

## Frankfurt Cases: The Fine-grained Response Revisited

Forthcoming in *Philosophical Studies*; please cite published version

### 1. Introduction

Consider the following familiar bit of science fiction.

*Assassin*: A nefarious neurosurgeon named Black wants Jones, a trained assassin, to decide to kill Smith. Black is willing to force Jones's hand if need be, but he would prefer that Jones make the decision to kill Smith on his own. So, he secretly implants a device in Jones's brain that enables him to control Jones's thoughts and behavior. The device is rigged to deterministically cause Jones to decide at time  $t$  to kill Smith, if, but only if, Jones does not decide on his own at  $t$  to kill Smith. There is, moreover, nothing Jones can do to prevent Black from carrying out this scheme. To Black's delight, Jones decides on his own at  $t$  to kill Smith, and so the coercive device never comes into play.<sup>1</sup>

This, of course, is an augmented version of a story originally sketched by Harry Frankfurt (1969, p. 835). Frankfurt and others contend that stories like this show that the principle of alternative possibilities (PAP, for short), a version of which states that a person is directly blameworthy for how he behaved at  $t$  only if it was within his power at or immediately prior to  $t$  to avoid behaving that way at  $t$ , is false. Their argument, in a nutshell, goes like this. Although Black and his mechanism are not among the causes of Jones's decision, their presence nevertheless renders

---

<sup>1</sup> For further details about how Black might accomplish all this, as well as a defense of the claim that scenarios like this are metaphysically possible, see Mele and Robb 1998 and 2003. There are numerous other "Frankfurt cases" in the literature. We focus here on those like *Assassin*, which were first developed by Mele and Robb 1998, because we think that they have the best chance of avoiding various difficulties often thought to plague other Frankfurt cases.

Jones powerless to avoid deciding at  $t$  to kill Smith. For either Jones decides on his own at  $t$  to kill Smith, or he decides as a result of Black's device; those, it seems, are his only options. Either way, though, he decides at  $t$  to kill Smith. However, because Jones decided on his own to kill Smith, without any "assistance" from the likes of Black and his coercive device, it seems that Jones could be directly blameworthy for his decision and subsequent actions, despite the fact that, through no fault of his own, it was not within his power to avoid deciding at  $t$  to kill Smith.

It is important to note that while Jones could not have avoided deciding at  $t$  to kill Smith, things did not have to go precisely the way they did either. Jones did not have to decide on his own at  $t$  to kill Smith; his decision to kill Smith could have instead been caused by Black's coercive device. As Frankfurt himself explains, "What action [Jones] performs is not up to him," though "it is in a way up to him whether he acts on his own or as a result of Black's intervention" (1969, p. 836). Several critics of Frankfurt's argument have seized on this point, insisting that it holds the key to showing where the argument goes wrong. Their claim is that what Jones is really blameworthy for is not deciding to kill Smith *per se*. What he is really blameworthy for, they contend, is deciding *on his own* to kill Smith, where "on his own" means, roughly, "not as a result of outside force or coercion." But, as we just noted, Jones could have avoided deciding on his own to kill Smith. So, if deciding on his own to kill Smith is what Jones is really blameworthy for, then, contrary to what Frankfurt and others claim, cases like *Assassin* do not provide us with scenarios in which someone is directly blameworthy what he did at  $t$  even though, through no fault of his own, the person could not have avoided doing it.<sup>2</sup>

We will refer to this response to Frankfurt's argument as the "fine-grained response" because it insists that a correct assessment of cases like *Assassin* requires being very precise

---

<sup>2</sup> Peter Van Inwagen (1978, p. 224, n. 24) was the first to suggest this sort of response. See also van Inwagen 1983, p. 181. Naylor 1984 subsequently developed the response in greater detail. More recent defenders of it include O'Connor 2000, Robinson 2012, and Speak 2002.

about what agents in those examples are blameworthy for. Although a number of philosophers, including some who are otherwise unsympathetic to Frankfurt's argument against PAP, have dismissed the fine-grained response, we believe there is a good deal to be said on its behalf. We will argue, in particular, that reflection on cases involving omissions undermines the main objections to the response and also provides the groundwork for an argument in support of it.

## 2. The Robustness Objection

We begin with an objection of John Martin Fischer's. The fine-grained response is a version of what Fischer (1994, p. 134) dubs "the flicker of freedom strategy," all versions of which rely in one way or another on the observation that there is a residual alternative possibility—a flicker of freedom—remaining in Frankfurt cases.<sup>3</sup> While Fischer acknowledges this residual alternative, he contends that it is irrelevant to the agent's responsibility. In his view,

it is not enough for the flicker theorist to analyze the relevant range of cases in such a way as to identify an alternative possibility. Although this is surely a first step, it is not

---

<sup>3</sup> According to Eleonore Stump, the flicker strategy "requires the supposition that doing an act-on-one's-own is itself an action of sorts," one that is distinct from the action the agent would have performed had the neuroscientist's device been among the causes of the agent's behavior. She then argues that this supposition is either "confused and leads to counterintuitive results; or, if the supposition is acceptable, then it is possible to use it to construct [Frankfurt cases] in which there is no flicker of freedom at all." (1999, pp. 301-302). This objection, however, runs together the fine-grained response with another version of the flicker strategy, which we might call the act-individuation version. According to the act-individuation version, Jones is indeed blameworthy for the decision to kill Smith that he made on his own. However, proponents of the act-individuation version of the flicker strategy contend that Jones could have avoided making that token decision, for while Black's device would have caused him to make a decision to kill Smith, that decision wouldn't have been identical to the one he made on his own, owing to its radically different causal history. Unlike proponents of the act-individuation approach, proponents of the fine-grained approach are not committed to saying that the decision Jones made on his own in the actual sequence of events is distinct from the one he makes in the counterfactual sequence of events in which his decision is caused by Black's coercive device. The fine-grained version of the flicker strategy therefore does not require the assumption that doing an act-on-one's-own is itself an action of sorts. To suggest that it does would be to conflate it with the act-individuation version of the flicker strategy. But once we clearly distinguish these two versions of the flicker strategy, we can see that Stump's criticism of the flicker strategy has no force against the fine-grained version of the strategy, as that version does not turn on what she regards as the implausible assumption that doing an act-on-one's-own is a distinct action. For further discussion of this issue, see [redacted for blind review].

enough to establish the flicker of freedom view, because what needs also to be shown is that these alternative possibilities *play a certain role* in the appropriate understanding of the cases. That is, it needs to be shown that these alternative possibilities *ground* our attributions of moral responsibility. (1994, p. 140)

According to Fischer, though, the sorts of alternatives identified by proponents of the flicker strategy are “not sufficiently *robust* to ground the relevant attributions of moral responsibility,” because they are not ones in which the agent *freely* does otherwise (1994, p. 140). He concludes that these alternatives are therefore irrelevant in and of themselves to moral responsibility.

We agree with much of what Fischer says here. We agree, in particular, that flicker theorists need to do more than simply identify an alternative possibility in the Frankfurt cases. They must also show that the alternative possibility helps ground moral responsibility. Where we take issue with Fischer’s position is his claim that the sorts of alternatives to which proponents of the fine-grained response advert are insufficiently robust to help ground responsibility. We contend that part of what makes Jones blameworthy for deciding on his own to kill Smith is that he could have avoided deciding on his own to kill Smith. We will argue that Fischer’s objection to this claim is unsuccessful and that there is good reason to think the claim is true.

A robust alternative is one that helps ground an agent’s moral responsibility; it is an alternative that is relevant *per se* to an explanation of why the agent is morally responsible for what he did.<sup>4</sup> Why think that the sorts of alternative possibilities to which proponents of the fine-grained response advert are not robust in this sense? According to Fischer, for an alternative to be robust, it must be one in which the agent acts freely or at least freely refrains from doing

---

<sup>4</sup> Our definition of a robust alternative is pretty much the standard one. For a slightly different use of the term “robust alternative,” see Mele 2006, p. 92.

something. But Fischer argues that the only alternative possibility available to the featured agent in cases like *Assassin* (viz., the alternative in which the agent's decision is caused by the coercive device) is not one in which the agent acts freely, nor is it one in which the agent freely refrains from acting on his own. Hence, Fischer concludes that that alternative is irrelevant per se to whether the agent is blameworthy for what happened in the actual sequence of events.

We wish to challenge Fischer's claim that, in the alternative sequence of events in which Jones's decision is caused by Black's device, Jones does not freely avoid deciding on his own to kill Smith. We contend that, in the alternative sequence, Jones does freely avoid deciding on his own. Attention to different ways an agent might freely avoid doing something supports our contention.

Sometimes when an agent freely omits or refrains from doing something, he first freely does something else in an effort to bring it about that he does not perform the action from which he wishes to refrain. Here is a case in point, a variant of which we will return to a bit later.

*Sloth*: John is walking along the beach when he sees a child struggling in the water. He believes that he could rescue the child with little effort, but not wanting to expend the energy it would take, decides not to even attempt to rescue the child. The child drowns.<sup>5</sup>

In this case, John first freely chooses not to rescue the child, and this choice results in his freely not rescuing her. Now, if this were the only way for an agent to freely omit or refrain from doing something, Fischer's claim that, in the alternative sequence of events, Jones does not freely avoid deciding on his own to kill Smith would be compelling, for, as Fischer (1994, p. 143) observes,

---

<sup>5</sup> This example is from Fischer and Ravizza 1998, p. 125.

in the alternative sequence, Jones does not first choose the possibility of not deciding on his own to kill Smith; he simply does not decide on his own.

It is, however, possible for an agent to freely omit doing something without first choosing the possibility of not doing it. The following case will help illustrate the point.

*Indecision*: Marla is deliberating about whether to attend a party this evening. Part of her wants to go; it will be a fun party, and she knows that she will have a good time. Another part of her, though, would prefer a quiet evening at home. At  $t$ , where  $t$  is some instant during the period of time that Marla is deliberating, Marla omits to decide to attend the party. To be clear, she does not decide at  $t$  not to attend the party, nor does she decide not to make a decision at  $t$ . She simply fails to decide at  $t$  one way or another.

Did Marla *freely* avoid deciding at  $t$  to go to the party? It depends, of course, but we see no reason why it could not be the case that she freely avoided deciding at  $t$  to attend. To drive the point home, let us add a few details to the case. Suppose that Marla's failure to decide at  $t$  did not result from coercion, manipulation, or any other freedom-subverting factor. Suppose, further, that Marla had it within her power to decide at  $t$  to attend the party and that it was up to her whether she decided at that moment to attend. Given these additional details, it seems that there are important respects in which Marla did freely avoid making a decision at  $t$  to attend the party, even though she did not (freely) choose at or prior to  $t$  not to make a decision at  $t$ .<sup>6</sup>

Cases like *Indecision* illustrate an important point. Not all omissions are the result of some prior choice or action; a person can omit to perform an action that he is considering performing without first choosing not to perform that action. There is a parallel here with

---

<sup>6</sup> Randolph Clarke (2014, pp. 96-97) makes similar observations about a different sort of case.

decisions. An agent typically need not do something to bring it about that he decides to *A*; he can just decide to *A*. Similarly, an agent often need not do anything to bring it about that he does not *A*; he can just not *A*.<sup>7</sup> This is especially true of omitting decisions. Typically when an agent omits or refrains from making a specific decision at *t*, the agent does not first decide not to make that decision at *t*. Instead, the agent either makes a different decision at *t*, or, as in cases like *Indecision*, makes no decision at all. And, as in *Indecision*, as long as the agent had it within his power at the time to *A* and his failure to *A* was not the result of any freedom-subverting factor, it appears that the agent freely avoids *A*-ing.

This fact casts significant doubt on Fischer's claim that, in the alternative sequence of events, Jones does not freely avoid deciding on his own to kill Smith. Jones had it within his power to decide on his own to kill Smith, and his failure in the alternative sequence to decide on his own was not a result of coercion, manipulation, or any other freedom-subverting factor.<sup>8</sup> But once we acknowledge all this, it seems we should also acknowledge that Jones freely avoids deciding on his own in the alternative sequence, even though he does not choose the possibility of not deciding on his own to kill Smith.

According to Fischer, in order for an alternative possibility to be robust, it must be such that the agent freely does something other than, or at least freely avoids doing, what he actually did. We have argued that, in cases like *Assassin*, the agent could have freely avoid deciding on his own to kill Smith, which, in turn, strongly suggests that that alternative possibility is sufficiently robust to ground the agent's responsibility for deciding on his own to kill Smith.

---

<sup>7</sup> Michael Robinson (2014, pp. 439-440) also makes this point.

<sup>8</sup> To be sure, Jones's *decision* in the alternative sequence was a product of manipulation. But this does not show that his failure to decide on his own was a result of manipulation. Recall that whether the device causes Jones's decision is contingent upon whether Jones decides on his own or not. Moreover, as even Frankfurt acknowledges, it was up to Jones whether he decided on his own to kill Smith. Jones's failure to decide on his own in the alternative sequence was thus not triggered by coercion but was itself a trigger of the coercion. Cf. [redacted for blind review].

### 3. The Moral Luck Objection

Central to the fine-grained response is the claim that Jones is not to blame for deciding to kill Smith owing to the fact that, through no fault of his own, he could not help deciding at  $t$  to kill Smith. Notice, though, that the features of the situation that prevent Jones from doing otherwise (viz., Black and his coercive mechanism) are not among the causes of Jones's actual decision to kill Smith. Proponents of the fine-grained response thus seem committed to the claim that features of an agent's environment that are irrelevant to an explanation of why the agent behaved as he did can nevertheless be relevant to whether the agent is blameworthy for his behavior. But some philosophers find this claim objectionable. Linda Zagzebski, for example, has suggested that this claim introduces a problematic sort of moral luck into the picture. She says:

It is only an accident that Black exists, and if he had not existed [Jones] would have had alternate possibilities. And if he had had alternate possibilities he would have done the very same thing in the same way. He is, therefore, just as responsible as he would have been if he had had alternate possibilities. To say otherwise is to permit [Jones] too great a degree of positive moral luck. He can't get off the moral hook *that* easily. (2000, p. 245)

The idea here seems to be this: to allow external circumstances such as the presence of Black and his device, the obtaining of which in no way affect what Jones does or why he does it, and which are, from Jones's perspective at least, just happenstance, to affect whether Jones is blameworthy

for deciding to kill Smith is to introduce an unacceptable form of moral luck. Surely, the thought goes, Jones cannot get off the hook for killing Smith because of such fortuitous circumstances.

There are two things to say in response Zagzebski's suggestion. First, it is important to keep in mind that proponents of the fine-grained response are not suggesting that Jones be exonerated entirely, only that he is not to blame for deciding to kill Smith. But this claim is consistent with there being other events and states of affairs in the story for which Jones can be held accountable. We can therefore agree that Jones is not off the hook in this case. Indeed, as we shall argue momentarily, proponents of the fine-grained response can even agree that Jones is "just as responsible" (i.e., worthy of just as much blame) as he would have been had it been within his power to do otherwise than decided to kill Smith. Second, while we concede that the fine-grained response introduces a certain kind of moral luck into the picture, the sort of luck it introduces, we shall argue, is unobjectionable. We develop both of these points in turn.

While proponents of the fine-grained strategy insist that Jones is not blameworthy for deciding to kill Smith, it is important to notice that, in saying this, we are not claiming that he is completely above reproach, nor are we proposing that he be completely exonerated. From the fact that Jones is not to blame for deciding to kill Smith, it does not follow that all negative moral assessments of him and his behavior would be inappropriate, nor does it follow that he is entirely blameless. Jones clearly is not the sort of guy you would want dating your daughter. He decided to kill someone in cold blood, despite knowing that doing so was seriously morally wrong. Disapprobation of him and his behavior therefore seems warranted. However, this leaves open the question of whether he is to blame for his decision and subsequent actions.

Disapprobation is not the same as blame. We can have a negative moral assessment of someone because of something the person did or failed to do while leaving it open whether he is

to blame for his immoral behavior. We can, for example, rightly judge a person to be insensitive, malevolent, or careless based on something the person did without deeming him blameworthy for the action in question or for the bad character traits manifested in the action. It simply might not be the person's fault that he is the way he is or that he behaved as badly as he did.

Even if Jones is not to blame for deciding to kill Smith, most everyone should be happy to acknowledge that there may be other states and events in *Assassin* for which Jones could be to blame. For instance, he may be to blame for seriously contemplating killing Smith in the first place, for being the sort of person who could decide to take the life of another person despite being aware of decisive moral reasons against doing so, and, if proponents of the fine-grained response are to be believed, for deciding on his own to kill Smith. By insisting that Jones is not blameworthy for deciding to kill Smith, proponents of the fine-grained strategy are therefore not committed to the further claim that all negative assessments of Jones and his behavior are unwarranted, nor are we committed to letting Jones off the hook. Indeed, a defining feature of our position is that Jones may very well be blameworthy, though not for deciding to kill Smith.

Having said this, however, we are happy to acknowledge that the fine-grained response does introduce a certain sort of moral luck into the picture. But as we shall now argue, the sort of luck it introduces is not especially uncommon, nor is it particularly objectionable.

Luck, it seems, can affect which events and states of affairs a person is blameworthy for without affecting the amount of blame of which the person is worthy. Fischer, who, you will recall, is no friend of the fine-grained response, illustrates the point using the following example.

*Broken Phone:* Smith witnesses a man being mugged outside his building. He knows he could easily dial 911, but, not wanting to be inconvenienced, decides to let sleeping dogs

lie. Unbeknownst to Smith, however, and through no fault of his own, his telephone was not working. So he could not have called the police even if he had tried.<sup>9</sup>

As Fischer rightly observes, although Smith is no doubt worthy of blame for *something* in this scenario—e.g., for deciding not to call the police, or for not trying to call them—he most certainly is not to blame for failing to call the police, i.e., for not successfully contacting them. Notice, though, that it is largely a matter of luck that Smith gets off the hook for failing to contact the police. After all, it was a fluke that the phone was not working properly, and if it had been working properly, Smith presumably would have been blameworthy for not contacting the authorities. Smith is thus extraordinarily lucky to escape blame for not calling the police. As Fischer goes on to point out, however, it does not follow that Smith is worthy of less blame than he would have been had he been blameworthy for not calling the police. Smith is worthy of just as much blame as he would have been had he been blameworthy for not contacting the police. It is just that he is blameworthy for fewer states of affairs than he otherwise would have been.

Examples like *Broken Phone* helpfully illustrate how luck can affect which events and states of affairs a person is blameworthy for without affecting how much blame the person is worthy of. Smith is very lucky to escape blame for not calling the police, and yet his good fortune in this matter does not make him worthy of less blame. Fischer puts the point this way: “whereas a certain kind of moral luck applies to the specification of the *content* of moral responsibility, it does not apply to the *extent* or *degree* of blameworthiness” (1986, p. 256).<sup>10</sup>

---

<sup>9</sup> See Fischer 1986, pp. 254-256. The example is originally due to van Inwagen 1983, pp. 165-166.

<sup>10</sup> Zimmerman 2002 draws a similar distinction between the degree of an individual’s responsibility and the scope of responsibility. It is worth pointing out that this distinction appears to provide a promising way to handle a wide range of cases involving moral luck. Consider, for example, the problem of distinguishing the culpability of a murder from the culpability of an attempted-but-luckily-unsuccessful murderer. It is intuitively plausible, some might say, that (since the difference between them is just a matter of luck) both agents deserve the same amount of blame. But it is also intuitively plausible that the murder is blameworthy for killing the victim, while the attempted

We contend that it is just the sort of luck Fischer highlights in cases like *Broken Phone* that is on display in Frankfurt cases like *Assassin*. As luck would have it, Jones is not to blame for deciding to kill Smith, though he may be to blame for deciding on his own to kill Smith. But, as we have seen, it does not automatically follow that Jones is worthy of less blame than he would have been had he also been blameworthy for deciding to kill Smith. All that follows is that Jones is blameworthy for fewer events and states of affairs than he otherwise would have been. Proponents of the fine-grained response can thus grant that Jones is worthy of *just as much blame* as he would have been in the absence of Black and his device, even though, as chance would have it, he is not to blame for the same states of affairs that he would have been blameworthy for in their absence. So, while defenders of the fine-grained strategy may disagree with Frankfurt and others about which events and states of affairs Jones is blameworthy for, we can all agree that he is not off the hook and, indeed, that he is worthy of just as much blame as he would have been had it been within his power to avoid deciding at *t* to kill Smith.

#### **4. The Artificial Separation Objection**

At the heart of the fine-grained response is the claim that Jones is blameworthy for deciding on his own to kill Smith but not for deciding to kill Smith. Some have worried that this claim slices things a bit too thin, and that it cannot plausibly be maintained that Jones is blameworthy for deciding on his own to kill Smith but not for deciding to kill Smith. According to Michael Otsuka, for example, this strategy “is controversial, since it is arguable that one needs to draw too fine a distinction in order to maintain that Jones is blameworthy for killing Smith on his own while at the same time denying that he is blameworthy for killing Smith” (1998, p. 690). In a

---

murder is not. Fischer’s distinction between the content and degree of blameworthiness accounts for this. The successful murder is blameworthy for an additional state of affairs, but he is not worthy of more blame.

similar vein, Robert Kane contends that the flicker strategy “artificially separates” moral responsibility for doing something on your own from moral responsibility for doing it. In general,” he says, “if we are responsible for doing something on our own, we are responsible for doing it.” And the same is true of Jones, he thinks. He insists that if Jones acted on his own, there is no reason to say that he is not morally responsible for his action (1996, p. 41).

The fine-grained response does indeed slice things pretty thin. There is no denying that. But what, exactly, is objectionable about this? *Jones decides on his own to kill Smith* and *Jones decides to kill Smith* are two related but nevertheless distinct states of affairs. Why should it not be feasible to suppose that Jones is morally responsible for the former state of affairs but not the latter? It may be true, as Kane claims, that, in general, a person who is morally responsible for *A-ing-on-his-own* is also responsible for *A-ing simpliciter*. But that is neither here nor there. From the fact that two things typically go together we cannot infer that they cannot be prized apart.<sup>11</sup>

To see this, consider the following case, which we discussed earlier.

*Sloth*: John is walking along the beach when he sees a child struggling in the water. He believes that he could rescue the child with little effort, but not wanting to expend the energy it would take, decides not to even attempt to rescue the child. The child drowns.

Ordinarily, an agent who is blameworthy for not trying to *A* and for deciding not to *A* is also blameworthy for not *A-ing*. *Sloth* is a case in point. John is blameworthy for not trying to save the child, for deciding not to save her, and also for not saving her. But we cannot generalize from relatively ordinary cases like this to all cases. Sometimes an agent is blameworthy for not trying to *A*, and for deciding not to *A*, and yet is not blameworthy for failing to *A*, as illustrated by cases

---

<sup>11</sup> Robinson (2012, p. 184) makes a similar point.

like *Broken Phone*. Recall that, in that case, Smith is not to blame for failing to contact the police, even though he is to blame for not trying to call them and for deciding not to call them.

The point can also be illustrated by reflection on the following variation on *Sloth*.

*Sharks*: John is walking along the beach when he sees a child struggling in the water. He believes that he could rescue the child with little effort, but not wanting to expend the energy it would take, decides not to even attempt to rescue the child. The child drowns. Unbeknownst to John, there is a school of sharks hidden beneath the water. Had John tried to rescue the child, the sharks would have eaten him, thwarting his rescue effort.<sup>12</sup>

Here again John is to blame for not trying to save the child and for deciding not to save her. But unlike in *Sloth*, it seems clear that in this case he is not blameworthy for failing to save the child.

Notice, moreover, that another judgment seems plausible: while John is not blameworthy in *Sharks* for failing to save the child, he might be blameworthy for failing *on his own* to save the child. He could have acted in such a way that the sharks (an outside force) would have prevented him from rescuing the child. Furthermore, we may suppose that John correctly believed that his decision not to enter the water would result in his failing on his own to save the child. Given all this, it seems plausible that John is to blame for failing on his own to save the child. However, we cannot infer from this fact that John is also blameworthy for failing to rescue the child.

Ordinarily, if an agent is blameworthy for failing to try to *A*, and if he is blameworthy for failing on his own to *A*, he is also blameworthy for failing to *A*. But, as cases like *Broken Phone* and *Sharks* illustrate, we cannot generalize from ordinary cases to all the relevant cases. An agent can be blameworthy for not trying to *A* and for failing on his own to *A* without being

---

<sup>12</sup> This example is from Fischer and Ravizza 1998, p. 125.

blameworthy for failing to *A*. Similarly, it may be true that, in general, a person who is morally responsible for *A*-ing-on-his-own is also responsible for *A*-ing *simpliciter*. But by itself, apart from further argumentation, that fact does not support Kane's conclusion that if an agent is morally responsible for *A*-ing-on-his-own, then the agent is also morally responsible for *A*-ing.

According to the fine-grained response, Jones is blameworthy for deciding on his own to kill Smith but not for deciding to kill Smith. It is instructive to note that ordinary moral judgments are often no less fine-grained, no less precise about what, exactly, a person is morally responsible for. "It's not what you said, it's how you said it" is a familiar accusation. So is "he did the right thing but for the wrong reason." But if a person can be criticized for performing an action for the wrong reason but not for the action per se, we are again left to wonder why it would be implausible to say that an agent is to blame for *A*-ing-on-his-own but not for *A*-ing?

Perhaps the worry is that the distinction between blameworthiness for *A*-ing and blameworthiness for *A*-ing-on-one's-own is somehow ad hoc or unmotivated. This possibility is suggested by Kane's remark that the fine-grained response "artificially separates" responsibility for performing an action from responsibility for performing the action on one's own. But if that is the worry, it is unfounded, for there is a principled reason to suppose that blameworthiness for *A*-ing and blameworthiness for *A*-ing-on-your-own come apart in Frankfurt cases like *Assassin*.

A plausible explanation of why blameworthiness for not trying to *A* and blameworthiness for not *A*-ing come apart in cases like *Broken Phone* and *Sharks* is that while the featured agent in those examples had at least some control over whether he tried to *A*, through no fault of his own, the agent had no control over whether he *A*-ed. To see this, consider *Broken Phone* again. Normally, having it within your power to *try* to call the police and having it within your power to successfully call them go together, which, in turn, helps explain why, normally, someone who is

blameworthy for failing to try to call the police is also to blame for not calling them. But in a case like *Broken Phone*, while the agent had some control over whether he tries to call the police, through no fault of his own, he had no control over whether he calls them, which plausibly explains why he can be blameworthy for failing to try to call the police, despite the fact that he is not to blame for failing to call them. Similar remarks apply, *mutatis mutandis*, to *Sharks*.

We contend that a similar explanation is available for why blameworthiness for *A-ing-on-one's-own* and blameworthiness for *A-ing* come apart in cases like *Assassin*, even though they typically go together. Through no fault of his own, Jones had no control over whether he decided to kill Smith, though he apparently did have at least some control over whether he decided on his own to kill Smith. Put somewhat differently, Jones had no say concerning whether the state of affairs *Jones decides to kill Smith* obtains, though he did have some say over whether the state of affairs *Jones decides on his own to kill Smith* obtains. According to proponents of the fine-grained response, it is this fact that grounds the difference in blameworthiness for the two states of affairs. The difference is therefore not artificial or unmotivated as Kane's remarks suggest; it is grounded in the difference in control Jones had over the two states of affairs in question. Because Jones, through no fault of his own, had no control over whether the state of affairs *Jones decides to kill Smith* obtains, he is not blameworthy for it. However, since he evidently did have some control over whether the state of affairs *Jones decides on his own to kill Smith* obtains, there is no obvious barrier to saying that he is blameworthy for this second state of affairs.

In ordinary cases in which there is no evil neuroscientist like Black waiting in the wings to make sure that things happen in a particular way, we tend to assume that the agent had control over whether he *A-ed* and thus over whether he *A-ed* on his own. This plausibly explains why blameworthiness for *A-ing* and blameworthiness for *A-ing-on-our-own* typically go together. But

there is arguably a morally significant difference between such ordinary cases and Frankfurt cases like *Assassin*. In the ordinary cases, the agent presumably had it within his power to avoid *A-ing* and thus had it within his power to avoid *A-ing-on-his-own*. So in those cases we have no reason to suppose that the agent is morally responsible for *A-ing-on-his-own* but not for *A-ing*. But things are importantly different in cases like *Assassin*. In those cases, the agent, through no fault of his own, could not help *A-ing*, which is why he is not to blame for *A-ing*. However, the agent apparently could have helped *A-ing-on-his-own*. Hence, the reason for thinking that he is not to blame for *A-ing* cannot be extended to show that he is not to blame for *A-ing-on-his-own*. Proponents of the fine-grained response therefore have a principled explanation of why blameworthiness for *A-ing* and blameworthiness for *A-ing-on-our-own* tend to go together, even though they come apart in some cases of the sort Frankfurt and his supporters adduce.

Proponents of the fine-grained strategy thus have good reason to resist the claim that if Jones decided on his own then there is no reason to deny that he is blameworthy for his decision. No doubt such claims gain purchase from the fact that the phrase “Jones decided on his own to kill Smith” can be interpreted in such a way that it is equivalent to “Jones decided of his own free will to kill Smith.” And, of course, if Jones had decided of his own free will to kill Smith, then, barring other exculpating considerations, he would have been blameworthy for his immoral decision. But the temptation to interpret the phrase in this way should be resisted. According to the fine-grained response, when we say that Jones decided on his own to kill Smith, we are making an observation about the etiology of Jones’s decision, viz., that it was not the result of outside coercion or force by the likes of Black and his device. By itself, though, this observation does not imply anything about whether Jones acted freely or about whether he is morally responsible for his action. A person can act “on his own” in the sense relevant to the fine-grained

response and yet not act freely. If you accidentally kill your neighbor by spooning a deadly poison into her tea that you blamelessly mistook for sugar, you will have killed your neighbor on your own. However, it would clearly be a mistake to say that you killed her of your own free will.<sup>13</sup>

## 5. The Case for the Fine Grained View

So far we have argued that reflection on certain cases in which an agent omits to perform some action casts doubt on each of the main objections to the fine-grained response. We believe that reflection on some of those same cases also suggests a reason to favor the response.

The scope of an agent's moral responsibility in cases like *Broken Phone* and *Sharks* is sensitive to what the agent can and cannot do. In those cases, the agent has it within his power to try to *A*, but, through no fault of his own, he cannot *A*. Consequently, it seems that the agent is not blameworthy for the fact that he fails to *A*. The agent may, however, be blameworthy for related facts, such as the fact that he does not even try to *A* or the fact that he decides not to *A*, since the agent did have at least some control over whether those related facts obtain.

We find it plausible that there should be symmetry in how lacking the power to *A* impacts the scope of responsibility in cases like *Broken Phone* and *Sharks* and how lacking the power to *avoid A-ing* impacts the scope of responsibility in Frankfurt cases like *Assassin*. If lacking control over whether certain facts obtain restricts the scope of the agent's responsibility in *Broken Phone* and *Sharks* in the way that it intuitively seems to, it is only natural to suppose that

---

<sup>13</sup> One might worry that the fine-grained approach will make moral responsibility too fine-grained. For example, will it turn out that agents are responsible for states of affairs such as their *A-ing*-given that the entire history of the universe can truly be described as follows: *e1* occurred, *e2* occurred, etc.? Our answer is that an agent can only be responsible for facts when the agent satisfies certain epistemic requirements. Since normal agents do not have the requisite epistemic access to the complete history of the universe, they will not be responsible for extremely fine-grained facts about that history. Notice, though, that Jones was presumably aware that he decided on his own to kill Smith. He did not know about Black's device, to be sure. So he was not aware that he had a choice about deciding as a result of Black's device. But he was presumably aware at the time that he was deciding on his own (i.e., not as a result of coercion or force). Hence, he can be responsible for making the decision to kill Smith on his own.

a similar lack of control would restrict the scope of the agent's responsibility in cases like *Assassin* in roughly the way that proponents of the fine-grained response suggest it does. We thus find it plausible to suppose that the agent in *Assassin* is not blameworthy for *A*-ing, given that, through no fault of his own, he could not help *A*-ing at *t*. The agent may, however, be blameworthy for related facts, such as the fact that he *A*-ed on his own, since the agent evidently did have at least some control over whether that related fact obtains.

Those who deny this symmetry between the two sorts of cases will have to say that, in a case like *Sharks*, lacking control over whether certain facts obtain has an impact on which facts the agent is blameworthy for, but that, in a case like *Assassin*, a similar lack of control has no impact on which facts the agent is blameworthy for. But this would be a rather puzzling asymmetry. Why should lack of control over whether certain facts obtain affect which facts the agent is blameworthy for in cases like *Sharks* but not in Frankfurt cases like *Assassin*?

This question becomes even more acute once we recognize that the distinctive feature of Frankfurt cases does not provide a basis for the alleged asymmetry. The main thing that distinguishes cases like *Assassin* from other scenarios in which an agent cannot avoid *A*-ing is that the circumstances that render the agent powerless to avoid *A*-ing in *Assassin*—viz., the presence of Black and his coercive mechanism—are not among the causes of the agent's *A*-ing. But this feature of the Frankfurt cases does not help to ground an asymmetry between them and omissions cases like *Sharks*, since the circumstances that render the agent powerless to *A* in a case like *Sharks* are not among the causes of the agent's not *A*-ing. Just as Black's coercive mechanism plays no role in the production of Jones's decision in *Assassin*, so too the sharks play

no role in producing John's omission in *Sharks*. Their role, like that of Black's mechanism, is merely to rule out alternatives. In this respect, then, the cases appear to be symmetrical.<sup>14</sup>

There is a related point to be made here. The presence of the sharks does not appear to be among the causes of John's omission in *Sharks*, and yet, their presence does seem to have an impact on whether John is blameworthy for that omission. As we noted earlier, while John is blameworthy in *Sharks* for not trying to save the child and for deciding not to save her, he is not to blame for omitting to save her. But had the sharks been absent (as in *Sloth*, for example), John presumably would have been blameworthy, not only for not trying to save the child and for deciding not to save her, but also for not saving her. The fact that the presence of the sharks restricts the scope of John's responsibility in this way, even though their presence plays no role in producing his omission, provides further support for the fine-grained response. For once we grant that the presence of the sharks has an impact on the scope of John's responsibility in *Sharks*, even though their presence is not among the causes of his omission, it seems much more plausible to suppose that the presence of Black and his device impacts the scope of Jones's responsibility in *Assassin* in roughly the way that the fine-grained response suggests it does, the fact that Black and his device are not among the causes of Jones's decision notwithstanding.

The fine-grained response thus enables us to deal with action cases like *Assassin* in the very same way that we deal with omission cases like *Sharks*. Moreover, as far as we can tell, it allows us to do so without undue costs. (As we have seen, proponents of the fine-grained response can accommodate the intuition that an agent in Frankfurt cases like *Assassin* is just as blameworthy as he would have been in the absence of Black and his device, and we can do so in

---

<sup>14</sup> Cf. [redacted for blind review]

a principled, non ad-hoc fashion, without having to abandon plausible principles like PAP.) And this, we contend, gives us at least some reason to endorse the fine-grained response.<sup>15</sup>

## References

Author 1, Article

Author 2, Article

Clarke, R. 2014. *Omissions: Agency, Metaphysics, and Responsibility*. New York: Oxford University Press.

Fischer, J.M. 1994. *The Metaphysics of Free Will: An Essay on Control*. Oxford: Blackwell.

Fischer, J.M. 1986. "Responsibility and Failure," *Proceedings of the Aristotelian Society* 86: 251-270.

Fischer, J. M. and Ravizza, M. 1998. *Responsibility and Control: A Theory of Moral Responsibility* (New York: Cambridge University Press).

Frankfurt, H. 1969. "Alternate Possibilities and Moral Responsibility," *Journal of Philosophy* 66: 829-839.

Kane, R. 1996. *The Significance of Free Will*. New York: Oxford University Press.

Mele, A. 2006. *Free Will and Luck*. New York: Oxford University Press.

Mele, A. and Robb, D. 1998. "Rescuing Frankfurt-style Cases," *Philosophical Review* 107: 97-112.

Mele, A. and Robb, D. 2003. "Bbs, Magnets and Seesaws: The Metaphysics of Frankfurt-style Cases," in David Widerker & Michael McKenna (eds.), *Moral Responsibility and Alternative Possibilities: Essays on the Importance of Alternative Possibilities*, Ashgate, 107-126

---

<sup>15</sup> Acknowledgements.

- Naylor, M. 1984. "Frankfurt on the Principle of Alternate Possibilities," *Philosophical Studies* 46: 249–258.
- O'Connor, T. 2000. *Persons and Causes: The Metaphysics of Free Will*. New York: Oxford University Press.
- Otsuka, M. 1998. "Incompatibilism and the Avoidability of Blame," *Ethics* 108: 685-701.
- Robinson, M. 2012. "Modified Frankfurt-type Examples and the Flickers of Freedom," *Philosophical Studies* 157: 177-194.
- Robinson, M. 2014. "The Limits of Limited Blockage Frankfurt-style Cases," *Philosophical Studies* 169: 429-446.
- Speak, D. 2002. "Fanning the Flickers of Freedom," *American Philosophical Quarterly* 39: 91-105.
- Stump, E. 1999. "Alternative Possibilities and Moral Responsibility: The Flicker of Freedom," *The Journal of Ethics* 3: 299-324.
- van Inwagen, P. 1978. "Ability and Responsibility," *The Philosophical Review* 87: 201-224.
- van Inwagen, P. 1983. *An Essay on Free Will*. New York: Oxford University Press.
- Zagzebski, L. 2000. "Does Libertarian Freedom Require Alternative Possibilities?" *Philosophical Perspectives* 14:231-248.
- Zimmerman, M. 2002. "Taking Luck Seriously," *Journal of Philosophy* 99: 553-576.