomenological Research* v101(2), pp432-455
- Finkelstein, C. (2003) "Is risk a harm?" *University of Pennsylvania Law Review* v151(3), pp963-1001
- Frick, J. (2015) "Contractualism and social risk" *Philosophy and Public Affairs* v43(3), pp175-223
- Fried, C. (1970) *An Anatomy of Values* (Cambridge, MA: Harvard University Press)
- Gardiner, G. (2021) "Relevance and risk: How the relevant alternatives framework models the epistemology of risk" *Synthese* v199(1-2), pp481-511
- Gerwith, A. (1981) "Are there any absolute rights?" *Philosophical Quarterly* v31(122), pp1-16
- Hayenhjelm, M. and Wolff, J. (2012) "The moral problem of risk impositions: A survey" *European Journal of Philosophy* v20(S1), e26-e51
- Holm, S. (2016) "A right against risk imposition and the problem of paralysis" *Ethical Theory and Moral Practice* v19(4), pp917-930
- Horton, J. (2017) "Aggregation, complaints and risk" *Philosophy and Public Affairs* v45(1), pp54-81
- Kumar, R. (2015) "Risking and wronging" *Philosophy and Public Affairs* v43(1), pp27-51
- McCarthy, D. (1997) "Rights, explanation and risks" *Ethics* v107(2), pp205-225
- McKerlie, D. (1986) "Rights and risk" *Canadian Journal of Philosophy* v16(2), pp239-252
- Nelkin, D. and Rickless, S. (2014) "Three cheers for double effect" *Philosophy and Phenomenological Research* v89(1), pp125-158
- Nozick, R. (1974) *Anarchy, State and Utopia* (New York: Basic Books)
- Oberdiek, J. (2004) "Lost in moral space: On the infringing/violating distinction and its place in the theory of rights" *Law and Philosophy* v23(4), pp325-346

- Oberdiek, J. (2012) "The moral significance of risking" *Legal Theory* v18(3), pp339-356
- Pritchard, D. (2015) "Risk" *Metaphilosophy* v46(3), pp436–461
- Railton, P. (1985) "Locke, stock and peril: Natural property rights, pollution and risk" reprinted in *Facts, Values and Norms* (Cambridge: Cambridge University Press, 2009)
- Rowe, T. (2021) "Can a risk of harm itself be a harm?" *Analysis* v81(4), pp694-701
- Scanlon, T. (2008) *Moral Dimensions: Permissibility, Meaning, Blame* (Cambridge, MA: Harvard University Press)
- Smith, M. (2016) *Between Probability and Certainty: What Justifies Belief* (Oxford: Oxford University Press)
- Smith, M. (2018) "The logic of epistemic justification" *Synthese* v195(9), pp3857-3875
- Smith, M. (2022) "The hardest paradox for closure" *Erkenntnis* v87(4), pp2003-2028
- Smith, M. (forthcoming) "Decision theory and de minimis risk" *Erkenntnis*
- Song, F. (2019) "Rights against high-level risk impositions" *Ethical Theory and Moral Practice* v22(3), pp763-778
- Thomson, J. (1977) "Some ruminations on rights" *University of Arizona Law Review* v19(1), pp45-60
- Thomson, J. (1983a) "Some questions about government regulation of behavior" in Machan, T. and Johnson, M. eds. *Rights and Regulation: Ethical, Political and Economic Issues* (Cambridge, MA: Ballinger)
- Thomson, J. (1983b) "Imposing risks" in Gibson, M. ed. *To Breathe Freely: Risk Consent and Air* (Totowa, NJ: Rowman and Allanheld)
- Thomson, J. (1990) *The Realm of Rights* (Cambridge, MA: Harvard University Press)
- Thomson, J. (1999) "Physician assisted suicide: Two moral arguments" *Ethics* v109(3), pp497-518