

# Increasing the Risk That Someone Will Die Without Increasing the Risk That You Will Kill Them<sup>1</sup>

Thomas Byrne

UCLA / MIT

**Pre-Print. Forthcoming at *Philosophy and Phenomenological Research*.**

1. When Ben dies, it might be the case that

Annie increased the risk that Ben would die,

and it might be the case that

Annie increased the risk that she (Annie) would kill Ben.

Often those increases will go hand-in-hand—for example, when Annie fires a gun in Ben’s direction—but they needn’t. Consider:

ROULETTE: Gunslinger is determined to play a single round of Russian roulette with Ben. Gunslinger has four bullets and a revolver with eight chambers. Before Gunslinger proceeds, Annie swaps Gunslinger’s revolver for one with only six chambers. Gunslinger unwittingly loads the four bullets into the six-chambered revolver, vigorously spins the cylinder and, when it stops spinning, he aims at Ben and pulls the trigger: Ben is shot dead.

By swapping the revolvers, Annie increased the risk that Ben would die—from four-in-eight to four-in-six—even though Annie didn’t increase the risk that *she* would kill Ben (there was never any risk of that).

Similarly, when Ben’s house floods, it might be the case that Annie increased the risk that Ben’s house would flood and it might be the case that Annie increased the risk that she (Annie) would flood Ben’s house. These can also come apart:

EMISSIONS: Last year, Annie flew her private jet across the Atlantic, emitting 10 tonnes of carbon-dioxide (CO<sub>2</sub>). The more CO<sub>2</sub> in the atmosphere, the more of the

---

<sup>1</sup> I’m grateful to Allison Balin, Kevin Dorst, Caspar Hare, Agustín Rayo, Haley Schilling, Kieran Setiya, Brad Skow, Jack Spencer, Roger White, Jake Wojtowicz, members of MIT’s Matti group, and an anonymous reviewer for this journal.

sun's energy that is trapped in the atmosphere; the more of the sun's energy that is trapped in the atmosphere, the greater the risk of an extreme weather event occurring at any given time or place (floods, fires, droughts, tornados, etc.). This week, a freak storm flooded Ben's house.

By emitting that CO<sub>2</sub>, Annie increased the risk that Ben's house would flood, even though Annie didn't increase the risk that *she* would flood Ben's house (there was never any risk of that).

Here is a third example. When Ben's vase breaks, it might be the case that Annie increased the risk that Ben's vase would break and it might be the case that Annie increased the risk that she (Annie) would break Ben's vase. These can come apart, too:

VASE: Ben keeps his vase on a side table in the corridor. While attending a chaotic party at his house, Annie carefully moves the vase to the edge of the table. An hour later, someone bumps into the table, knocking the vase to the ground.

By moving the vase to the edge of the table, Annie increased the risk that Ben's vase would break, even though Annie didn't increase the risk that *she* would break the vase. (There was never any risk of that—she was *very* careful.)

There's a pattern here. For agent *A*, thing *X* and action *V*,  
when *X* *V*s (e.g. the vase breaks),

it might be the case that

(i) *A* increased the risk that *X* would *V* (e.g. that the vase would break),

and it might be the case that

(ii) *A* increased the risk that she would *V* *X* (e.g. that she would break the vase).

When (ii) is true, but (i) is not, then I shall say that

(iii) *A* *merely* increased the risk that *X* would *V*

and I will call any action that merely increases the risk that X will V a *mere risk imposition*.<sup>2 3</sup> (What about cases in which X doesn't V? I discuss those in §3.)

Mere risk impositions have converses. Just as A might increase the risk that X will V without increasing the risk that A will V X, so too might A decrease the risk that X will V without decreasing the risk that A will V X. When they do, I shall say that

(iv) A *merely* decreased the risk that X would V.

For example:

Gunslinger\* is determined to play Russian roulette with Ben\*. Gunslinger\* has four bullets and a revolver with eight chambers. Before Gunslinger\* proceeds, Annie\* steals one of Gunslinger\*'s bullets. Gunslinger\* loads the remaining three bullets into his eight-chambered revolver, vigorously spins the cylinder and, when it stops spinning, he aims at Ben\* and pulls the trigger: Ben\* is shot dead.

By stealing the bullet, Annie\* decreased the risk that Ben\* would die—from four-in-eight to four-in-six—even though Annie\* didn't decrease the risk that *she* would kill Ben\* (there was never any risk of that). That is: Annie\* merely decreased the risk that Ben\* would die.

What happens when a mere risk imposition meets its converse? For example:

ROULETTE OFFSET: Gunslinger is determined to play Russian roulette with Ben. Gunslinger has four bullets and a revolver with eight chambers. Annie swaps his eight-chambered revolver for one with six chambers. She then takes one of Gunslinger's bullets. Gunslinger loads the three bullets into the six-chambered revolver, vigorously spins the cylinder and, when it stops spinning, he aims at Ben and pulls the trigger: Ben is shot dead.

By swapping the revolvers, Annie merely increased the risk that Ben would die—from four-in-eight to four-in-six. By stealing the bullet, Annie merely decreased the risk that Ben would die—four-in-six to

---

<sup>2</sup> Strictly, one should distinguish between the transitive form of V (e.g. Annie *breaks<sub>T</sub>* the vase) and the intransitive form (e.g. the vase *breaks<sub>I</sub>*). I ignore this since it would smudge the prose to include it, yet nothing turns on it.

I introduced this way of speaking in my (2021). There I said that each instance of A V-ing X (A killing B, A flooding B's house, etc.) is an instance of the same metaphysical relation—one that I dubbed *MAKING* (just as, for example, each instance of A causing X to V is an instance of the metaphysical relation, *viz.* causation). And I argued that *MAKING* aligns with the central deontological constraint: e.g. it's impermissible to *MAKE* someone die (/kill them) in order save the lives of five others. Readers familiar with that paper might think of this paper as continuing that same project—this time focusing on constraints in the absence of *MAKING*.

<sup>3</sup> Barry and Cullity (2022) mention this distinction and dub it the difference between “attributable” and “non-attributable” harms. To my mind, that language suggests that it's an epistemic distinction—concerning what we can attribute to whom—which it isn't. It also puts *harming* centre-stage, but there is widespread disagreement about what is and isn't a harm (a disagreement I think best avoided when possible).

three-in-six. Since the degree by which Annie increased the risk (one-in-six) is no greater than the degree by which Annie decreased the risk (also one-in-six), I will say that Annie has *offset* her mere risk imposition.

How should we think about such cases? I will say that mere risk impositions come apart from other risk impositions in morally important ways. I'll say that offsetting a mere risk imposition has important moral implications, while offsetting other risk impositions does not. That conclusion has implications for the morality of carbon offsetting and they will be addressed in §6. It also has more general applications and those will be discussed in §8.

## 2. Let's first return to the simpler case we began with:

ROULETTE: Gunslinger is determined to play a single round of Russian roulette with Ben. Gunslinger has four bullets and a revolver with eight chambers. Before Gunslinger proceeds, Annie swaps Gunslinger's revolver for one with only six chambers. Gunslinger unwittingly loads the four bullets into the six-chambered revolver, vigorously spins the cylinder and, when it stops spinning, he aims at Ben and pulls the trigger: Ben is shot dead.

By swapping the revolvers, Annie merely increased the risk that Ben would die. What is the moral status of her doing so? I take it to be a datum that, in swapping the revolvers, Annie wronged Ben—she did a *terrible* thing to him.

That Annie wronged Ben also follows from a plausible principle connecting the moral status of mere risk impositions to the moral status of something more familiar. The principle I have in mind is that:

other things equal, if A's V-ing X would wrong B, then A's merely increasing the risk that X would V wrongs B.

Applied to ROULETTE, that principle says that if Annie's killing Ben wrongs Ben, then Annie's merely increasing the risk that Ben would die wrongs Ben; and since Annie's killing Ben would certainly wrong Ben, that principle returns that Annie's merely increasing the risk that Ben would die wrongs him, too. It's an attractive principle and I suspect it's true ("other things equal" is a *very* forgiving

clause), but I won't argue for it since the particular claim that Annie wronged Ben in ROULETTE is sufficient for my purposes here and, as I said, I take that to be a datum.

(As an aside, notice that while the facts of the case fix that Annie wrongs Ben, they don't fix whether or not she acted permissibly. After all, the case doesn't tell us why Annie swapped the revolvers: perhaps a villain had kidnapped her children and insisted that she trade them for an eight-chambered revolver, lest he kill them. Even if that were why she did it, it wouldn't change the fact that she wrongs Ben—she still did a terrible thing to *him*. In any case, that Annie wrongs Ben will be my focus, here.)

By swapping the revolvers in ROULETTE, Annie wronged Ben. In virtue of what is that so? The simple answer is that Annie wronged Ben in virtue of increasing the risk that he would die. I think the simple answer is correct, but it's worth ruling out some other candidates.

(a) *That Ben wouldn't have died if Annie hadn't swapped the revolvers.* But it isn't true that Ben wouldn't have died if Annie hadn't swapped the revolvers. Suppose that Annie hadn't, in fact, swapped the revolvers: what would have happened? Well, Gunslinger would instead have loaded those four bullets into his eight-chambered revolver and he would have vigorously spun that gun's cylinder, before aiming at Ben and pulling the trigger—that much we know. But whether Ben would have been killed depends on whether there would have been a bullet under the firing pin; and that depends on exactly how Gunslinger would have spun that revolver:

if he had spun it with 8.54N of force, then (we can suppose) a bullet would have been under the firing pin,

if he had spun it with 8.55N of force, then a bullet wouldn't have been under the firing pin,

if he had spun it with 8.52N of force, then a bullet wouldn't have been under the firing pin,

and so on.

But there just isn't any fact of the matter of exactly how Gunslinger would have spun the eight-chambered revolver if Annie hadn't swapped the revolvers. In the language of possible worlds (and following the standard approach to evaluating conditionals such as these), the possible world in which Annie doesn't swap the revolvers and Gunslinger spins his eight-chambered revolver with 8.54N of

force is no closer to the actual world than the possible world in which Annie doesn't swap the revolvers and Gunslinger spins with 8.55N of force. And so there is no fact of the matter as to whether Ben would have been killed had Annie not swapped the revolvers and, *a fortiori*, it's not true that Ben wouldn't have died had she not done so.<sup>4</sup>

What goes for ROULETTE here also goes for EMISSIONS: there is no fact of the matter as to whether Ben's house would have flooded this week had Annie not flown the Atlantic. The weather is extremely sensitive to past conditions, since tiny changes in those conditions rapidly multiply. Indeed, they multiply so rapidly that while forecasters can pretty accurately predict tomorrow's weather, their forecasts for more than ten days hence are no better than guesses. Given this sensitivity, whether Ben's house would have flooded this week had Annie not flown the Atlantic depends on *exactly* what Annie would have done instead, yet there is no fact of the matter. For instance, it might be that if Annie hadn't taken her trip, then she would have spent the weekend gardening instead, but the weather in the vicinity of Ben's house this week is dependent upon exactly how she would have gardened instead:

if Annie had gardened in exactly *this* way (buffeting *these* molecules with her trowel at *this* time and in *this* way etc.), then (we can suppose) there wouldn't have been a storm near Ben's house this week,

if Annie had gardened in exactly *that* way, then there would have been a storm near Ben's house this week,

if Annie had gardened in exactly *that other* way, then there wouldn't have been a storm near Ben's house this week,

and so on.

Yet just as there is no fact as to exactly how Gunslinger would have spun the barrel had Annie not swapped the revolvers, nor is there a fact as to exactly how Annie would have gardened had she not flown the Atlantic; *a fortiori* it's not true that Ben's house wouldn't have been flooded had Annie not taken her trip.

What goes for ROULETTE and EMISSIONS also goes for VASE: there is no fact of the matter as to whether the vase would have broken if Annie hadn't moved it. Whether it would have broken depends

---

<sup>4</sup> Much of this section is heavily indebted to Hare (2011).

on whether someone would have walked down the corridor in just the right way: sufficiently off-course, sufficiently forcefully, with sufficient carelessness (drunkenness?), etc. Yet movements that precise through a chaotic party are similarly sensitive to past conditions since tiny changes in those conditions rapidly multiply. Given this sensitivity, whether the vase would have broken had Annie not moved it depends on exactly what Annie would have done instead yet—just as before—this is no fact of the matter. This time, I leave it to the reader to imagine their own counterfactuals.

(b) *That Annie caused Ben to die.* It's not at all clear whether, by swapping the revolvers, Annie caused Ben to die. Sometimes common-sense is a good guide to whether *c* causes *e*, but I'm sceptical that common-sense has anything much to say about this sort of case. So to convince ourselves one-way-or-other, we'd have to instead see what the various competing accounts of causation said about the case, before deciding which of those accounts to believe—no small task.<sup>5</sup> But thankfully it's not one we have to complete here. Since even if we did convince ourselves that Annie caused Ben to die, there's no doubt that mere causation is insufficient for Annie to wrong Ben: A sends B to the shops and, en route, B accidentally hits and kills C; A caused C to die (by causing B to kill her), yet clearly A does not wrong C.<sup>6</sup>

(c) *That Annie affected whether Ben would die (where an action affects whether Ben would die just in case it's indeterminate whether Ben would have died had that action not been performed).* Annie did affect whether Ben would die, but so too did, say, the bystander who caught Gunslinger's eye just before he spun the barrel (since there is no fact of the matter as to how Gunslinger would have spun the barrel had he not caught the bystander's eye); but that bystander certainly doesn't wrong Ben.

I have run out of candidates and so I return to the simple answer I began with: Annie wronged Ben in ROULETTE in virtue of having increased the risk that he would die.

---

<sup>5</sup> We'd have to decide because they don't all agree. For example, Lewis (1986) would say that Annie doesn't cause Ben to die, while Frick (ms) would say that she does.

<sup>6</sup> A reviewer asks whether I was too hasty here and suggests an alternative diagnosis: Annie wronged Ben in virtue of *intentionally* causing him to die (which is left untouched by A's not wronging C, since A didn't intentionally cause C to die). If that's right, then we would have to settle whether Annie did in fact cause Ben to die, after all. However, I don't think it is right since wronging and intentions are independent: if I kill you with my car then I wrong you (I do a *terrible* thing to you), regardless or whether it was intentional.

3. When I introduced mere risk impositions that X would V, I did so with three examples and, in each of those examples, X did indeed V: Ben died in ROULETTE, Ben's house flooded in EMISSIONS and Ben's vase broke in VASE (a torrid week for Ben!). I did it that way because this paper is about wrongings and while I take it to be a datum that Annie wrongs Ben in ROULETTE, it's unclear to me whether Annie would wrong Ben in a variation of ROULETTE in which Ben doesn't die—and the same goes for those variations of EMISSIONS and VASE in which Ben's house and vase are unaffected.

Some people think that Annie would wrong Ben in those cases: they think that so-called “pure” risk impositions do wrong (call them the “affirmers”). Others disagree: they think pure risk impositions cannot wrong and, instead, all that could be said is that Annie would have risked wronging Ben (call them the “deniers”).<sup>7</sup> It's unclear to me who's right because it's unclear to me what the cash value of their disagreement is: what is the difference between Annie's pure risk imposition *wronging* Ben and Annie's pure risk imposition *risking wronging* him? Sometimes, the mark of A's wronging B is that A owes B compensation, but to employ that mark here would be to beg the question against the affirmers since compensation plainly isn't owed in cases of pure risk imposition (what could it be compensation for?). So, again, what is the cash value of the affirmers' and the deniers' disagreement? As I say, it's unclear to me.

In any case, neither side is excluded here. I will only consider cases in which X does V and the deniers can proceed as if those are the only cases that matter, while the affirmers can proceed by extending everything said about those cases to those case of pure risk.

#### 4. Recall:

ROULETTE OFFSET: Gunslinger is determined to play Russian roulette with Ben. Gunslinger has four bullets and a revolver with eight chambers. Annie swaps his eight-chambered revolver for one with six chambers. She then takes one of Gunslinger's bullets. Gunslinger loads the three bullets into the six-chambered revolver, vigorously spins the cylinder and, when it stops spinning, he aims at Ben and pulls the trigger: Ben is shot dead.

---

<sup>7</sup> For discussion of pure risk see, e.g., Thomson (1986) and Oberdiek (2012).



Did Annie wrong Ben? (I again set aside the question of whether Annie acted permissibly since the details of the case do not settle it: if it was costless for her to steal more bullets, then surely it was impermissible for her not to do so; on the other hand, if she needed the eight-chambered revolver to save the life of another and it wasn't possible for her to steal more bullets, then I would think she acted permissibly.)

The simple answer is that Annie didn't wrong Ben. And here's the simple explanation why: the risk that Ben would die (three-in-six; one-in-two) was no higher than it would have been had Annie not gotten involved (four-in-eight; one-in-two)—and no one else intervened, and there weren't any confounding factors, etc. Again, I think the simple answer and its explanation are both correct, but again it's worth considering other candidates.

*(a) through (c), from above:* the same considerations given above *vis-a-vis* Annie wronging Ben in ROULETTE apply equally here *vis-a-vis* Annie wronging Ben in ROULETTE OFFSET.

*(d) That Annie didn't do more to help (that she didn't, e.g., steal all Gunslinger's bullets).* Perhaps Annie could have done more to help, but the same holds for, say, a bystander who similarly failed to steal all of Gunslinger's bullets; but that bystander certainly doesn't wrong Ben (failures to aid are not wrongings). (One possible difference between Annie and the bystander *vis-a-vis* (d) is that, since Annie has already stolen one bullet, it would be costless for her to take another, while it might be costly for the bystander to do so. *Possible* but not necessary: we can imagine that the rest of the bullets are in Gunslinger's pocket.)

*(e) That Annie performed some action that merely increased the risk that Ben would die.* This candidate differs from (a) in that it seeks to ground the fact that Annie wronged Ben entirely in her swapping the revolvers (an *action* that increased the risk that Ben would die), setting aside whatever she might have done later—in this case, setting aside that she later stole one of Gunslinger's bullets. This won't do. Imagine a case in which Annie first swaps the revolvers and then steals, say, two of his bullets (or three...or all of them!): in such a case, the risk that Ben would die is lower (potentially much lower) than it would have been had Annie not gotten involved. Plainly, Annie would not wrong

Ben in such a case and so the fact that Annie performed some action that merely increased the risk that Ben would die cannot be sufficient for Annie to wrong Ben.

I have again run out of candidates and so I again return to the simple answer I began with: Annie didn't wrong Ben in ROULETTE OFFSET. Why not? Because the risk that Ben would die was no higher than it would have been had Annie not gotten involved: Annie offset her mere risk imposition.

This is not a trivial result. Notably, the same does not hold for risk impositions, in general. Consider:

Annie fires a bullet into the air. She then prevents someone else from firing a bullet into the air. By firing her bullet into the air, Annie increased the risk that Ben (walking nearby, without a helmet) would die; by preventing the other bullet being fired into the air, Annie decreased the risk that Ben would be killed by the same degree. Alas, the bullet falls on Ben's head, killing him.

Even though Annie decreased the risk that Ben would die by the same degree that she increased the risk that he would die—even though, that is, Annie offset the risk that she imposed that Ben would die—Annie clearly wronged Ben. After all, Annie killed Ben! (Note how the same holds even if Annie stops someone else from firing two bullets into the air—or three, or four. In which case, Annie would have decreased the risk that Ben would die, but she still wronged Ben: she killed him!)

Similarly, consider a case from a recent paper by Christian Barry and Garret Cullity (2022):

A boating lake provides the water for a nearby town. The boats release a toxic discharge into the lake. Annie drives a boat, releasing a certain amount of toxin into the lake. Later, she uses a filter to remove the same amount of toxin from the lake that her boating released into it. Ben's water supply comes from the lake and he consumes some of the toxin released by Annie's boat.

Even though Annie offset the risk she imposed that Ben would be poisoned, Annie clearly wronged Ben—she poisons him!<sup>8</sup> (Barry and Cullity themselves don't address whether Annie wrongs Ben. They only say that she acts permissibly. That might be true, but only under certain fillings-in of the details: e.g. it presumably wouldn't be permissible to discharge vast amounts of toxin into a lake into order to go pleasure boating, even if you removed other toxins from the lake.)

---

<sup>8</sup> We shouldn't be distracted by the epistemic problem of determining whether Annie poisoned Ben and, in turn, whether Annie wronged Ben: A wronged B when she sneezed on him, infecting him with black death, even though no one at the time had any way of knowing that she had done so.

This all suggests a thesis:

(\*) offsetting a mere risk imposition prevents that mere risk imposition from wronging, while offsetting risk impositions in general does not.

It's an attractive thesis and it accounts for why Annie didn't wrong Ben in ROULETTE OFFSET, but did wrong Ben in the two cases, above. (In what sense does offsetting "prevent" the mere risk imposition from wronging? In the same sense that paying for an item before removing it from the shop "prevents" that removal from wronging the owner.)

We should wonder why mere risk impositions come apart from risk impositions in general in the way (\*) says. My best guess is that it's because mere risk is *fungible*: its units are interchangeable in every way. Electronic money is also fungible and it behaves similarly. Suppose, for example, that Philanthropist first transfers \$1000 into PETA's account and, later, hacks into PETA's account and transfers \$1000 elsewhere. Only a confused fur-trader could complain of Philanthropist that she enriched PETA. That's because electronic money is fungible and so it's nonsense to say of any particular unit of money in PETA's account that it came from some donor or other. Instead, the only facts in the vicinity are quantitative ones: that the balance is higher (/lower) than it would have been were it not for a certain credit (/debit); that, e.g., the balance is \$20 higher than it would have been but for so-and-so's donation. Yet, given that Philanthropist performed both transfers and, as a result, the balance is the same as it would have been had Philanthropist performed neither transfer, there is no such quantitative fact. *A fortiori* there is nothing for the fur-trader to complain about.

My guess is that the very same goes for mere risk. Mere risk is fungible and it's similarly nonsensical to say of, e.g., ROULETTE that any given unit of risk that Ben would die is "Annie's risk" (what could that even mean?). Instead, all we can say in that particular case is that the risk that Ben would be killed is higher than it would have been had Annie not acted as she did.<sup>9</sup> However, in

<sup>9</sup> Of course, there are other things we can say, but what matters is whether they make a moral difference. For example, we can say that the risk that Ben would be killed *by Annie's revolver* is higher than it would have been, but does that make a moral difference? I don't think so. What matters, I would think, is whether Ben is killed (fixing how painful a death it is). More pressingly, if it morally matters whether he is killed by Annie's revolver (as opposed to Gunslinger's), then it should also matter whether he's killed at 12:00 (or at 12:01), or in *this* exact spot (or *that* exact spot), or that the bullet penetrates exactly *here* (or *there*) and so on. But those latter three aspects of Ben's death are easily affected: recall the bystander who catches Gunslinger's eye before he pulls the trigger—he affects the time of death, Ben's exact location at the time of death and, in turn, the location of the bullet's penetration. So if Annie wrongs Ben in virtue of him being killed by her revolver, then it seems the bystander also wrongs Ben in virtue of the time, place and bullet-location of his death. But, as before, the bystander certainly doesn't wrong Ben. Alternatively, one must explain why the revolver used implicates Annie, but these other aspects don't implicate the bystander. I don't see how that explanation could go.

ROULETTE OFFSET we can't even say that: given that Annie both swapped the revolvers and stole the bullet and, as a result, the risk that Ben would be killed was the same as it would have been had Annie performed neither of those actions, there is no such quantitative fact. On the other hand, when Annie, e.g., fired the bullet into the air, there is something else we can say that *isn't* fungible—namely, that Annie killed Ben.

5. Just as we wondered why mere risk impositions come apart from risk impositions in general, we should also wonder when they come apart. I've set the matter aside until now, because this paper's working cases haven't raised the issue: it's plain in ROULETTE that Annie increased the risk that Ben would die, just as it's plain that she didn't increase the risk that she would kill him: similarly for EMISSIONS, and VASE (and their offset counterparts).

What goes for those three also goes for the majority of other cases and that's because we are instinctively expert at determining whether something is or isn't a killing (or a flooding or a breaking, etc.). Indeed, this expertise allows us to draw distinctions *vis-a-vis* killing (etc.) even when common-sense causation sees no difference. For example, if A beats B and leaves him immobile in a field, before he later dies of exposure when the temperature drops, does A kill B? *Yes*. If A beats B and leaves him immobile in a field, before he is later struck by lightning when a storm rolls in, does A kill B? *No*. (Even though, in both cases, it seems like A causes B's death by causing the weather to kill him.) I've discussed this expertise at length, elsewhere (Byrne, 2021 §3).

That instinctive expertise gets us a long way, but it does eventually fail (and I'll turn to examples of its doing so, presently). One response to such failure is the search for an analysis of killing (and of flooding and of breaking, etc.—or perhaps a single analysis covering them all) which will tell us whether A V'd X (or risked V'ing X) for any particular case. I have little hope for such analyses—both here and more generally (*ibid.* §5). Instead, I think the best we can do is take the hard cases as they come and, slowly, build up some theory—no small task. I close this section doing some of that: starting with a case where it's very unclear whether A increases the risk that she will V X, before turning to a simpler case which might help.

To that end, suppose:

HELMET: Annie steals Ben's helmet. With no reasonable alternative, Ben decides to cycle home, regardless. On the way, he is hit by a car and dies of a head injury. Had Ben been wearing a helmet when he was hit, he wouldn't have died.

By stealing Ben's helmet, Annie increased the risk that Ben would die, but did she increase the risk that she would kill him? Did she kill him? To my mind, this is a *very* difficult question; moreover, I don't think it's one that can be answered by staring deeper into the case itself—so let's consider a different one, instead.

Gunslinger\* is practicing. He loads his revolver with blanks, vigorously spins the cylinder and, when it stops spinning, he aims at Ben\* and pulls the trigger. Unbeknownst to Gunslinger\*, Annie\* had replaced one of his blanks with a live round and, unfortunately, that round was in front of the firing pin when the trigger was pulled. Ben\* dies.

It seems to me that Annie\* not only increased the risk that Ben\* would die, but that she also increased the risk that she would kill him—and she did kill him. This is not a mere risk imposition. (Similarly, consider the case in which the stagehand loads the prop gun with live rounds; or the case in which he swaps the dummy razor for a real one just before the actor is to “slice” his throat. The stagehand kills both times.)

What separates the preceding case from ROULETTE, above, is that here Annie\* *created* the risk that Ben\* would die—she didn't merely increase it. After all, before she intervened, the risk that Ben\* would die was zero (on the other hand, in ROULETTE there was already some risk that Ben would die, before Annie swapped the revolvers). My suspicion is that this makes all the difference as to whether Annie kills Ben.

Although, that's not quite right since there was always some risk that Ben\* would die: heart attacks and freak accidents do happen. (A natural idea is that those risks are sufficiently negligible that they can be ignored, but that won't help here. Notice that I didn't specify the size of Gunslinger\*'s revolver and that's because it doesn't matter: there might have been five, fifty or fifty-thousand blanks in the revolver alongside Annie\*'s live round, but if that live round kills Ben\*, then Annie\* kills him—regardless of how negligible a risk it might have been.) What's important, then, it seems to me, is that

Annie\* created the risk that Ben\* would die *by a particular, metaphysically relevant, means* and he did die by that means: that is why Annie\* killed Ben\*.

Unfortunately, it's very hard to say what that particular means is—even though we have no problem determining what it isn't. It isn't, for example, that Ben\* be killed by that revolver (Annie\* would surely still have killed Ben\* if, at the last moment, Gunslinger\* had moved the bullets into a different revolver). Nor can it be that Ben\* be killed by that particular live round (it would have changed things if Gunslinger\* had instead thrown all the bullets at Ben\*, who then happened to choke on the live round—just as he would a blank; the live round would have killed him, but Annie wouldn't have). Or perhaps it's that Ben\* be killed by Gunslinger\* (although whether Annie created that risk turns on whether Gunslinger\* might have snapped and killed Ben\* otherwise—which doesn't seem relevant). Perhaps all we can say is that it's that Ben\* be killed by that live round in a Russian-roulette-sort-of-way.

That isn't fully satisfying, but it gives us something with which to return to HELMET. We know that Annie, by stealing Ben's helmet, increased the risk that he would die—but die by what particular means? If we can answer that question, we can then ask whether Annie created the risk that he would die by that particular means (versus merely increasing it): if she did, then the thoughts of the preceding paragraphs would return that she killed him; if not, then they would return that she didn't. The natural candidate is that Ben would die by being hit by a car and if that's right then Annie didn't create that risk (there was already some risk that he would die—helmets are not impervious), and therefore Annie didn't kill Ben. (This thesis makes a prediction: if helmets were impervious such that their wearers were immune to car crashes, then Annie would have killed Ben here. That's borne out since Annie would, I think, have killed Ben in such a case.)

So perhaps Annie doesn't kill Ben in HELMET. If that's right, and if the reasoning that led us there is right, then it will similarly lead us towards categorising other hard cases, too: the case in which A steals B's anti-venom, before B is bitten by a viper; or in which A delays B's departure such that B has to drive during a storm and crashes; or perhaps even the case in which A persuades B to play Russian roulette with C; and so on.

But that's all tentative. And even if it's all correct, it leaves much unanswered: what, for instance, is the particular metaphysically relevant means by which A increases the risk that X will V in any given case; and what makes it metaphysically relevant? It also leaves open whether there are other sorts of hard cases that do not lend themselves to that same treatment (I suspect there are). But, as I said, these are hard questions and answering them is no small task—and it's not one that can be completed here.

6. Let's return to the thesis just introduced:

(\*) offsetting a mere risk imposition prevents that mere risk imposition from wronging, while offsetting risk impositions in general does not.

That thesis has implications; I discuss one here.

Our CO<sub>2</sub> emissions contribute to climate change and to the suffering of those affected by it: they contribute by trapping the sun's energy in the atmosphere and, in turn, by increasing the risk that those affected by it would be affected. John Broome (2012) argues that by contributing to climate change in that way, we wrong those affected by it. Suppose he's right.<sup>10</sup>

Broome then goes on to argue that we can avoid wronging those affected by climate change if we offset our emissions by, e.g., capturing the same amount of CO<sub>2</sub> from the atmosphere that we emitted into it in the first place: if we emitted 10 tonnes, then we offset those emissions if we capture 10 tonnes from the atmosphere—and so on.

Broome's view is not popular.<sup>11</sup> Here is a tidy way of putting what is found to be objectionable about it, from Caspar Hare (2013). Consider:

BARRELS OFFSET: Aggie has a factory on the west side of the river and waste-processing plant on the east side. It costs a lot of money for Aggie to transport the waste from the factory across the river. Aggie's factory throws a barrel of toxic waste into the river. Later, Aggie's waste-processing plant retrieves a barrel of

---

<sup>10</sup> It's not obvious he's right. Some people will say that while Annie might have contributed to that suffering, her contribution to the suffering of any one individual will be so small as to be imperceptible to that individual; and then they'll say if no individual can perceive Annie's contribution to their suffering, then Annie can't have wronged them. That's a complicated discussion. For my part, I suspect Broome is right: I suspect that it's the fact that Annie wrongs those people that explains why, other things equal, Annie shouldn't fly her private jet around the world.

<sup>11</sup> See for example Goodin (1994), Cripps (2016), Monbiot (2006).

waste from the river that an upstream farm had thrown into the river. The barrel from Aggie's factory goes downstream and ruins Bertie's crop.

By throwing one barrel into the river and then retrieving a second barrel, Aggie hasn't increased the risk that Bertie's crop would be destroyed. Even so, Hare (rightly) says that Aggie wrongs Bertie—she ruined his crop! Hare then claims that what goes for Aggie offsetting the waste she throws into the river also goes for our offsetting the CO<sub>2</sub> we emit into the atmosphere and, therefore, that Broome is wrong about carbon offsetting.

With (\*) in hand, we can see that the final step in Hare's argument is mistaken. Compare the preceding case with the following:

EMISSIONS OFFSET: Last year, Annie flew her private jet across the Atlantic, emitting 10 tonnes of CO<sub>2</sub>. When she landed, she captured 10 tonnes of CO<sub>2</sub> from the atmosphere. The more CO<sub>2</sub> in the atmosphere, the more of the sun's energy that is trapped in the atmosphere; the more of the sun's energy that is trapped in the atmosphere, the greater the risk of an extreme weather event occurring at any given time or place (floods, fires, droughts, tornados, etc.). This week, a freak storm flooded Ben's house.

Here, Annie merely increased the risk that Ben's house would flood. BARRELS OFFSET wasn't like that since, in throwing the barrel into the river, Aggie not only increased the risk that Bertie's crop would be ruined, but she additionally increased the risk that she (Aggie) would ruin Bertie's crop. Rightly understood as such, Broome's claim about carbon offsetting and Hare's claim about BARRELS OFFSET sit on opposing sides of the line drawn by (\*) and that line makes all the difference in the world *vis-a-vis* offsetting and wronging since offsetting a mere risk impositions does prevent wronging (as in EMISSIONS OFFSET), while offsetting risk impositions in general does not (as in BARRELS OFFSET).

7. When I introduced ROULETTE OFFSET in §1, I said that Annie offset her mere risk imposition since the degree by which Annie merely increased the risk (one-in-six) is no greater than the degree by which Annie merely decreased the risk (also one-in-six). There's actually more subtlety to offsetting than that introduction acknowledged—let's turn to it.

ROULETTE OFFSET has certain structural properties, some of which have already been made plain:



- (I) Annie merely increased the risk that Ben would die by some degree,
- (II) Annie merely decreased the risk that Ben would die by that same degree,

While (I) and (II) are necessary for Annie to have offset her mere risk imposition, they are not sufficient. Consider an extreme case:

DOUBLE ROULETTE: Gunslinger-1 is determined to play Russian roulette with Ben. He has four bullets and an eight-chambered revolver. Annie steals two of Gunslinger-1's bullets. Gunslinger-1 loads his remaining two bullets into his eight-chambered revolver, spins, aims and fires: nothing happens. Later that day, Gunslinger-2 is determined to play Russian roulette with Ben. He also has four bullets and an eight-chambered revolver. Annie swaps his eight-chambered revolver for one with only six chambers. Gunslinger-2 loads his four bullets into the six-chambered revolver, spins, aims and fires: Ben is shot dead.

(I) and (II) hold because in stealing the two bullets Annie merely decreased the risk that Ben would die from 48/64 to 16/64, while in swapping the revolvers Annie merely increased the risk that Ben would die back from 16/64 to 48/64. Even so, I take it to be a second datum that Annie wrongs Ben here. By swapping the revolvers and increasing the risk that he would die, Annie did a *terrible* thing to Ben—even though Annie had decreased the risk that he would die by the same degree earlier that day. (Again, I set aside the question of whether Annie acted permissibly since, again, the details of the case do not settle it.)

Since offsetting a mere risk imposition prevents that mere risk imposition from wronging, Annie cannot have offset her mere risk imposition here. (I might instead have reasoned as follows: since Annie has offset here mere risk imposition, it can't be that such offsetting is sufficient to prevent wronging—and I might then have adjusted (\*) accordingly. As I see it, this is a terminological matter and I think it simpler to focus on the nature of offsetting, than on the nature of (\*).)

What is it that explains why Annie offset her mere risk imposition in ROULETTE OFFSET, but didn't in DOUBLE ROULETTE? To answer that question, we should consider other structural properties of ROULETTE OFFSET and DOUBLE ROULETTE, but this time ones that they don't share. One such property is ordinal, in that in ROULETTE OFFSET

- (III) Annie increased the risk that Ben would die *before* Annie decreased the risk that Ben would die,

while, in DOUBLE ROULETTE, Annie increased the risk *after* she decreased the risk. Another property is temporal, in that in ROULETTE OFFSET

(IV) Annie increased and decreased the risk that Ben would die more-or-less simultaneously,

while in DOUBLE ROULETTE those actions were hours apart. Another property concerns those means introduced in §5, in that in ROULETTE OFFSET

(V) Annie increased the risk that Ben would die by being shot by Gunslinger and Annie decreased the risk that Ben would die by being shot by Gunslinger.

The underlined clauses pick out the means by which Annie increased and decreased the risk that Ben would die and what matters *vis-a-vis* (V) is that they are the same. On the other hand, in DOUBLE ROULETTE, Annie decreased the risk that Ben would die by being shot by Gunslinger-1, yet increased the risk that Ben would die by being shot by Gunslinger-2—and those are different means.

My suspicion is that the ordinal property (III) is irrelevant here (would ROULETTE OFFSET have been any different had Annie instead stolen the bullet a second before swapping the revolvers?). On the other hand, I think the temporal and means properties (IV) and (V) are relevant, but not because they themselves make the difference between ROULETTE OFFSET and DOUBLE ROULETTE, but because they are necessary for a sixth property that does.

The property I have in mind concerns what I will call the *manifestation of risk*; where the risk that, e.g., Gunslinger-1 will kill Ben manifests at the moment Gunslinger-1 pulls the trigger. Understood as such, in ROULETTE OFFSET

(VI) Annie increased and decreased the risk that Ben would die before any corresponding risk that Ben would die had manifested.

(I.e. Annie both swapped Gunslinger's revolvers and stole his bullet before a trigger was pulled.) On the other hand, in DOUBLE ROULETTE Annie decreased the risk that Ben would die after the risk that Gunslinger-1 would kill Ben (a risk that Annie increased) had manifested. (I.e. Annie swapped Gunslinger-2's revolver after Gunslinger-1 had pulled his trigger.)

I think that (VI) makes all the difference between ROULETTE OFFSET and DOUBLE ROULETTE *vis-a-vis* offsetting. I think that because when I consider further variations of the cases, (VI) rightly divides them between those in which Annie doesn't wrong Ben and those in which she does. For

example, (VI) rightly places EMISSIONS OFFSET alongside ROULETTE OFFSET since the best estimates say it takes years for emissions to affect the climate. I leave it to the reader to consider other cases.

However, it's less clear to me why (VI) makes all the difference, but my suspicion is that it concerns something much more general than offsetting—namely, when it's proper to evaluate an agent's distinct actions independently and when it isn't. Often it won't make any difference either way, but here is an example where it does:

A charity shop operates using an honesty box. Annie enters the shop and takes a liking to certain dress, priced \$10. Annie drops \$10 into the honesty box and exits the shop carrying the dress.

Here Annie performed two distinct actions, she

(A) dropped \$10 into the honesty box,

(B) removed the dress from the shop.

When we want to evaluate Annie's time in the charity shop, we are making a mistake if we do so by first evaluating (A) and then evaluating (B). We are making a mistake if we evaluate as follows: "in performing (A), Annie did something good (she donated \$10 to charity); in performing (B), Annie did something bad (she took a \$10 from charity); yet the good and the bad cancel each other out and, on the whole, Annie did something morally neutral." We are making a mistake because while it's true that Annie did something morally neutrally—namely, she simply bought a dress—it's false that Annie did anything good or anything bad. To avoid that mistake, we have to evaluate (A) and (B) together, as a single action—namely, the action that is buying a dress.

Compare that with the case in which, e.g., Annie visits the shop twice: on Monday she sees nothing she likes, yet drops \$10 into the honesty box, regardless; she returns on Friday, sees a \$10 dress she likes and removes it from the store. Here Annie did do something good (namely, donating \$10 to charity) and something bad (namely, stealing a \$10 dress from the shop) and we would be making a mistake if we combined those two together to conclude that Annie's actions that week were morally neutral. They weren't morally neutral: for starters, Annie stole a dress (from charity!).

In the former case, but not the latter, we are making a mistake when we evaluate Annie's actions independently of each other. Why?

My suspicion is that whatever answers that question will similarly answer why Annie didn't wrong Ben in ROULETTE OFFSET, but did wrong him in DOUBLE ROULETTE and it will do so by appealing to (VI). That is, it will say that since (VI) holds of ROULETTE OFFSET, we would be making a mistake to evaluate

Annie's swapping Gunslinger's revolver

independently of

Annie's stealing one of Gunslinger's bullets.

On the other hand, it will say that since (VI) doesn't hold of DOUBLE ROULETTE, we would be making a mistake if we *didn't* evaluate

Annie's stealing two of Gunslinger-1's bullets

independently of

Annie's swapping Gunslinger-2's revolver.

And when Annie's actions in DOUBLE ROULETTE are evaluated independently, it's clear that Annie wronged Ben since, in swapping Gunslinger-2's revolver, she increased the risk that Ben would die. That is the right result.

There's a further question: why does (VI) make that difference to whether Annie wronged Ben? Just as there is the further question: why does the time delay between Annie's dropping the money in the box and removing the dress from the shop make a difference to whether Annie stole the dress? They are interesting questions and I suspect they share an answer, but it's not an answer that's required here.

Instead, what's important is that

(\*) offsetting a mere risk imposition prevents that mere risk imposition from wronging, while offsetting risk impositions in general does not.

While I've focused on a few specific examples of mere risk impositions here, my suspicion is that they are not a rare thing. Indeed, my suspicion is that once we are live to them, we will see them cropping up in various places—particularly when an agent's action affects others only in virtue of its being part of some complex or collective system. And, in turn, my suspicion is that we'll see that (\*)'s applications are correspondingly varied.

8. As evidence for that variety, I end by speculatively considering one quite different sort of example —meat-eating.

Peter Singer (1975) might have convinced us that factory farming is impermissible, but it's not obvious how that conclusion is supposed to bear upon the permissibility of, e.g., me buying a chicken from the supermarket. After all, I am not factory farming and the chicken is long-dead by the time I buy it. This is the so-called "I don't make a difference" objection to ethical vegetarianism.

Shelly Kagan (2011) says it's because I should *expect* my buying the chicken to make a difference and he provides the following toy model (which I've simplified further):

Whenever the butcher sells its 20th chicken, it orders 20 new chickens from the slaughterhouse. And whenever the slaughterhouse slaughters 20 chickens, it orders 20 new chickens from the hatchery. So when Annie buys a chicken, she knows that 19/20 times, her purchase won't make any difference, but 1/20 times (those times when her chicken is the 20th purchased) it will make a big difference: it will result in the hatching, rearing and eventual slaughtering of 20 chickens. Accordingly, Annie rightly *expects* that her purchase will make the difference of exactly one chicken being hatched, reared and slaughtered.

Kagan says that it's that expectation that explains why it's impermissible for Annie to purchase the chicken. Let's suppose that he's more-or-less correct.

For Kagan, that chickens are slaughtered and hatched in batches of twenty (as opposed to of fifty or fifty-thousand) isn't important since the expectation will always sum out to exactly one chicken. However, once the number becomes large enough—and there's every reason to think that in today's meat industry, the number is enormous—something important changes in the model: there are no longer any determinate facts about how Annie's purchase affects things, but, instead, only risk facts.

It's hard to say exactly what the relevant number is, but its size isn't that important. What's important is that it's big enough such that the timeframe by which it operates—the approximate time it takes for the number to be reached—is not one of days, but instead one of weeks. Since it takes around 42 days to rear a chicken to slaughter-weight, we can assume that even the least forward-thinking factory farm operates on a timeframe of at least that length.

Now, suppose that Annie bought her chicken five weeks ago and Ben ate exactly one chicken this week: would Ben have eaten exactly one chicken this week had Annie not bought hers? It depends on the sort of meat-eater that Ben is. If he eats chicken exactly once per week (like my grandmother who would eat chicken every Wednesday and only on Wednesdays), then the answer is likely *yes*. But if Ben isn't so set in his ways and instead eats according to his whims and fancies and moods and where he happens to be and what happens to look good on the restaurant's menu and...so on (as most of us do), then there will be no fact of the matter as to whether Ben would have eaten exactly one chicken this week.

The shortest route to seeing why goes via the weather. We've already seen that the weather is extremely sensitive to past conditions and so the weather today depends on *exactly* how Annie behaved five weeks ago. Similarly, what the weather would have been like today (and everyday for around the previous three weeks) had Annie not bought that chicken depends on *exactly* how Annie would have instead behaved, yet there just is no fact of the matter. It might be the case that she would instead have eaten pasta, but there's no fact as to whether she would have eaten it exactly *this* way or exactly *that* way, etc., yet it is those preciser facts that fix the weather for the last three weeks.

Readers might ask themselves the following question: if the weather had been different for the last three weeks (raining at different times, raining more, raining less, sunnier, slightly warmer, etc.), would you have eaten for dinner last night what you did in fact eat for dinner last night? If you're anything like me (or Ben), then you have no idea what you would have eaten for dinner last night and such ignorance is appropriate: if you're anything like me (or Ben), then there is no fact of the matter.

Just as there's no fact of the matter as to what Ben would have eaten last night had Annie not bought that chicken five weeks ago, nor is there any fact of the matter as to what he would have eaten (for much of) this week. In turn, there is no fact of the matter as to how much chicken he would have eaten this week had Annie not bought that chicken (and the same goes for the previous two-or-so weeks, too). And what goes for Ben goes for everyone else like Ben: there's no fact of the matter as to how much chicken they would have eaten in the last three-or-so weeks had Annie not bought her chicken five weeks ago.

If that's all correct, then Kagan's explanation of why it's impermissible for Annie to buy a chicken is incorrect—in letter, not spirit. Recall, Kagan says that  $1/n$  times that Annie buys a chicken, that purchase will result in the hatching, rearing and eventual slaughtering of  $n$  chickens. But we now see that that's incorrect: it will never be the case that if Annie hadn't purchased a particular chicken, then  $n$  chickens that were hatched, reared and slaughtered wouldn't have been. Instead, all we can say is that a number of chickens were hatched, reared and slaughtered in the time following Annie's purchase and, by making that purchase, Annie increased the risk that those chickens would be hatched, reared and slaughtered.<sup>12</sup> I suggest that it is that mere risk imposition—one that wrongs those chickens—that explains why Annie's purchase was impermissible.

That all means that (\*) applies to Annie's purchase of the chicken: Annie can offset her purchase. That could happen in various ways, for example:

Annie sees Ben about to buy a chicken at the supermarket. Annie pays him \$10 not to buy a chicken and to instead have pasta that evening. Annie buys a chicken. (And, recall, Ben is the sort of person whose dietary preferences today are independent of whatever he ate yesterday: so his abstaining today doesn't increase the risk that he'll eat chicken tomorrow.)

In buying the chicken, Annie merely increased the risk that a greater number of chickens will be hatched, reared and slaughtered in the future. By paying Ben not to buy a chicken, Annie merely decreased that same risk—and by the same degree. Annie has offset her purchase. So if (if!) what accounts for why it's impermissible to buy dead chickens is that it increases the risk that more chickens will be hatched, reared and slaughtered, then Annie has acted permissibly.

As I said, that is speculative—far more so than the discussion that preceded it. However, its purpose is not to convince you that Annie's chicken purchase in that final case is permissible, but instead to support my suspicion that, once we are live to mere risk impositions, we will see them cropping up in various places throughout moral theory.

---

<sup>12</sup> The precise degree by which Annie's purchase increased that risk will depend on both the particular details of the size, timeframe and delay that the factory farm operates on. And pending how those details are fixed, the degree of increased risk might also depend upon when in the delay Annie purchased the chicken. I set the question aside since that degree isn't important for what follows (since it will always be offset).

## References

- Barry, C., & Cullity, G. (2022). Offsetting and Risk Imposition. *Ethics*, 132(2), 352–381. <https://doi.org/10.1086/716871>
- Broome, J. (2012). *Climate Matters: Ethics in a Warming World*. W. W. Norton.
- Byrne, T. (2021). MAKING Metaphysics. *Philosopher's Imprint*, 21(20). <http://hdl.handle.net/2027/spo.3521354.0021.020>
- Cripps, E. (2016). On Climate Matters: Offsetting, Population, and Justice. *Midwest Studies In Philosophy*, 40(1), 114–128. <https://doi.org/10.1111/misp.12050>
- Frick, J. (Submitted). *Chancy Causation and the Problem of Aggregate Events*.
- Goodin, R. E. (1994). Selling Environmental Indulgences. *Kyklos*, 47(4), 573–596. <https://doi.org/10.1111/j.1467-6435.1994.tb02067.x>
- Hare, C. (2011). Obligation and Regret When There is No Fact of the Matter About What Would Have Happened if You Had not Done What You Did. *Noûs*, 45(1), 190–206. <https://doi.org/10.1111/j.1468-0068.2010.00806.x>
- Hare, C. (2013). *Symposium on the work of John Broome*.
- Kagan, S. (2011). Do I Make a Difference? *Philosophy & Public Affairs*, 39(2), 105–141. <https://doi.org/10.1111/j.1088-4963.2011.01203.x>
- Lewis, D. (1986). Postscripts to 'Causation'. In D. Lewis (Ed.), *Philosophical Papers Vol. II*. Oxford University Press.
- Monbiot, G. (2006, October 18). Paying for our sins. *The Guardian*. <https://www.theguardian.com/environment/2006/oct/18/green.guardiansocietysupplement>
- Oberdiek, J. (2012). THE MORAL SIGNIFICANCE OF RISKING. *Legal Theory*, 18(3), 339–356. Cambridge Core. <https://doi.org/10.1017/S1352325212000018>
- Scanlon, T. (1998). *What We Owe to Each Other*. Belknap Press of Harvard University Press. <https://books.google.com/books?id=ZzvdtwEACAAJ>
- Singer, P. (1975). *Animal Liberation*. Random House.
- Thomson, J. J. (1986). Imposing Risks. In *Rights, restitution, and risk: Essays, in moral theory*. Harvard University Press.