



The Power of Care: Reply to Sliwa

Daniel Morgan, Philosophy, University of York, UK,
danielmorgan1917@gmail.com

Nina Simone sang: “I’m just a soul whose intentions are good.” Paulina Sliwa (2019) defends a subtle, worked out picture on which Simone’s excuse turns out to be the basic form of an excuse. But good intention is a straitjacket that most excuses resist. There are cases where someone has an excuse, and the intention they have is just *neutral*. There are cases where someone has an excuse, and the intention they have is actually *bad*. In both sorts of case, the agent’s excuse must consist in something other than a good intention. I argue that it consists in their caring as they ought. The basic form of an excuse is “She cared; her heart was in the right place.”



The Power of Care: Reply to Sliwa

Daniel Morgan

To excuse someone is to concede that there was something wrong about what they did; excusing is different from justifying. Still, if they have an excuse, then that is some kind of point in their favour. What kind of point? On Sliwa's picture of excuses, the kind of excuse that Nina Simone sang about having – "I'm just a soul whose intentions are good" – is actually the basic form of an excuse. I offer counterexamples to this picture, and show how the counterexamples can be accommodated on an alternative picture that appeals to the notion of care. Most excuses don't fit the straitjacket Sliwa's account makes for them. But they are captured by an account that appeals to care.

§I sets out Sliwa's account. §II sets out a counterexample involving an excused agent whose only relevant intention is a *neutral* intention. §III sets out a counterexample involving an excused agent whose only relevant intention is a *bad* intention. §IV sketches an account of excuses on which the basic form of an excuse is: "She cared; her heart was in the right place."

I. THE GOOD INTENTION ACCOUNT

Sliwa's summary of her account, which she calls "The Good Intention Account," is as follows:

Here is the account, in a nutshell. When someone commits a wrong, we may generally infer that they lacked a morally adequate present-directed intention. Excuses block that inference. Excuses are considerations that show that the agent's wrongdoing does not reflect her lack of a morally adequate present-directed intention.¹

Sliwa follows Bratman in taking intentions to be "mental states characterized by their distinct functional role in action and planning."² Intentions are *stable*: once formed, they tend to persist. Intentions are *controlling*: unless revised, intentions tend to lead to action. Sliwa further assumes that intentions are *coherent*: one cannot have "two simultaneous conflicting present-tense intentions and act on one of them without giving up or revising the other."³

¹ Paulina Sliwa, "The Power of Excuses," *Philosophy & Public Affairs* 47, no. 1 (2019): 45, <https://doi.org/10.1111/papa.12139>.

² *Ibid.*, 46.

³ *Ibid.*, 46.

Supposing one thinks, as Sliwa does, that having an excuse is always a matter of what one's motives are, one option is a *lacking a bad motive* view. Another option is a *having a good motive* view. Sliwa highlights this distinction and stresses that she is going for the second option. She writes: "I suggest that excuses point to the presence of a morally adequate motive, rather than the absence of a reprehensible one. Second, I suggest that it's a particular type of motive that matters for excuses: a morally adequate present-directed intention."⁴

Sliwa fleshes out "present-directed" as requiring that the intention "concern what to do *right now*."⁵ Present-directed intentions thus contrast with intentions targeted at specific points in the future, like the intention to wear black *tomorrow*, or general resolutions, like the resolution to *only ever* wear black.⁶ In support of the focus on present-directed intention, Sliwa says: "When it comes to excuses, it is present-directed intentions that matter. When confronted with wrongdoing, we care about the wrongdoer's motivation at the moment of action—not her general resolutions and plans."⁷

Sliwa makes the colloquial notion of "good" intention more precise by reference to the idea of *moral adequacy*. Regarding intention *adequacy* in general, she writes: "To assess an intention for adequacy is to assess it as a plan with respect to a given goal. How adequate a plan is with respect to that goal depends on how conducive it is to realizing it – for example, how conducive my intention to go to Japan is to getting me there."⁸ So, adequacy goes with conduciveness, and the primary barrier to conduciveness Sliwa highlights is *lack of specificity*. For example, the intention to go to Japan is highly non-specific, leaving open as it does all details of one's route.

Sliwa says that a *morally adequate* intention is conducive to some relevant *moral* goal and that it's for the ethicists to say what moral goals there are.⁹ Merely having the intention "to do the right thing" is not morally adequate, even assuming that doing the right thing is one moral goal, since it is so non-specific. It is analogous to the intention to go to Japan. Sliwa thinks that to have a moral-ignorance-related excuse one needs something more like the intention to go to Japan *by taking a train to Heathrow*.¹⁰ This seems very plausible, and very separable from Sliwa's intention-centric framework. But the key takeaway for the moment is that "morally adequate" doesn't mean morally "not bad." Moral adequacy of intention requires *positive conduciveness* towards a *relevant moral goal*.

⁴ Ibid., 45.

⁵ Ibid., 46.

⁶ Ibid., 46.

⁷ Ibid., 47.

⁸ Ibid., 51.

⁹ Ibid., 52.

¹⁰ Ibid., 52.

II. NEUTRAL INTENTION CASES

I think you can be excused, and not have a morally adequate present-directed intention (a “good intention”) that excuses you, because your only present-directed intention is morally *neutral* rather than good. That is, I think, often the case with what Sliwa calls “I didn’t do it on purpose” excuses. Sliwa sets out her account’s story about such cases in this passage:

Consider a common way of making an excuse: “I’m so sorry, I didn’t mean to offend,” or “I didn’t intend to give away a secret – I thought this was common knowledge,” or “I was just trying to...”. They take the form of an appeal: yes, I did act wrongly – I broke the promise, gave away the secret, caused the offense – but unintentionally so.

The Good Intention Account explains why one’s wrongdoing being unintentional can constitute an excuse. To say that it was unintentional is just to say that it did not reflect the lack of a morally adequate intention. Rather, something went wrong in putting my morally adequate intention into action: I slipped, tripped, lost my balance, or I was ignorant about some crucial feature of the situation.

Consider Jones who stirs a heaping spoonful of poison into Smith’s tea. Since he has every reason to believe that the white powder is sugar, the poisoning is unintentional. In light of this, the wrongful deed does not reflect a lack of morally adequate present-directed intention. Jones is guided by the present-directed intention to sweeten Smith’s tea; there is nothing untoward about that. The poisoning reflects Jones’ ignorance about the content of the sugar bowl.¹¹

I’ll discuss the Smith/Jones case, and I’ll focus on two candidates for what morally adequate present-directed intention might be excusing Jones. The first candidate, which Sliwa considers, is the intention to *be sweetening Smith’s tea*. The second candidate, which she does not consider, is the intention to *not be killing Smith*.

I’ll argue that the problem with the tea-sweetening intention is that Jones’s intending this benefit for Smith is incidental to his excuse. We can easily think of a relevantly similar case in which the goal the agent is working towards is morally neutral, in so far as it doesn’t involve a benefit for anyone other the agent themselves, yet the agent is still excused. I’ll argue that the problem with the intention to not be killing Smith is that there is no reason to think Jones has this intention. That’s because not killing Smith is not the *motive* or *point* of anything Jones does.

Any putative “good intention” that Sliwa’s theory might appeal to in the Smith/Jones case will have one or the other of these two problems. Which problem a

¹¹ Ibid., 49.

candidate intention has will depend on whether the intention is directed at *producing a small benefit* or *avoiding a great harm*. There is room for variation in how we specify the benefit and the harm. For example, we might focus on “benefiting Smith,” generically, rather than “sweetening Smith’s tea.” Or we might focus on “not poisoning Smith” rather than “not killing Smith.” But the basic problem, I think, is that it’s irrelevant to Jones’s excuse that he intends to be producing the small benefit, and there is no reason to think that he intends to not be producing the great harm. The most illuminating description of Jones’s present-directed intentions as he heaps in the poison is that there is “nothing untoward” about them (to use Sliwa’s own revealing phrase), not that they are *good*.

II.A An Intention to Be Sweetening Smith’s Tea

Sliwa’s candidate for a good intention that Jones has is the intention to be *sweetening Smith’s tea*. This is positively conducive towards what might be regarded as a moral goal, albeit a rather trivial one, something like *giving Smith some pleasure* or *giving Smith a thing he wants* or *slightly benefiting Smith*.

Consider though another very similar case. Two actors perform a poisoning scene. The first actor assumes they are stirring a harmless, non-sweetening, powder into the other actor’s cup. This assumption of harmlessness is reasonable since the powder has been handed to them by the officer in charge of props. Tragically, the powder is in fact poison and the first actor unintentionally kills the second. The actor, I stipulate, does not intend to *sweeten the tea* of their fellow actor, nor to confer any other *benefit* on them, or on anyone else. Just like Jones, they have an excuse. But the goal that is the source of that excuse can’t involve any kind of *benefit*. No story about Jones’s excuse that targets a *benefit* Jones intends to produce for Smith is going to extend to the actor as well. The actor’s positive intention – to act out a scene – is morally neutral, not mildly altruistic. But the right story about Jones will extend to the actor, because the actor is every bit as excused as Jones is, and their excuse is of the very same kind.

II.B An Intention to Not Be Killing Smith

Perhaps Sliwa focused on the wrong intention, by the lights of her own theory. There is an alternative candidate. Jones is being excused for the wrong of killing Smith. This makes the goal of *not killing Smith* uniquely relevant. Perhaps the present-directed intention that excuses Jones is the intention to *not be killing Smith right now*.

The problem is that we lack any good reason, independent of Sliwa’s theory of excuses, to think Jones has formed this intention. If we ask, “Why is Jones doing what he is doing?”, Sliwa’s own answer – “He intends to be sweetening Smith’s tea” – sounds right. The answer “He intends to not be killing Smith” does not sound right.

What explains this difference, I think, is that not killing Smith is not the *point* or *motive* of anything Jones does, whereas sweetening his tea is.

Notice that the point here is not about what conscious deliberations Jones engaged in, or failed to. The following claim of Sliwa's seems highly plausible: "future-directed intentions and resolutions will often automatically give rise to present-directed intentions."¹² Here, "automatically" means "without conscious deliberation." Let's assume that Jones has the standing commitment to never kill anyone. Let's agree that this standing commitment could give rise to a more specific present-directed intention – an intention to not be killing Smith right now – without Jones performing any conscious deliberation. The question is whether we have a reason to think this kind of transition, from standing commitment to present-directed intention, has occurred in the case at hand. My answer is that, because not killing Smith isn't the *motive* of anything Jones does, we lack any reason to think that it has.

Contrast with a different case in which conscious deliberation is also absent. I am *driving to get to the shops* and I pass a school. Without any conscious deliberation, in particular without conscious thinking about the possibility of killing any children, I *slow down* to 15mph. Is there reason in this case to think that, as I drove past the school, I had a present-directed intention to not kill any of the children who study there? There is. It's supplied by the fact that, although not killing any children may not be the point or motive of my *driving to the shops*, it is the point or motive of my *reducing my speed*.

Jones's case is not like this. Not killing Smith is not only not the ultimate point of the stirring in of the powder, but, additionally, in no way shapes *how* the powder is stirred in. There is no aspect of how Jones is stirring in the powder the motive or point of which is to not kill Smith. We could make Jones's case like the driving past the school case. Suppose Jones is a superhero who could in principle kill Smith by stirring in the powder at mega-speeds causing a potentially fatal tea-tsunami. An intention to not kill Smith might lead him to stir at a more human-friendly speed. But, as naturally imagined, Jones is not like that.

A standing commitment *can* give rise to a present-directed intention without any conscious deliberation. But that doesn't mean that, for every standing commitment an agent has, and for everything they do, we should think that they have the present-directed intention to not right now be violating that standing commitment by doing that thing. So, abstracting from the specifics of the case, there is no reason to think Jones intended to not be killing Smith by stirring powder into his tea.

Focusing on the specifics of the case, there is still no reason to think Jones intended to not be killing Smith, since there's nothing in how Jones does what he does that is

¹² Ibid., 47.

shaped by the goal of not killing Smith. That Jones *in fact* killed Smith does explain why a not-killing intention has a salience to us that intentions regarding other harms – e.g., bringing Smith’s 20 year of sobriety to an end by unknowingly loading his tea with cocaine – does not. But there’s no more reason to think that Jones had the intention to not be killing Smith, than there is to think that he had the intention to not be getting Smith high.

Jones aimed to benefit Smith a little, and instead harmed him a lot. He has some kind of excuse. If the excuse consists in a present-directed intention, the intention will either be one directed at producing the small benefit, or at not producing the great harm. But the fact that Jones intends to produce a small benefit is irrelevant to his excuse. His intention might as well be neutral (and we can easily think of a similar case, the actor case, in which the excused agent’s intention is neutral). And there is no reason, independent of Sliwa’s theory, to think that Jones has formed the intention to not be producing the great harm.

III. BAD INTENTION CASES

In this section, I discuss a different kind of problem for Sliwa’s account. It involves a kind of case where someone is excused (partially) despite the fact that the only relevant intention they have is a *bad* intention. As in the previous section, I draw on one of Sliwa’s own examples: “The agent who gives in to police interrogation and falsely testifies against an innocent person acts wrongly under duress.”¹³ Intuitively, police duress can be some kind of excuse for giving false and damaging testimony. Since the only kind of excuse the Good Intention Account can allow for involves good intention, Sliwa has to treat the case as one in which the testifier retains a relevant good intention – e.g., to not testify against the innocent person – even as the duress leads them to act against that intention, and on something that is not an intention. Sliwa calls this a “momentary desire” (2019:55). She fleshes out the case by imagining the momentary desire the false testifier’s good intention is circumvented by is a desire to “just make the interrogation stop” (2019: 55).

No doubt Sliwa’s testifier has an excuse based on their good intention. But does good intention *have* to be on the scene in a case of false testifying where there is an excuse? The candidate excused testifier I am interested in never gets as far as forming the good intention to not testify against the innocent person. A fortiori this intention is not present-but-circumvented while they give the false testimony. They testify against the innocent person. But they are, even as they testify, highly conflicted: they testify

¹³ Ibid., 54.

despite strongly desiring not to, and they strongly desire not to because they care about the bad consequences that testifying against someone may have for that person. Why then do they testify? Because they care about something else even more. This other thing is something that it's morally entirely permissible for them to care about, but not permissible to prioritise over the extremely weighty moral goal of not testifying against an innocent.¹⁴ For example, they care even more about not being beaten up. This leads them to decide to testify falsely, which they then do.

It's obvious that there's a *morally relevant difference* between the action on this occasion of this person, and the action of the person who testifies against an innocent because it just doesn't bother them in the least to do so. The action of the first exhibits something that falls in the range of normal human weakness (where exactly it falls will depend on the details); the second exhibits chilling callousness. But if there is a *morally relevant difference* between the two agents' actions, and neither agent is *justified* in doing what they are doing, then there's nothing for the difference to be other than the one having an *excuse* (a partial one, intuitively), and the other not.

Excuses with the "strong but outweighed morally positive motivation" structure that this person's excuse exemplifies are plausibly *common*. For example, they seem common in addiction, because addictive desires can be so strong. Suppose an addict steals from a family member and to do so they engage in protracted and planned activity that exhibits some finesse and even ingenuity (e.g., trying out various birthday-related password combinations on their family member's computer, sending the right amount of money to their dealer's account and remembering to add the extra fee for urgent deliveries). It seems implausible, in a case like this, that they could be acting *against* present-directed intentions: e.g., that, even as they try the fourth birthday combination, they have a sustained *intention* "to not be stealing" that is circumvented by something that is less than an intention.¹⁵ But their action nevertheless seems *morally better* than that of an agent who steals from a family member *without it bothering them in the least*. The addict's extreme desire to use is plausibly overpowering a weaker, but still strong, desire to not steal, and the fact that they have this strong desire to not steal that needs to be overpowered is morally in their favour, compared to the other agent. Granted that their action is not *justified*, what kind of point in their favour could this be if not an *excuse*?

¹⁴ That the false testifier is *not* justified distinguishes this case from the putative counterexample to her view Sliwa considers at *ibid.*, 61. The difference preempts the reply she gives to that case: the agent isn't excused but they are justified.

¹⁵ The discussion at *ibid.*, 55–59 suggests that Sliwa would emphatically agree.

IV. CARING AS ONE OUGHT

There is an account that can deal with problem cases I have been discussing. Here it is:

CARE: An agent has an excuse when, and to the extent that, their wrong-doing does not reflect a failure of care.

Caring goes with being motivated *generically* as opposed to having one *specific* kind of motivational state (e.g., intention). If something is motivating to one, then one cares about it. Sometimes one has a corresponding intention. Sometimes one doesn't.

I don't take a position on some tricky questions one might have about care. To care might just *be* to be motivated. Alternatively, caring might be a state that is distinct from and *explains* (by grounding, or causing) being motivated. Caring might be exhausted by its motivational aspect, or it might have aspects additional to motivation. These are interesting questions, but not ones I think that need to be answered to see that CARE does a better job with Jones, and my false testifier (the one who acts on, not against, intention), than the Good Intentions Account does.

Caring can be, and mostly is, *dispositional*. It's not that, at the moment he stirs in the white powder, there is *something it's like* for Jones to care about not killing other people, to care about not killing Smith, to care about not poisoning Smith, and to care about not *now* poisoning Smith by moving his right arm thusly. He cares about at least some of these things in virtue of the fact that if, as he put the powder in the tea, someone grabbed his arm and said the powder was poison, he would be shocked, he would unhesitatingly stop doing what he is doing. These dispositional tests constitute Jones as someone who *cares* about not poisoning Smith (but, notice, they don't constitute him as someone who has formed an *intention* to not be poisoning Smith).

Failure of care has two subtypes. One is not caring about something one ought to care about (at all or as much as one ought). The other is caring about something one ought not to (at all or more than one ought).¹⁶ It's the first kind of failure that is relevant to Jones and to my false testifier. An example of the second kind of failure would be someone using a racially offensive term in a second language, having been told that it is the respectful term for the relevant group of people. Their use of the term is excused because it fails to reflect their being positively motivated by the goal of giving racial, or any, offence.

¹⁶ This mirrors what Arpaly and Schroeder, in their account of blameworthiness, call the distinction between indifference and ill-will; see Nomy Arpaly and Timothy Schroeder, *In Praise of Desire* (Oxford: Oxford University Press, 2013), <https://doi.org/10.1093/acprof:oso/9780199348169.001.0001>.

CARE says that Jones is excused since his action does not reflect a failure to care about killing Smith. In fact, there's no level of care about not killing that Jones's killing Smith reflects his failure to attain, precisely because he's unaware that what he is doing is killing Smith. Hence, he seems like he might have a *complete* excuse, depending on how the details of the case are filled in.

CARE says that my false testifier is excused if, and to the extent that, their false testimony fails to reflect a failure to care about not testifying against the innocent person. But there is some level of care about not testifying against an innocent person that this person does not have. If they had *that* level of care they wouldn't have testified falsely. CARE therefore predicts that their excuse is a *partial* excuse, which it intuitively is.

Why does CARE tie having an excuse to whether a wrong-doing *reflects* a failure of care as opposed, more simply, to whether the wrong-doer fails to care? Consider Jones* who would love to kill, but who kills Smith in complete ignorance of the fact that he is doing that, unaware that the powder is poison. A bad fellow no doubt, but not on account of what he has done on this occasion: he seems to have the same excuse in killing Smith as Jones has. That is because neither agent's killing Smith *reflects* a failure of care, even though one of them does fail to care. When considering excuses, we usually don't just care about which psychological properties someone *has*, at the time they act; we care about which properties are *reflected in* (or "expressed in" or "manifested in" or "in the appropriate way, causing" or "explaining") their action.¹⁷

The primary goal of this paper has been to provide an assessment of Sliwa's interesting project of tying excuses to intentions, which to my knowledge it has yet to receive. CARE is relevant to that assessment because, even if there were cases – e.g., Jones, my false testifier – that the Good Intention Account finds hard to deal with, that would be far less conclusive if those cases were hard for any account to deal with. That isn't how things are. CARE seems to say exactly the right thing about these cases. I don't claim that CARE is a particularly novel view. I think that the basic idea behind CARE is a natural starting point for a story about excuses that we would need to be argued away from. With a view to further clarifying CARE, I close by highlighting two points of contrast between CARE and an account in the literature that CARE has much in common with, Arpaly and Schroeder's.

¹⁷ The word "reflect" appears in the formulation of Sliwa's account I began with, on p. 45, although she does not say what work she wants it to do. She also sometimes drops "reflects"-talk in favour of "presence"-talk, e.g., "Whether someone has an excuse is a matter of their moral psychology: the presence or absence of morally adequate intentions."; "Excuses point to the presence of a morally adequate motive." Sliwa, "Excuses," 59, 45.

The first contrast is that CARE is non-committal on something Arpaly and Schroeder want to be committed about. Is the key thing – for having an excuse (my framing), or for blame being mitigated (Arpaly and Schroeder’s) – to be motivated by right-making features of actions, or by the fact that those actions are right? Arpaly and Schroeder commit to the first option. But there’s nothing in CARE that requires that, since one can care about the fact that something is right, just as one can care about features that are, in fact, right-making features. So, there’s no reason why someone who defends CARE has to say that moral ignorance is never an excuse. They can even borrow Sliwa’s nice story about *when* moral ignorance is an excuse. The story would be in terms of how specific the thing the agent cares about is (just caring about “doing the right thing” is never enough). What sort of *attitude* (e.g., intending vs. caring) is central in an account of excuses is one question. What sort of *content* (e.g., *de dicto*, *de re*) is central is another, on the face of it, orthogonal question. CARE is only addressing the first question.

Second, a natural way of framing the contrast between Arpaly and Schroeder and Sliwa is that the former praise (i.e., make central to theorising about some morally important phenomenon) desire and the latter praises intention. But there’s no sense in which CARE praises desire at the expense of intention. Desire and intention are on a par, as far as CARE is concerned, since they are both closely connected with caring, which is what CARE says is morally relevant. Furthermore, CARE is sensitive to *degree of care* (“to the extent that”) and in so far as *outright intending* is more correlated with caring a lot than *merely desiring*, CARE can even allow that knowing that someone got as far as intending, and didn’t merely desire, something morally positive they ultimately failed to do (e.g., to not testify against an innocent person) can be evidence relevant to the strength of their excuse.

V. CONCLUSION

Given the features intentions have (stability, controllingness etc), it makes sense that intentions play the distinctive role in *planning* that they do. But is there any reason to expect that the “planning role,” and the “excusing role,” should be played by one and the same part of the psyche? Aspects of mind other than how things are with one’s planning states seem like they ought to be potentially relevant to one’s moral status, to one’s credit or discredit, in a case where one does wrong. On my picture, this is because one may be excused by what one cares about.

Acknowledgements

This paper has been much improved by points made by Keith Allen, Lucy Campbell, Hannah Carnegie-Arbuthnott, Greg Currie, Stephen Everson, Alexander Greenberg, Alex Gregory, Owen Hulatt, David Ingram, Jess Isserow, Mary Leng, Chris Jay, Paul Noordhof, Matthew Ratcliffe, Louise Richardson, Léa Salje, Ashley Shaw, Rob Trueman, Pekka Väyrynen, Gerald Lang, Christian Piller and two anonymous editors for this journal. Special thanks are due to Owen Hulatt for commenting on a number of drafts.

Competing Interests

The author has no competing interests to declare.

