

Resultant Luck*

I. Introduction

The focus of this paper is the concept of resultant (moral) luck. Resultant luck is one of the four varieties of moral luck originally identified by Nagel in his (1979).¹ It is the variety of moral luck that most plainly seems to illustrate the fact that commonsense morality is committed to the existence of moral luck and, also, for many, it is the most obviously problematic kind of luck. Resultant luck is usually characterized quite loosely, as moral luck (good or bad) “in the way our actions or projects turn out,” or as moral luck (good or bad) “about the consequences, or results, of our acts.” It is common to illustrate the concept of resultant luck by appeal to examples. For instance, we judge a reckless driver who kills someone as a result of his bad driving much more harshly than another reckless driver who is fortunate enough not to kill anyone simply because no one happens to run into his path. The “fortunate” reckless driver (the one who doesn’t come out responsible for any deaths) is morally lucky; the “unfortunate” reckless driver (the one who comes out responsible for a death) is morally unlucky. More generally, the point is made that what results from an agent’s acts is, quite commonly, beyond the agent’s control, or not fully within the agent’s control; still, we hold agents morally responsible for some of those results. Thus we hold agents

* Thanks to Juan Comesaña, Dan Hausman, Shaun Nichols, Derk Pereboom, Nico Sillins, Erin Taylor, an anonymous referee for *Philosophy and Phenomenological Research*, and audiences at the University of Arizona and Cornell University for helpful comments and discussion.

¹ See also Williams (1981). The term “resultant luck” is originally from Zimmerman (1987). Another common label is “consequential luck.” The other kinds of moral luck identified by Nagel are: luck about our constitution or dispositions, luck about the causes of our acts, and luck about the circumstances in which we act.

morally responsible for things that are not fully within their control. This is the standard, rough picture of resultant luck.²

The standard picture of resultant luck makes use of two important notions: the notion of being in control (or fully in control) of an outcome and the notion of being morally responsible for an outcome. In what follows I will assume that we have an intuitive grasp of these notions and I will leave them unanalyzed. That is to say, I will assume that we have an intuitive grasp of the sense in which, e.g., the reckless drivers are not (fully) in control of how things turn out, and I will assume that we also have an intuitive grasp of what it would be for agents to be morally responsible for outcomes in the world (not just for, say, their acts of will). In this paper I will aim to show that, even assuming that we can make sense of these important notions, the standard picture of resultant luck is too rough, simplistic, and it fails to capture the totality of the phenomenon of resultant luck. Thus my goal here is to shed light on the concept of resultant luck, by offering a characterization of that concept (the ordinary concept of resultant luck, or a somewhat cleaned up and pruned version of it) that captures the different sources of complexity behind it. In the process, I draw some distinctions that I think are helpful to get a better grip on the concept, I reveal some ambiguities permeating it, and I uncover some forms of resultant luck that have largely gone unnoticed. A secondary goal of the paper is to examine some interesting differences that exist between two particular varieties of resultant luck, resultant luck by action and by inaction.

Although my aim here is not to defend the possibility of resultant luck or to reject it, achieving a clearer understanding of the concept of resultant luck (getting a grip on the kind of thing to which we would be committed if we were to be committed to resultant luck) can of

² Note that resultant luck is a kind of moral luck that concerns specifically the concept of moral responsibility but not other moral concepts like the concept of moral obligation.

course be helpful, and perhaps even necessary, towards determining whether the phenomenon of resultant luck is really possible. The existing literature on moral luck doesn't come close to providing a full understanding of the phenomenon; hence this paper attempts to fill this gap.

II. Causal and non-causal resultant luck

Resultant luck is supposed to be moral luck about the results or consequences of acts. A preliminary question to consider is, what kinds of consequences of acts can give rise to resultant luck?

The cases that are typically offered as examples of resultant luck, such as the reckless driver example mentioned above, quite generally involve *causal* consequences of acts: the fortunate reckless driver is morally lucky because he doesn't cause any harm; the unfortunate reckless driver is morally unlucky because he causes harm. Now, on the face of it, just like there could be causal resultant luck (resultant luck about the causal consequences of acts), there could also be non-causal resultant luck (resultant luck about the non-causal consequences of acts).³ Here is one way in which one could argue for the existence of this type of resultant luck. Take the unfortunate reckless driver, who causes harm as a result of his bad driving—say, he kills a pedestrian. Imagine that the pedestrian he killed was married to a woman called “Carla.” Then, when the pedestrian dies, Carla becomes a widow. Moreover, it seems that Carla's widowing is also a consequence of the reckless driver's actions (she became a widow because of his bad driving), as well as something for which he is

³ Daniel Statman has noted this, in passing (see his “Introduction” to Statman (1993), n.37). He draws attention to an example by Rescher (1993): a villain plans to burglar his grandfather's house while he's on vacation; unbeknownst to the villain, however, the grandfather has just died and he has inherited everything in the house, so he ends up taking what is really his. Statman seems to think that this is an example of non-causal resultant luck, but he doesn't explain his reasons for believing this.

responsible.⁴ But, is Carla's widowhood a *causal* consequence of the reckless driver's actions? According to some views of causation and events, it isn't. For example, according to some such views (e.g., Lewis (1986a)), only predominantly *intrinsic* events can be causes and effects. On these views, Carla's widowhood isn't a causal consequence of the reckless driver's actions because it's too extrinsic: whether it obtains depends heavily on what happens in some other spatiotemporal region.

Imagine that we agree with those views and think that the driver's behavior doesn't cause Carla's widowhood. Then, in what sense does it result in it? Presumably, what one would have to say is that the driver's behavior results in the widowhood by virtue of a combination of both causal and non-causal connections. The driver's behavior causes the pedestrian's death, which, in turn, non-causally results in Carla's widowhood.⁵ What does the non-causal connection between the pedestrian's death and Carla's widowhood consist of? Following Kim (1974), one could say that it's a relation of "non-causal dependence": although the pedestrian's death doesn't cause Carla's widowhood, Carla's widowhood logically follows from the pedestrian's death, in conjunction with the fact that the pedestrian was married to Carla. (Note that this is an asymmetric relation: Carla became a widow because the pedestrian died, not vice-versa.)

At any rate, all I would like to suggest here is that we should be open to the possibility of non-causal resultant luck. I gave one possible example of this phenomenon. But, even if this example failed, there could be others. On the face of it, our acts can have non-causal

⁴ Why not think that Carla's widowhood is just identical with (numerically the same event as) the pedestrian's death? A reason for resisting this is that they seem to have different properties, in particular, whereas the death only involves the pedestrian, Carla's widowhood presumably involves Carla.

⁵ Note that the pedestrian's death and Carla's widowhood occur simultaneously. This is a reason to think that the relation between them is not causal, for instantaneous causation is not easy to come by (see Lewis (1986a), p. 263).

consequences in addition to causal consequences, and we can be just as responsible for them as we are for their causal consequences. Moreover, it seems that, just as we may not fully control what causally results from our acts, we also may not fully control what non-causally results from them. So we should allow for the possibility of this under-recognized variety of resultant luck. The distinction between causal and non-causal resultant luck will be useful in what follows; I will refer back to it on a couple of occasions.

III. An ambiguity about “how things turned out”

In this section I reveal an important ambiguity present in the way in which resultant luck is usually characterized, and I explain how I think that we should resolve that ambiguity.

As I said, resultant luck is typically characterized by appeal to examples. In particular, in order to illustrate the concept of resultant luck one usually concocts two possible scenarios that are the same with respect to what the agent can control but that differ with respect to the outcome that ensues. For instance, in the example above the two relevant scenarios seem to be one where the agent drives recklessly and harm ensues and another one where the agent drives in the same way but no harm ensues. A natural thought, then, is to think that what makes the reckless driver scenarios examples of resultant luck is the fact that in each case the agent’s act results in an outcome that could easily have been different, where the agent didn’t control which outcome actually came about. As we were imagining the unfortunate reckless driver scenario, the driver ran over a pedestrian who happened to cross the street exactly while he was driving by. Imagine that, had the pedestrian noticed the nice-looking lady across the street, he would have been distracted for a few seconds and would have started to cross the street a few seconds later than he actually did; as a result, he wouldn’t have died. In other

words, there are some nearby possible worlds where the driver behaves in the exact same way but the pedestrian doesn't die. It is tempting to think that this is what makes the unfortunate reckless driver scenario a case of resultant luck.

But this way of understanding resultant luck is misguided, and illustrating the phenomenon of resultant luck by appeal to scenarios of the kind described above is, at best, misleading. Imagine that, if the pedestrian had started to cross the street a few seconds later as a result of his brief distraction, then a truck whose driver just passed out would have run him over. In a scenario of that kind, the outcome that the agent's act resulted in couldn't easily have been different. In nearby possible worlds where the driver behaves in the same way but the pedestrian acts differently, the pedestrian still dies (and, we may imagine, in a similar way, at around the same time, etc.).⁶ But it still seems that the reckless driver is morally unlucky in this type of scenario: he is morally responsible for the pedestrian's death given that he ran over him, and despite the fact that he had no full control of the fact that he ran over him. Hence the fact that there can be resultant luck in scenarios of this kind shows that the characterization of resultant luck suggested above is wrong: what makes the reckless driver scenarios examples of resultant luck is not the fact that the agent's act results in an outcome that could easily have failed to occur, or the fact that the agent didn't control which outcome came about.

So, in what sense is resultant luck moral luck about "how things turn out"? The answer, I think, is that the expression "how things turn out" is ambiguous. The ambiguity in question is of the product/process kind. On the one hand, we can use "how things turn out" to

⁶ This is an example of "causal redundancy" (sometimes also loosely referred to as "overdetermination"). There is causal redundancy in the sense that there is an outcome that is actually caused in a certain way, but that would have still been caused in some other way if things had been slightly different (see Lewis (1986b), postscript E).

refer to the *outcome* that actually comes about (this is the way the expression is normally understood and, perhaps, the most natural way of understanding it). On this first interpretation, what fails to be fully in the agent's control is what actually happens after the agent acts. Let's call this the "outcome-driven" reading of "how things turn out." On the other hand, we can use the expression to refer to how the agent winds up being *related to* the outcome. On this second interpretation, what fails to be fully in the agent's control is (not just what actually happens but) whether the agent's behavior *results in* what actually happens. Given that this second interpretation zeroes in on the agent's relation to the outcome instead of on the outcome itself, I'll call it the "agent-driven" reading of "how things turn out."

Now, according to commonsense morality, what matters to an agent's moral responsibility is the agent's relationship to an outcome, not whether the outcome occurs after the agent acts.⁷ For instance, in the truck version of the reckless driver case, we would hold the driver responsible if he caused the pedestrian's death, even if the pedestrian would still have died in virtue of his being hit by the truck, and we would not hold him responsible if he didn't cause the pedestrian's death, even if the pedestrian still died in virtue of his being hit by the truck. It seems clear, then, that we should adopt the agent-driven view over the outcome-driven view.

So, how could we formulate an account of resultant luck motivated by these considerations (an "agent-driven account" of resultant luck)? There are two preliminary remarks to be made before we offer a tentative formulation. Firstly, an agent can be subject to resultant luck with respect to different outcomes, and in virtue of having behaved in different

⁷ Some hard-core Utilitarian views might reject this idea. But, again, our concern here is the ordinary concept of resultant luck. In the next section I discuss an example of a different kind that gives even stronger reason to embrace the agent-driven view and to reject the outcome-driven view.

ways. This suggests that we should take the focus of our analysis to be not a property of an agent but a three-place relation between an agent, an outcome, and a particular behavior by the agent. Secondly, there are two different ways in which an agent can be subject to resultant luck: by being morally responsible for an outcome and by failing to be morally responsible for an outcome. It will be useful, then, to break up the account into two parts: the first part will specify the conditions under which an agent that is responsible for an outcome is subject to resultant luck and the second part will specify the conditions under which an agent that is not responsible for an outcome is subject to resultant luck.

Now, what should each part say? In both cases, it seems natural to distinguish four different clauses. First, a “responsibility clause” stating that the agent is/is not morally responsible for the outcome in question. This is the clause in virtue of which the account would be an account of resultant *moral* luck, as opposed to some other kind of luck (say, epistemic luck). This would be followed by a “consequence clause” stating that the agent’s behavior resulted/didn’t result in the outcome. It’s in virtue of the consequence clause that the account would be an account of *resultant* moral luck, as opposed to some other variety of moral luck (say, constitutive moral luck). Then it seems that we would need a “link clause” to tie the responsibility and consequence clauses together. After all, a case of resultant luck is not just a case where an agent is or is not responsible for an outcome that he brought about or didn’t bring about, but a case where the agent is or is not responsible for an outcome, in some important sense, *because* he brought it about or didn’t bring it about. (For example, the key difference between the fortunate and unfortunate reckless drivers, in virtue of which one of them is morally lucky and the other one isn’t, is that one of them brings about a death and the other one doesn’t.) Finally, the last element of the account would be a “control clause” stating

that the agent didn't control the fact that his behavior resulted/didn't result in the outcome. This is the clause in virtue of which the account would be an account of resultant moral *luck*, as opposed to some other kind of moral phenomenon.

Now, what kind of link (between the responsibility and consequence clauses) should the link clauses postulate? The simplest way to go seems to be to understand the link clauses as positing an explanatory connection, i.e. as stating that the fact that the consequence clauses hold explains the fact that the responsibility clauses hold. I will go along with this suggestion for now, but let me note that, on the face of it, this can only be right if we understand "explains" as meaning: is at least part of the explanation of. For, when an agent is morally responsible for an outcome, the fact that his behavior resulted in the outcome is only part of the explanation of his responsibility; other facts—"epistemic" facts such as the fact that the agent could foresee that the outcome would ensue as a result of his behavior—also account for his responsibility.

So here is a first attempt at formulating an account of resultant luck ('A' stands for an agent, 'O' for an outcome, and 'B' for a behavior by A):

Resultant Luck (First Pass):

A is subject to resultant luck with respect to O in virtue of B iff:

Either:

- (I) (a) A is morally responsible for O
- (b) B resulted in O
- (c) (b) explains (a)
- (d) The fact that (b) holds is not fully in A's control

Or:

- (II) (a) A is not morally responsible for O
- (b) It is not the case that B resulted in O
- (c) (b) explains (a)
- (d) The fact that (b) holds is not fully in A's control

For the reasons explained above, clauses (Ib) and (IIb) are intended to apply to causal and non-causal results alike, clause (IIb) is meant to be consistent with O's occurring (I used "It is not the case that..." to try to flag this), and, finally, clauses (Ic) and (IIc) are meant to be consistent with other facts entering in the explanation of the agent's moral responsibility for O, or lack thereof.

Now, as we will see in due course, this account won't do because it is both too broad and too narrow. It is too broad in that it classifies cases where there is no resultant luck as scenarios of resultant luck. This is due to a problem with the link clauses of the account, which I'll take up in section V. And it is also too narrow in that it leaves out some genuine cases of resultant luck. This is due to a problem with the consequence clauses, which I'll take up in section VI. But, first, in the next section I will discuss scenarios of resultant luck of a different kind that the agent-driven account successfully captures but that an outcome-driven account overlooks. They are scenarios of resultant luck *by omission*. As we will see, scenarios of this kind are interesting in their own right, since the way in which they give rise to resultant luck is importantly different from the way in which standard scenarios give rise to it.

IV. Actual and counterfactual resultant luck

Consider the following scenario (from Fischer and Ravizza (1998)):

Sharks: While walking by the beach, Bad Samaritan sees a child drowning. He thinks he could easily jump into the water and save him but decides not to do so. The child dies. Unbeknownst to Bad Samaritan, the water is infested by sharks. Had he jumped in, the sharks would have attacked him and prevented him from saving the child.

Sharks is a scenario of omission: to the extent that Bad Samaritan is responsible for anything in this case, it's because of what he *doesn't* do, not because of what he does. In addition, *Sharks* seems to be a case of (good) moral luck: Bad Samaritan is lucky that the water was infested by sharks because the fact that it was infested by sharks seems to relieve him of at least some moral responsibility: although he is presumably responsible for not trying to save the child (given that, as far as he could tell, he could save him), he is not responsible for his failure to save him, or for the child's death, because he couldn't have saved him. In particular, *Sharks* seems to be a case of resultant moral luck. For it seems that Bad Samaritan is morally lucky because his behavior didn't contribute to the death's occurrence: the child's death wasn't one of its consequences (in contrast, in possible worlds where the water is shark-free, the child's death does seem to be one of the consequences of Bad Samaritan's behavior, as I will explain momentarily).⁸

⁸ Could one argue that *Sharks* is a case of circumstantial luck, not of resultant luck? (Circumstantial luck is another variety of moral luck identified by Nagel; see n.1.) After all, one might argue, Bad Samaritan is morally lucky given the circumstances he was placed in (because the sharks are present and thus he couldn't have saved

Is Sharks an instance of causal resultant luck or of non-causal resultant luck? This depends on what the right view on omissions and causation turns out to be. If, as some people think, omissions can be causes, then Sharks is arguably an instance of causal resultant luck. For, if omissions can be causes, then there is arguably a causal difference between Sharks and a shark-free scenario where Bad Samaritan also fails to intervene. Had there been no sharks in the water (or other obstacles to rescuing the child), then Bad Samaritan's failure to intervene would have been a cause of the child's death (for, if *any* omission is a cause, it seems that this one, in particular, is). By contrast, in the actual case in which the sharks are present, Bad Samaritan doesn't cause the child's death by failing to intervene. On the other hand, if, as other people think, omissions cannot be causes, then Sharks is an instance of non-causal resultant luck. For it still seems extremely plausible to say that, when there are no sharks, Bad Samaritan's inaction *results in* the child's death, or that it has it among its consequences, and that this isn't the case when the sharks are in the water.⁹

As far as I am aware, the literature on moral luck contains no discussion of the phenomenon of resultant luck by omission. But reflecting on this phenomenon can help us get a better grip on the concept of resultant luck. First of all, given the standing debate about the causal powers of omissions, the existence of resultant luck by omission makes it even more

the child). But note that this is also true of the reckless driver cases (and other paradigmatic cases of resultant luck): whether the pedestrian happens to be in the driver's path is similarly part of the circumstances, and it equally determines whether the agent is morally responsible. A case of circumstantial luck is not *just* a case where the circumstances matter to the agent's responsibility, but one where they matter in a particular way: they put the agent in a position where he must make certain decisions that he wouldn't otherwise have to make. For instance, if you are never in a situation where you have reason to think that you could save a drowning child, you won't ever have to make the decision whether to jump into the water to rescue a drowning child; *this* is circumstantial luck.

⁹ For example, one can say that Bad Samaritan's inaction results in the child's death in the shark-free scenario because in that scenario the child's death counterfactually depends on his behavior (since, had he jumped in, he would have succeeded in saving the child). By contrast, his inaction doesn't result in the child's death in Sharks because the child's death doesn't counterfactually depend on his behavior in that scenario. (See Dowe (2001) for a view of this kind. Dowe calls the non-causal relation that omissions can participate in, and in virtue of which they can have consequences or results, "quasi-causation.")

evident that we should be open to the possibility of non-causal resultant luck.¹⁰ Also, scenarios of resultant luck by omission like Sharks show ever so clearly that the outcome-driven account of resultant luck is misguided. For the outcome-driven account plainly fails to classify Sharks as a case of resultant luck. Bad Samaritan is morally lucky in this case, but not because the outcome could easily have been different. His moral luck is obviously not explained by the fact that the child could easily have survived, since this isn't a fact. Rather, and quite on the contrary, his moral luck is explained by the fact that the child wouldn't have survived even if he had tried to save him.

In addition, scenarios of resultant luck by omission are interesting because they have the following (related) feature. According to the agent-driven account, both the agent in Sharks and the agent in the fortunate reckless driver case are less blameworthy because their behavior didn't have harm as a consequence. But notice that what grounds *this* fact (the fact that the behavior doesn't issue in harm) seems to be importantly different in each case. We think that Bad Samaritan's behavior in Sharks doesn't issue in the child's death because he couldn't have saved the child.¹¹ But this isn't the reason we think that the fortunate reckless driver's behavior doesn't issue in harm: we don't think that his behavior doesn't issue in harm because there is a pedestrian whose life he couldn't have saved. Rather, we think that his behavior doesn't issue in harm simply because he didn't run over any pedestrians.

This is an important distinction. I suggest that we capture it by distinguishing two varieties of resultant luck, which (a bit misleadingly, and for lack of a better name) I'll call

¹⁰ Also, elsewhere I have argued that one can be morally responsible for an outcome by omission without causing it, even if omissions are causally efficacious (see Sartorio (2004)). If I am right, this is another potential source of non-causal resultant luck.

¹¹ This is very rough. I discuss this, and how it accounts for the agent's lack of moral responsibility, in Sartorio (2005a).

actual luck and *counterfactual luck*. Actual luck scenarios are scenarios where the agent's resultant luck is due to actual facts (what actually happens or doesn't happen), and counterfactual luck scenarios are scenarios where the agent's resultant luck is due to counterfactual facts (what happens or doesn't happen in close possible worlds, different from the actual world). The fortunate reckless driver case is a scenario of actual luck: what gets the reckless driver off the hook is the fact that no pedestrian gets in his path—an actual fact. In contrast, Sharks is a scenario of counterfactual luck: what gets Bad Samaritan off the hook is the fact that the sharks *would* have attacked him if he had jumped in—a counterfactual fact.¹²

There are many examples of resultant luck by omission like Sharks: if I couldn't have rescued the cat from the tree, I am typically not responsible for the nervous breakdown suffered by the cat's owner; if I couldn't have rescued the miner from the mine before it collapsed, I am typically not responsible for the miner's death from asphyxiation, etc.¹³ This suggests that “negative” responsibility (responsibility for something by omission) is importantly liable to counterfactual luck: the agent's moral responsibility significantly hinges on what *would* have happened if the circumstances had been different.

Now, “positive” responsibility (responsibility for something by commission) doesn't seem to hinge on the counterfactual facts in the same way, or to the same extent. Consider, for example, the following case:

¹² One reason this terminology is a bit misleading is that there is a sense in which in both cases there is moral luck because of some actual fact, in particular, an actual *causal* fact, or a fact about what the agent's behavior actually *results in*. However, I trust that the terminology still tracks an intuitive distinction: intuitively, what gets Bad Samaritan off the hook in Sharks is a fact about what the sharks would have done—not a fact about what the sharks actually *do*—whereas what gets the fortunate reckless driver off the hook is something that the pedestrian does (or doesn't do)—not something that he *would have done*.

¹³ See, notably, van Inwagen (1978).

Backup Assassin: A husband gets frustrated with the married life and decides to kill his wife. He designs a murder plan and kills her. Unbeknownst to him, the Mafia wanted his wife dead and they had sent an assassin to kill her. Had the husband not killed her himself, the assassin would have.

Here the husband's moral responsibility for his wife's death isn't determined by the counterfactual facts. The fact that the wife would still have died if the husband hadn't killed her is irrelevant. Intuitively, the husband is morally responsible for his wife's death regardless of what would have happened if he hadn't killed her. Again, one can think of a variety of cases of positive responsibility involving backup assassins, evil neuroscientists waiting in the wings, etc., and in all of these cases it seems that the counterfactual facts are irrelevant to the agent's moral responsibility for the outcome.¹⁴ This suggests that, at least as far as counterfactual luck goes, omissions are more subject to resultant luck than actions.¹⁵

Interestingly, however, there is a flip side to this phenomenon. There are, again, two opposed varieties of resultant luck: actual luck and counterfactual luck. And, although omissions are more subject to *counterfactual* luck than actions, they are less subject to *actual* luck than actions. To see this, consider the following variations on the killer husband scenario and the drowning child scenario:

¹⁴ See, notably, Frankfurt (1969). The cases offered by Frankfurt ("Frankfurt-style cases") are different from Backup Assassin in that they involve an early kind of causal redundancy (redundancy at the level of choices) instead of a late kind of redundancy (redundancy just at the level of outcomes). This difference isn't important for our purposes here.

¹⁵ Here I will bypass the question of whether this asymmetry between negative and positive responsibility generalizes to all actions and omissions (some people believe that it doesn't; I discuss this in my (2005a)). What seems clear is that it is at least much more common for omissions to be subject to counterfactual luck than it is for actions. This claim is compatible with the existence of some omission cases where the agent is *not* subject to counterfactual luck as well as with the existence of some action cases where the agent *is* subject to counterfactual luck. (In fact, I have argued for the existence of the latter type of case in Sartorio (2005b)).

Attempted Murder: The evil husband devices a plan to murder his wife. He poisons her food and then leaves the house. When the wife is about to eat the poisoned food, a fire breaks up unexpectedly and she runs out of the house. This time there are no backup assassins waiting in the wings. The wife never gets to try the poisoned food, which is consumed in the fire.

Rescued Child: Bad Samaritan notices the child drowning in the water and, again, decides not to help him. This time there are no sharks in the water, or any other obstacles to his rescuing the child. Right after he leaves the scene, another person walks by, sees the child, and rescues him.

It is clear that, in Attempted Murder (an action case), what actually happens matters a great deal to the husband's moral responsibility. Given that his plan is thwarted and the wife never dies, he is only responsible for trying to kill her, but not for killing her. Now, what is the bearing of the actual facts on Bad Samaritan's moral responsibility in Rescued Child (an omission case)?

The answer to this question is less obvious. At first sight, it might seem that the actual facts are relevant in this case too, and in precisely the same way. Given that the child doesn't die, Bad Samaritan cannot be held responsible for the child's death, just like the husband in Attempted Murder cannot be held responsible for his wife's death, since she doesn't die. Still, I think that there is an important sense in which the actual facts matter less in the omission case than in the action case. For there is something such that Bad Samaritan is responsible for

it in Rescued Child but the husband is not responsible for its counterpart in Attempted Murder. It's the fact that *he* didn't save the child, or the fact that the child wasn't saved by *him*.¹⁶ Even though the child survived in this case, Bad Samaritan wasn't the one who saved him. He had the opportunity to rescue him and no means of knowing that someone else would rescue him, and he still didn't do it. So he is to blame for not having rescued him. But notice: the positive counterpart in Attempted Murder of the fact that Bad Samaritan didn't save the child in Rescued Child would be the fact that the husband killed his wife. Now, the husband didn't kill his wife. So, clearly, he is not responsible for any such fact.

Let us examine this more closely. On the one hand, there is obviously something that both Bad Samaritan and the evil husband are responsible for: what they tried or didn't try to do (the evil husband is responsible for trying to kill his wife and Bad Samaritan is responsible for not trying to save the child). They are both also responsible for their evil or selfish intentions. On the other hand, there is something that neither Bad Samaritan nor the evil husband is responsible for: the occurrence of a particular death (the wife's or the child's). But what I am suggesting is that there is an important difference between the two cases in that, given that the wife didn't die, *all* the evil husband can be responsible for is his intentions or an unsuccessful attempt. In contrast, Bad Samaritan's blameworthiness exceeds the sphere of his will and of mere attempts. For, unlike the evil husband, he was successful in what he attempted (not) to do: he successfully omitted to save the child, and in circumstances where he should have saved him.

¹⁶ This fact is an outcome (or a way in which things can turn out) according to a liberal conception of outcomes (or of ways in which things can turn out) that I favor. Hence I take it to be something with respect to which Bad Samaritan can be subject to resultant luck.

What is at the root of this difference, I think, is a general fact about the nature of actions and omissions. Doing something entails that the thing is done. By contrast, failing to do something (at least something that has to do with the bringing about of an outcome in the external world) doesn't entail that the thing is not done; it only entails that it is not done *by the agent* (since it could be done by someone else). Boiling water entails that the water is boiled, but failing to boil water doesn't entail that the water is not boiled. Similarly, killing someone entails that the person dies, but failing to save a person doesn't entail that the person is not saved.¹⁷ So, whereas the evil husband didn't kill his wife, Bad Samaritan did (intentionally) omit to save the child.¹⁸

In virtue of this fact, Bad Samaritan ends up being less morally lucky than the evil husband. To be clear: I am not saying that he is less morally lucky in the sense that he is more blameworthy. If unsuccessful attempts to kill deserve more blame than successful attempts to not save (and it is plausible to think that this is the case), then the evil husband will still be more blameworthy than Bad Samaritan. Bad Samaritan is less morally lucky than the evil husband just in the sense that his responsibility is to a lesser extent *a matter of luck* than the evil husband's responsibility.

As we have seen, what reduces the evil husband's responsibility in the action scenario is an actual fact: the fact that his wife didn't die. The evil husband is more subject to actual luck given this fact than Bad Samaritan is given the (similarly actual) fact that the child didn't die. Hence we should conclude that, whereas (as noted before) negative responsibility is more

¹⁷ Weinryb discusses this asymmetry in his (1980). Intriguingly, Weinryb takes this to show that omissions don't have consequences and thus that we cannot be responsible for outcomes by omission. I fail to see why we should conclude that (which is, of course, consistent with someone else's being able to see why!).

¹⁸ What does it take for an agent to omit to A (intentionally)? This is a contended matter, but many philosophers of action would say that the following conditions are sufficient: intending not to A, not A-ing as a (non-deviant) result of intending not to A, and having the ability and the opportunity to A. All of these conditions are met in the Bad Samaritan case.

subject to counterfactual resultant luck than positive responsibility, positive responsibility is more subject to actual resultant luck than negative responsibility. There is a sense in which, for the agent to be morally responsible for how things turn out, the world has to cooperate more in the case of negative responsibility than in the case of positive responsibility, and another sense in which it has to cooperate less. It has to cooperate more to secure the relevant counterfactual facts, but it has to cooperate less to secure the relevant actual facts.

This concludes my discussion of resultant luck by inaction. Admittedly, a more thorough examination of this phenomenon (more than I can hope to have achieved here) would be desirable. But my main goal in this section has been to show that scenarios of resultant luck by inaction are an interesting focus of study since they are importantly different from traditional action scenarios and thus they give rise to some new issues and questions concerning moral luck. In the next section I return to the project of giving a more precise and accurate account of resultant luck.

V. Problem with the link clauses

Here is, again, the full statement of the agent-driven account from section III:

Resultant Luck (First Pass):

A is subject to resultant luck with respect to O in virtue of B iff:

Either:

- (I) (a) A is morally responsible for O
- (b) B resulted in O
- (c) (b) explains (a)

(d) The fact that (b) holds is not fully in A's control

Or:

(II) (a) A is not morally responsible for O

(b) It is not the case that B resulted in O

(c) (b) explains (a)

(d) The fact that (b) holds is not fully in A's control

In what follows I argue that this account misclassifies some cases. In this section I argue that it classifies some cases that aren't scenarios of resultant luck as scenarios of resultant luck; in the next section I argue that it also has the opposite flaw: it classifies some cases that are scenarios of resultant luck as scenarios where there is no resultant luck.

Imagine that I whistle a tune at home. A second later, the Queen of England falls and breaks her hip. Ordinarily, we wouldn't think that I am morally responsible for the Queen's breaking her hip. Why not? At least part of the explanation, it seems, is that my whistling the tune didn't result in her falling and breaking her hip: there is just no connection between my whistling the tune and her accident. So I am not responsible for the Queen's breaking her hip, and this is true (at least partly) in virtue of the fact that my whistling the tune didn't result in her breaking her hip.¹⁹ But, surely, I didn't control the fact that my whistling the tune didn't result in the Queen's breaking her hip. Someone could have set up a mechanism that would be triggered by my whistling and that would make her fall and break her hip. Given that I don't

¹⁹ Other facts besides the fact that my whistling didn't result in the Queen's breaking her hip also explain why I am not morally responsible (see below in the text). But this fact is at least part of the explanation, and that is enough for the relevant link clause to be met (see section III).

have control over this, I don't have control over whether or not my whistling the tune resulted in the Queen's breaking her hip.

So all the conditions in part II of the account are met, and thus the extant account of resultant luck entails that my not being morally responsible for the Queen's breaking her hip is a matter of luck. This generalizes. Consider all the outcomes for which I am not morally responsible and which I didn't help to bring about. For many such outcomes (presumably, all of those inside my light cone), we can imagine a process by which something I did or failed to do *would* have contributed to their occurrence, and in a way that was beyond my control. So the extant account of resultant luck entails, for all such outcomes, that my not being morally responsible for them is a matter of luck.

Am I subject to resultant luck in all these cases? On the face of it, this is very implausible. One could try to argue that the fact that the account has this consequence isn't enough to show that it is wrong. After all, part of the reason Nagel's original article was so influential is that it seemed to show that moral luck is everywhere. If moral luck in general is everywhere, perhaps resultant moral luck is everywhere too (and, in particular, in places where we hadn't expected to find it). However, I think that it would be a mistake to regard cases like the Queen scenario as cases of resultant luck. For, if we were to do that, we would be confusing resultant moral luck with luck of a non-moral kind.

Let me explain. I think that there might be an element of luck (however small or insignificant) in the Queen scenario. Namely, I am lucky, to some small extent, that my act of whistling a tune didn't result in the Queen's breaking her hip. Still, this doesn't seem to be an element of *moral* luck: I don't seem to be at all lucky for not being *morally responsible* for the Queen's breaking her hip. For, even in the unlikely event that my whistling the tune had

caused the Queen to break her hip, I still wouldn't have been morally responsible for her breaking her hip. This is so because no reasonable person in similar circumstances as mine could have predicted that whistling a tune would have such consequences. In other words, to the extent that there is luck in this case, it is not of a moral kind. Although it concerns my relationship to outcomes in the world, it doesn't concern my moral responsibility for them. Whatever exactly this kind of luck is, it isn't moral luck.²⁰

So, how could we avoid the implication that I am morally lucky in cases like the Queen scenario? The easiest thing to do seems to be to strengthen the link clauses.²¹ Again, the reason we don't think that I am morally lucky in the Queen scenario is that I still wouldn't be morally responsible in the unlikely event that my whistling the tune did result in the Queen's breaking her hip. For whatever epistemic conditions have to obtain for me to be morally responsible were presumably not met in this case, in particular, I couldn't possibly have foreseen that my whistling the tune would contribute to the Queen's breaking her hip. In other words, in this case the fact that my whistling the tune didn't bring about the outcome of the Queen's breaking her hip doesn't make the difference between my being morally responsible and my not being morally responsible.

Hence the idea would be to say that, in a genuine scenario of resultant luck where the agent is not morally responsible for an outcome, the fact that the agent's behavior didn't result in the outcome is what accounts for the agent's not being morally responsible, given that the epistemic conditions for moral responsibility are met. In other words, the fact that the

²⁰ Even though it isn't moral luck, this kind of luck could still warrant certain differences in our psychological responses (along the lines of Wolf's "nameless virtue" of caring about what results from our acts; see Wolf (2001)). For example, it could be that a feeling of regret towards the Queen's fall is more warranted if I caused it (and I find out about this) than if I didn't, even if I not morally responsible for her fall in either case.

²¹ Alternatively, one could add an extra clause (an "epistemic clause," stating that the epistemic conditions for moral responsibility are met).

consequence clause holds isn't just part of the explanation of the fact that the corresponding responsibility clause holds, as the link clauses of the original account say. Instead, the responsibility clause *rests on* the consequence clause, in the sense that, given that the epistemic conditions for moral responsibility are met, whether the agent's behavior brings about the outcome determines whether the agent is morally responsible for that outcome.²² If we modify the link clauses accordingly, this is the account that results:²³

Resultant Luck (Second Pass):

A is subject to resultant luck with respect to O in virtue of B iff:

Either:

- (I) (a) A is morally responsible for O
- (b) B resulted in O
- (c) (a) rests on (b)
- (d) The fact that (b) holds is not fully in A's control

Or:

- (II) (a) A is not morally responsible for O
- (b) It is not the case that B resulted in O
- (c) (a) rests on (b)
- (d) The fact that (b) holds is not fully in A's control

²² Note that "rests on" isn't synonymous with "is fully explained by." As I noted in section III, it would be false to say that claims about responsibility are fully explained by claims about what outcomes behaviors result in. Rather, the idea is to say that, in the circumstances (crucially, those circumstances include the fact that the epistemic conditions for responsibility are met) the responsibility claim depends on the consequence claim.

²³ For the sake of uniformity, I choose to modify both link clauses. This isn't really needed, since the problem to which the Queen scenario gives rise is only a problem for (IIc), not (Ic). But nothing important is lost by altering the formulation of both link clauses. At worst the account becomes redundant (since now both (Ia) and (Ic) entail that the epistemic conditions for moral responsibility are met).

For example, the fortunate reckless driver is subject to resultant luck because he is not responsible for any harm and his not being responsible for any harm rests on the fact that his behavior didn't result in any harm. This is so because the epistemic conditions that are required for moral responsibility were met, so the consequence clause makes the difference between his being responsible for some harm and his not being responsible for any harm. By contrast, I am not subject to resultant luck with respect to the outcome of the Queen's breaking her hip because my not being responsible for the harm doesn't rest on the fact that my behavior didn't result in the harm. This is so because the epistemic conditions for moral responsibility were not met, so the consequence clause doesn't make the difference between my being responsible for the harm and my not being responsible for it.

I have suggested that the original account should be revised by strengthening the link clauses. As a result, a difference between the revised account and the original account is that the link clauses play subtly different roles in them. As I explained at the time, the role of the link clauses in the original account was just to tie the responsibility and consequence clauses together. The link clauses play a more complex role in the revised account. Given that, as we have seen, it is only if we understand them in a specific way that the phenomenon that results is genuine moral luck, the link clauses also help render the account an account of *moral* luck (they do this in tandem with the responsibility clauses).

This wraps up my discussion of the link clauses. In the next and final section I discuss the consequence clauses.

VI. Problem with the consequence clauses

According to the extant account of resultant luck, the element that falls outside of the agent's control, in virtue of which there is resultant moral luck, is the existence of a consequence relation (causal or non-causal), which can obtain or fail to obtain between the agent's behavior (action or omission) and an outcome. The agent does his part (say, he decides to drive recklessly, or to not try to help the child, although he is aware of the consequences that such behavior might have), and then it all comes down to whether his behavior actually has the outcome as a consequence: if it does, he's morally responsible for the outcome; otherwise he's not (hence, his responsibility hinges on whether the behavior results in the outcome).

Now, I will suggest that this account is still deficient because it's incomplete: there is an element of a different kind that can also be out of our control and that can also give rise to resultant luck. We can see this by appeal to another evil-husband scenario, a variation on the ones introduced in section IV:

Choking with Water: Again, the evil husband wants his wife dead. He offers her a glass of water and gives her what he believes to be a glass of liquid poison. (He pours it from a bottle that he takes from a particular shelf in the cellar, where they usually keep that kind of poison.) But the bottle did in fact contain just water (the manufacturer simply filled the bottles with water and sold them as poison in order to make a larger profit). Unfortunately, however, the wife chokes on her own while drinking the water and dies immediately. There is nothing that the husband could have done at that point to save her life.

In Choking with Water, the evil husband does cause his wife's death, but we still wouldn't hold him morally responsible for her death (only for his evil intentions, or for trying to kill his wife). Why? The reason seems to be that the causal chain linking him to the death is "deviant": he ends up causing the harm, but not in the way he expected, and not in a way that any reasonable person could have foreseen.²⁴

Now, as usual, once deviance enters the picture, several new questions arise. First of all, what exactly is a "deviant" causal chain? How big of a departure from what the agent could foresee, or from what a reasonable agent could foresee, is required for a chain to be deviant? (And, what exactly is a "reasonable" agent?) But, more importantly, it is not even clear how deviance (roughly conceived) plays out in our ordinary moral assessments. As I have suggested, I think it is clear that in *some* cases of deviance (like Choking with Water) we don't hold agents morally responsible for the outcomes, due to the deviance of the causal chain, that is, due to some abnormality in the way the agent's behavior resulted in the outcome's occurrence. But it is unclear that we respond in the same way in all cases of deviance. Consider, for example, one last evil-husband scenario:

Slipping on Porch: The husband knows that his wife suffers from serious arachnophobia: if she were to see any spiders (especially any large specimens), she would be likely to have a nervous breakdown that would cause severe psychological

²⁴ Would we have the same reaction if the glass did in fact contain poison (but the wife still died from having choked with it, instead of from the poison itself)? I would still be inclined to say that the husband isn't responsible for the wife's death in this case, but I suspect that ordinary intuitions would be less clear than in Choking with Water. For discussion of scenarios where the agent is, presumably, relieved of responsibility for the outcome due to the deviance of the causal chain, see Feinberg (1970).

damage. The husband just came into the house, and saw a large spider on the porch.

So, hoping that his wife will see the spider on her way out, she persuades her to go run an errand. When she walks out of the house, the spider is already gone. However, as soon as she walks out the door, she slips on her own and falls, which causes some temporary pain and a few small bruises.

Here the causal chain to the harm is deviant too. But, for some reason, we are more tempted to hold the evil husband responsible for the harm in this case. We are tempted to think: she wouldn't have fallen had it not been for the fact that he persuaded her to get out of the house, and the only reason he did this is that he thought that she would thereby be harmed; so, we want to conclude, he is responsible for the fall and for the harm. What's puzzling is that the reasons we want to offer for holding him responsible in this case would also apply to Choking with Water, where we don't hold him responsible for the harm: in Choking with Water, the wife wouldn't have died had it not been for the fact that he gave her the glass with water, and the only reason he did this is that he thought that she would thereby be harmed. So, how could we try to account for the difference in our judgment about these two cases?

One possibility would be to say that we tend to hold the agent morally responsible for the harm in scenarios of deviance when the agent's fault is much greater than the actual harm. In Choking with Water, the fault and the harm are proportional (the husband tries to kill his wife and his wife actually dies); in Slipping on Porch the fault is greater than the harm (the husband tries to cause her to have a nervous breakdown, but she only ends up having some minor and temporary pain). Now, could this difference correspond to a genuine moral difference? Why would an agent be morally responsible for a harm he causes in a deviant way

only when the harm is minor (relatively to his fault)? Perhaps we shouldn't hold agents responsible even in those cases. Perhaps the only reason we feel tempted to hold agents responsible in those cases is that, in general, we want to hold people responsible for harms they caused when they are very much at fault, and, in cases where the harm is relatively minor, we are willing to overlook the fact that holding them responsible would be inappropriate or unjust, given that the injustice that we would be committing would also be minor.²⁵

At any rate, I am not going to try to settle this issue here. Our ordinary concept of responsibility seems to be particularly unclear, perhaps even inconsistent, in what concerns the role of deviance, and it is not my aim here to figure out the best way to resolve the obscurities and incongruities surrounding it. If one wanted to give an account of moral responsibility for outcomes, then one would have to look into the issue of deviance more carefully. My goal here, however, is not to give an account of responsibility but to identify different potential sources of resultant luck, i.e. ways in which lack of control over “how things turn out” can give rise to moral luck. Hence the conclusion that I would like to draw from the preceding discussion is only that there is an important source of resultant luck that is overlooked by the extant account of this concept. Even if it turned out that not all kinds of deviance give rise to resultant luck, I think it is safe to say that some do—in particular, I think

²⁵ Feinberg suggests something along these lines in Feinberg (1970), p. 434. An alternative explanation would be to say that we blame the agent in Slipping on Porch more than in Choking with Water because people more commonly slip on porches than they choke with water. Again, if this is what's behind our difference in judgment, it is hard to see how our reaction could be justified. For, although slipping on porches is more common than choking with water, the agent in Slipping on Porch had no particular reason to think that his wife would slip in that case. So, even if there is a difference of this kind, it doesn't seem robust enough to ground a moral difference. For an interesting empirical study about how people's ordinary attributions of moral responsibility respond to deviance, see Pizarro, Uhlmann, and Bloom (2003). (Their main conclusion is that deviance does have an important attenuating effect on people's attributions of moral responsibility, but that this attenuating effect is usually the result of “intuitive gut feelings of right and wrong,” not of rational deliberation.)

that scenarios like Choking with Water show quite conclusively that this is the case. If so, the extant account of resultant luck fails. For, as these scenarios show, an agent can be subject to resultant luck by failing to be morally responsible for an outcome due to factors beyond his control, although the reason he fails to be responsible is not that his behavior fails to have the outcome as a consequence. The kind of resultant luck illustrated by scenarios like Choking with Water is an under-recognized variety of resultant luck, which has its source, not in the fact that the agent's behavior results or fails to result in an outcome, but in the fact that the agent's behavior results or fails to result in an outcome in an atypical, unexpected, or otherwise abnormal way. This is also moral luck about "how things turned out," or resultant moral luck.²⁶

Following a customary practice in philosophy, and in order to explicitly flag the fact that it is not fully clear what the relevant sense of deviance is, I'll use the expression "in the right way" to refer to the way of resulting in an outcome that is normal or non-deviant in the relevant sense. This is, then, my proposed account of resultant luck:

Resultant Luck:

A is subject to resultant luck with respect to O in virtue of behavior B iff:

Either:

²⁶ The only work I know of that focuses on the phenomenon of deviance in connection with moral luck is Sverdlik (1988). Sverdlik argues against resultant luck in general and against deviance as a source of resultant luck in particular. He claims that, if an agent's blameworthiness for a harm he caused depended on the manner in which the harm came about (in particular, on whether it was way in which the agent intended it to happen), then any trivial departure from the intended manner would make the agent less blameworthy. And this, he thinks, is very implausible (for example, if a terrorist plants a bomb and the bomb goes off slightly later than planned, he thinks that we won't hold him any less responsible for the explosion). But, of course, to say that the manner in which the harm comes about matters is not necessarily to say that any detail matters. It might be, for example, that only considerable departures from the intended manner reduce the agent's responsibility. So I don't think that this is a good reason to rule out deviance as a source of resultant luck.

- (I) (a) A is morally responsible for O
- (b) B resulted in O in the right way
- (c) (a) rests on (b)
- (d) The fact that (b) holds is not fully in A's control

Or:

- (II) (a) A is not morally responsible for O
- (b) It is not the case that B resulted in O in the right way
- (c) (a) rests on (b)
- (d) The fact that (b) holds is not fully in A's control

The intended interpretation of II(b) is one according to which it can be satisfied in one of three ways: O fails to occur, O occurs but B still fails to result in O, and B results in O but still fails to do so in the right way. The agent can be subject to resultant luck in any of these scenarios.

To conclude: there are, in increasing order of generality, *three* potential sources of resultant luck for a given agent and an outcome. First, the outcome itself can occur or fail to occur. Second, and more generally, the agent can bring about the outcome or fail to bring it about. Third, and even more generally, the agent can bring about the outcome in the right way or fail to bring it about in the right way. Each of these elements can give rise to resultant luck because each of them may fail to be fully in the agent's control. First, agents may not fully control whether certain outcomes *occur*. Second, in cases where those outcomes do in fact occur, agents may still not fully control whether their behavior *results in* those outcomes.

And, finally, in cases where their behavior does in fact result in those outcomes, agents may still not fully control the *way* in which they do.

References

- Dowe, P. (2001) "A Counterfactual Theory of Prevention and 'Causation' by Omission," *Australasian Journal of Philosophy* 79, 2, 216-226.
- Feinberg, J. (1970) "Sua Culpa," in *Doing and Deserving: Essays in the Theory of Responsibility* (Princeton: Princeton U.P.), 187-221.
- Fischer, J. M. and M. Ravizza (1998) *Responsibility and Control* (Cambridge: Cambridge U.P.).
- Frankfurt, H. (1969) "Alternate Possibilities and Moral Responsibility," *Journal of Philosophy* 66: 829-39.
- Kim, J. (1974) "Noncausal Connections," *Noûs* 8, reprinted in *Supervenience and Mind* (New York: Cambridge U.P., 1993), 22-32.
- Lewis, D. (1986a) "Events," in *Philosophical Papers, Volume II* (New York: Oxford), 241-69.
- Lewis, D. (1986b) "Causation," in *Philosophical Papers, Volume II* (New York: Oxford), 159-213.
- Nagel, T. (1979) "Moral Luck," in *Mortal Questions* (New York: Cambridge U.P.), reprinted in Statman (1993), 57-71.
- Pizarro, D., E. Uhlmann, and P. Bloom (2003) "Causal Deviance and the Attribution of Moral Responsibility," *Journal of Experimental Social Psychology* 39: 653-60.
- Rescher, N. (1993) "Moral Luck," in Statman (1993), 142-166.

- Sartorio, C. (2004) "How to be Responsible for Something without Causing it," *Philosophical Perspectives* 18, 1: 315-36.
- Sartorio, C. (2005a) "A New Asymmetry between Actions and Omissions," *Noûs* 39, 3: 460-482.
- Sartorio, C. (2005b) "Causes as Difference-Makers," *Philosophical Studies* 123, 1-2: 71-96.
- Statman, D. (1993) (ed.) *Moral Luck* (New York: State University of New York Press).
- Sverdlik, S. (1988) "Crime and Moral Luck," *American Philosophical Quarterly* 25: 79-86.
- van Inwagen, P. (1978) "Ability and Responsibility," *Philosophical Review* 87: 201-24.
- Weinryb, E. (1980) "Omissions and Responsibility," *The Philosophical Quarterly* 30: 118: 1-18.
- Williams, B. (1981) *Moral Luck* (Cambridge: Cambridge U.P.).
- Wolf, S. (2001) "The Moral of Moral Luck," *Philosophic Exchange* 31: 4-19.
- Zimmerman, M. (1987) "Luck and Moral Responsibility," *Ethics* 97, reprinted in Statman (1993), 217-233.